



---

# Complete Software Guide for Junos<sup>®</sup> OS for EX4300 Switches, Release 14.1X53

Release

14.1X53



---

Published: 2015-01-21

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

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

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

*Complete Software Guide for Junos® OS for EX4300 Switches, Release 14.1X53*  
Release 14.1X53  
Copyright © 2015, Juniper Networks, Inc.  
All rights reserved.

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

#### YEAR 2000 NOTICE

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

#### END USER LICENSE AGREEMENT

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



# Table of Contents

	About the Documentation . . . . .	cix
	Documentation and Release Notes . . . . .	cix
	Supported Platforms . . . . .	cix
	Using the Examples in This Manual . . . . .	cix
	Merging a Full Example . . . . .	cx
	Merging a Snippet . . . . .	cx
	Documentation Conventions . . . . .	cxi
	Documentation Feedback . . . . .	cxiii
	Requesting Technical Support . . . . .	cxiii
	Self-Help Online Tools and Resources . . . . .	cxiii
	Opening a Case with JTAC . . . . .	cxiv
<b>Part 1</b>	<b>Getting Started with ELS</b>	
<b>Chapter 1</b>	<b>Getting Started with ELS . . . . .</b>	<b>3</b>
	Getting Started . . . . .	3
	Getting Started with Enhanced Layer 2 Software . . . . .	3
	Understanding Enhanced Layer 2 Software Support . . . . .	3
	Using the ELS Translator Tool . . . . .	4
	Configuring a VLAN . . . . .	5
	Configuring the Native VLAN Identifier . . . . .	6
	Configuring Layer 2 Interfaces . . . . .	6
	Configuring Layer 3 Interfaces . . . . .	6
	Configuring an IRB Interface . . . . .	7
	Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface . . . . .	7
	Enhanced Layer 2 CLI Configuration Statement and Command Changes . . . . .	8
	Enhanced Layer 2 CLI Configuration Statement and Command Changes . . .	17
<b>Part 2</b>	<b>System Setup</b>	
<b>Chapter 2</b>	<b>Overview . . . . .</b>	<b>29</b>
	Software Overview . . . . .	29
	Understanding Software Infrastructure and Processes . . . . .	29
	Routing Engine and Packet Forwarding Engine . . . . .	29
	Junos OS Processes . . . . .	30

<b>Chapter 3</b>	<b>Configuration</b> . . . . .	<b>33</b>
	Initial Configuration . . . . .	33
	Configuring the LCD Panel on EX Series Switches (CLI Procedure) . . . . .	33
	Disabling or Enabling Menus and Menu Options on the LCD Panel . . . . .	33
	Configuring a Custom Display Message . . . . .	34
	Configuring Date and Time for the EX Series Switch (J-Web Procedure) . . . . .	35
	Configuring System Identity for an EX Series Switch (J-Web Procedure) . . . . .	36
	Configuration Statements . . . . .	37
	arp (System) . . . . .	39
	authentication-key . . . . .	40
	auxiliary . . . . .	41
	boot-server (NTP) . . . . .	42
	broadcast . . . . .	43
	broadcast-client . . . . .	44
	console (System Ports) . . . . .	45
	default-address-selection . . . . .	46
	domain-name . . . . .	47
	gre-path-mtu-discovery . . . . .	47
	host-name . . . . .	48
	icmpv4-rate-limit . . . . .	48
	icmpv6-rate-limit . . . . .	49
	inet6-backup-router . . . . .	50
	internet-options . . . . .	51
	ipip-path-mtu-discovery . . . . .	52
	ipv6-duplicate-addr-detection-transmits . . . . .	52
	ipv6-path-mtu-discovery . . . . .	53
	ipv6-path-mtu-discovery-timeout . . . . .	53
	ipv6-reject-zero-hop-limit . . . . .	54
	lcd-menu . . . . .	55
	location (System) . . . . .	57
	menu-item . . . . .	59
	multicast-client . . . . .	62
	no-multicast-echo . . . . .	63
	no-ping-record-route . . . . .	64
	no-ping-time-stamp . . . . .	64
	no-redirects (IPv4 Traffic) . . . . .	65
	no-tcp-rfc1323-paws . . . . .	66
	no-tcp-rfc1323 . . . . .	66
	ntp . . . . .	67
	path-mtu-discovery . . . . .	68
	peer (NTP) . . . . .	69
	port-type . . . . .	70
	ports . . . . .	71
	power . . . . .	72
	processes . . . . .	73
	server (NTP) . . . . .	74
	tcp-drop-synfin-set . . . . .	75
	traceoptions (SBC Configuration Process) . . . . .	76
	trusted-key . . . . .	78

<b>Chapter 4</b>	<b>Administration . . . . .</b>	<b>79</b>
	Operational Commands . . . . .	79
	clear chassis display message . . . . .	81
	clear system reboot . . . . .	84
	configure . . . . .	88
	op . . . . .	90
	request chassis pic . . . . .	92
	request chassis routing-engine master . . . . .	96
	request system halt . . . . .	101
	request system logout . . . . .	107
	request system power-off . . . . .	108
	request system reboot . . . . .	113
	request system reboot . . . . .	118
	request system scripts convert . . . . .	120
	request system scripts refresh-from commit . . . . .	122
	request system scripts refresh-from event . . . . .	123
	request system scripts refresh-from op . . . . .	124
	request system storage cleanup . . . . .	125
	restart . . . . .	135
	set chassis display message . . . . .	146
	set date . . . . .	149
	show chassis fan . . . . .	150
	show chassis firmware . . . . .	163
	show chassis lcd . . . . .	174
	show configuration . . . . .	187
	show host . . . . .	190
	show ntp associations . . . . .	191
	show ntp status . . . . .	193
	show system firmware . . . . .	196
	show system reboot . . . . .	198
	show system software . . . . .	201
	show system storage . . . . .	209
	show system switchover . . . . .	217
	show system uptime . . . . .	221
	show system users . . . . .	226
	show system virtual-memory . . . . .	231
	show task replication . . . . .	289
	show version . . . . .	291
	show version fpc . . . . .	304
<b>Chapter 5</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>305</b>
	Troubleshooting Loss of the Root Password . . . . .	305

<b>Part 3</b>	<b>User Interfaces</b>	
<b>Chapter 6</b>	<b>Overview</b>	<b>311</b>
	Software Overview	311
	Understanding Software Infrastructure and Processes	311
	Routing Engine and Packet Forwarding Engine	311
	Junos OS Processes	312
	User Interfaces	313
	CLI User Interface Overview	313
	CLI Overview	313
	CLI Help and Command Completion	313
	CLI Command Modes	313
	EX Series Switches Hardware and CLI Terminology Mapping	315
<b>Chapter 7</b>	<b>Configuration</b>	<b>317</b>
	Configuration Tasks	317
	Using the J-Web CLI Terminal	317
	Configuring the Web Browser	317
	Setting Domain Name, Hostname, and Name Server	318
	Enabling SSH on your system	318
	Sample Configuration on an EX Series Switch	318
	Statement Hierarchies	319
	[edit access] Configuration Statement Hierarchy on EX Series Switches	322
	Supported Statements in the [edit access] Hierarchy Level	322
	Unsupported Statements in the [edit access] Hierarchy Level	324
	[edit accounting-options] Configuration Statement Hierarchy on EX Series Switches	325
	Supported Statements in the [edit accounting-options] Hierarchy Level	326
	Unsupported Statements in the [edit accounting-options] Hierarchy Level	326
	[edit chassis] Configuration Statement Hierarchy on EX Series Switches	327
	Supported Statements in the [edit chassis] Hierarchy Level	327
	[edit class-of-service] Configuration Statement Hierarchy on EX Series Switches	328
	Supported Statements in the [edit class-of-service] Hierarchy Level	329
	Unsupported Statements in the [edit class-of-service] Hierarchy Level	330
	[edit ethernet-switching-options] Configuration Statement Hierarchy on EX Series Switches	331
	Supported Statements in the [edit ethernet-switching-options] Hierarchy Level	331
	Unsupported Statements in the [edit ethernet-switching-options] Hierarchy Level	334
	[edit event-options] Configuration Statement Hierarchy on EX Series Switches	334
	Supported Statements in the [edit event-options] Hierarchy Level	334
	Unsupported Statements in the [edit event-options] Hierarchy Level	336

[edit firewall] Configuration Statement Hierarchy on EX Series Switches . . .	336
Supported Statements in the [edit firewall] Hierarchy Level . . . . .	337
Unsupported Statements in the [edit firewall] Hierarchy Level . . . . .	337
[edit forwarding-options] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	338
Supported Statements in the [edit forwarding-options] Hierarchy	
Level . . . . .	338
Unsupported Statements in the [edit forwarding-options] Hierarchy	
Level . . . . .	340
[edit interfaces] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	342
[edit interfaces ae] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	343
Supported Statements in the [edit interfaces ae] Hierarchy Level . . .	343
Unsupported Statements in the [edit interfaces ae] Hierarchy Level . .	346
[edit interfaces ge] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	346
Supported Statements in the [edit interfaces ge] Hierarchy Level . . .	347
Unsupported Statements in the [edit interfaces ge] Hierarchy Level . .	350
[edit interfaces gr] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	350
Supported Statements in the [edit interfaces gr] Hierarchy Level . . .	350
Unsupported Statements in the [edit interfaces gr] Hierarchy Level . .	353
[edit interfaces interface-range] Configuration Statement Hierarchy on EX	
Series Switches . . . . .	353
Supported Statements in the [edit interfaces interface-range] Hierarchy	
Level . . . . .	354
Unsupported Statements in the [edit interfaces interface-range]	
Hierarchy Level . . . . .	357
[edit interfaces lo] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	360
Supported Statements in the [edit interfaces lo] Hierarchy Level . . . .	361
Unsupported Statements in the [edit interfaces lo] Hierarchy Level . .	363
[edit interfaces me] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	363
Supported Statements in the [edit interfaces me] Hierarchy Level . . .	364
Unsupported Statements in the [edit interfaces me] Hierarchy	
Level . . . . .	366
[edit interfaces vlan] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	367
Supported Statements in the [edit interfaces vlan] Hierarchy Level . . .	367
Unsupported Statements in the [edit interfaces vlan] Hierarchy	
Level . . . . .	370
[edit interfaces vme] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	370
Supported Statements in the [edit interfaces vme] Hierarchy Level . .	370
Unsupported Statements in the [edit interfaces vme] Hierarchy	
Level . . . . .	373

[edit interfaces xe] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	373
Supported Statements in the [edit interfaces xe] Hierarchy Level . . .	373
Unsupported Statements in the [edit interfaces xe] Hierarchy Level . .	377
[edit poe] Configuration Statement Hierarchy on EX Series Switches . . . .	377
Supported Statements in the [edit poe] Hierarchy Level . . . . .	377
[edit policy-options] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	378
Supported Statements in the [edit policy-options] Hierarchy Level . .	378
Unsupported Statements in the [edit policy-options] Hierarchy	
Level . . . . .	389
[edit protocols] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	389
[edit protocols bfd] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	391
Supported Statements in the [edit protocols bfd] Hierarchy Level . . .	391
Unsupported Statements in the [edit protocols bfd] Hierarchy	
Level . . . . .	391
[edit protocols bgp] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	392
Supported Statements in the [edit protocols bgp] Hierarchy Level . .	392
Unsupported Statements in the [edit protocols bgp] Hierarchy	
Level . . . . .	400
[edit protocols connections] Configuration Statement Hierarchy on EX	
Series Switches . . . . .	401
Supported Statements in the [edit protocols connections] Hierarchy	
Level . . . . .	402
Unsupported Statements in the [edit protocols connections] Hierarchy	
Level . . . . .	402
[edit protocols dcbx] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	403
Supported Statements in the [edit protocols dcbx] Hierarchy Level . .	403
Unsupported Statements in the [edit protocols dcbx] Hierarchy	
Level . . . . .	403
[edit protocols dot1x] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	404
Supported Statements in the [edit protocols dot1x] Hierarchy Level . .	404
Unsupported Statements in the [edit protocols dot1x] Hierarchy	
Level . . . . .	405
[edit protocols igmp] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	406
Supported Statements in the [edit protocols igmp] Hierarchy Level . .	406
Unsupported Statements in the [edit protocols igmp] Hierarchy	
Level . . . . .	407

[edit protocols igmp-snooping] Configuration Statement Hierarchy on EX Series Switches . . . . .	407
Supported Statements in the [edit protocols igmp-snooping] Hierarchy Level . . . . .	407
Unsupported Statements in the [edit protocols igmp-snooping] Hierarchy Level . . . . .	408
[edit protocols isis] Configuration Statement Hierarchy on EX Series Switches . . . . .	408
Supported Statements in the [edit protocols isis] Hierarchy Level . . . . .	409
Unsupported Statements in the [edit protocols isis] Hierarchy Level . . . . .	411
[edit protocols lacp] Configuration Statement Hierarchy on EX Series Switches . . . . .	411
Supported Statements in the [edit protocols lacp] Hierarchy Level . . . . .	411
Unsupported Statements in the [edit protocols lacp] Hierarchy Level . . . . .	412
[edit protocols link-management] Configuration Statement Hierarchy on EX Series Switches . . . . .	412
Supported Statements in the [edit protocols link-management] Hierarchy Level . . . . .	412
Unsupported Statements in the [edit protocols link-management] Hierarchy Level . . . . .	413
[edit protocols lldp] Configuration Statement Hierarchy on EX Series Switches . . . . .	413
Supported Statements in the [edit protocols lldp] Hierarchy Level . . . . .	414
Unsupported Statements in the [edit protocols lldp] Hierarchy Level . . . . .	414
[edit protocols lldp-med] Configuration Statement Hierarchy on EX Series Switches . . . . .	415
Supported Statements in the [edit protocols lldp-med] Hierarchy Level . . . . .	415
Unsupported Statements in the [edit protocols lldp-med] Hierarchy Level . . . . .	415
[edit protocols mld] Configuration Statement Hierarchy on EX Series Switches . . . . .	416
Supported Statements in the [edit protocols mld] Hierarchy Level . . . . .	416
Unsupported Statements in the [edit protocols mld] Hierarchy Level . . . . .	417
[edit protocols mld-snooping] Configuration Statement Hierarchy on EX Series Switches . . . . .	417
Supported Statements in the [edit protocols mld-snooping] Hierarchy Level . . . . .	417
Unsupported Statements in the [edit protocols mld-snooping] Hierarchy Level . . . . .	418
[edit protocols mpls] Configuration Statement Hierarchy on EX Series Switches . . . . .	418
Supported Statements in the [edit protocols mpls] Hierarchy Level . . . . .	418
Unsupported Statements in the [edit protocols mpls] Hierarchy Level . . . . .	420

[edit protocols msdp] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	429
Supported Statements in the [edit protocols msdp] Hierarchy	
Level . . . . .	430
Unsupported Statements in the [edit protocols msdp] Hierarchy	
Level . . . . .	431
[edit protocols mstp] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	431
Supported Statements in the [edit protocols mstp] Hierarchy Level . .	432
Unsupported Statements in the [edit protocols mstp] Hierarchy	
Level . . . . .	433
[edit protocols mvrp] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	433
Supported Statements in the [edit protocols mvrp] Hierarchy Level . .	433
Unsupported Statements in the [edit protocols mvrp] Hierarchy	
Level . . . . .	434
[edit protocols neighbor-discovery] Configuration Statement Hierarchy on	
EX Series Switches . . . . .	434
Supported Statements in the [edit protocols neighbor-discovery]	
Hierarchy Level . . . . .	434
Unsupported Statements in the [edit neighbor-discovery] Hierarchy	
Level . . . . .	435
[edit protocols oam] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	435
Supported Statements in the [edit protocols oam] Hierarchy Level . .	435
Unsupported Statements in the [edit protocols oam] Hierarchy	
Level . . . . .	437
[edit protocols ospf] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	438
Supported Statements in the [edit protocols ospf] Hierarchy Level . .	438
Unsupported Statements in the [edit protocols ospf] Hierarchy	
Level . . . . .	440
[edit protocols ospf3] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	441
Supported Statements in the [edit protocols ospf3] Hierarchy	
Level . . . . .	441
Unsupported Statements in the [edit protocols ospf3] Hierarchy	
Level . . . . .	444
[edit protocols pim] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	444
Supported Statements in the [edit protocols pim] Hierarchy Level . . .	445
Unsupported Statements in the [edit protocols pim] Hierarchy	
Level . . . . .	447
[edit protocols rip] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	447
Supported Statements in the [edit protocols rip] Hierarchy Level . . .	448
Unsupported Statements in the [edit protocols rip] Hierarchy Level . .	449



[edit protocols ripng] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	450
Supported Statements in the [edit protocols ripng] Hierarchy Level . .	450
Unsupported Statements in the [edit protocols ripng] Hierarchy	
Level . . . . .	451
[edit protocols router-advertisement] Configuration Statement Hierarchy	
on EX Series Switches . . . . .	451
Supported Statements in the [edit protocols router-advertisement]	
Hierarchy Level . . . . .	452
Unsupported Statements in the [edit protocols router-advertisement]	
Hierarchy Level . . . . .	452
[edit protocols router-discovery] Configuration Statement Hierarchy on EX	
Series Switches . . . . .	452
Supported Statements in the [edit protocols router-discovery] Hierarchy	
Level . . . . .	453
Unsupported Statements in the [edit protocols router-discovery]	
Hierarchy Level . . . . .	453
[edit protocols rstp] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	453
Supported Statements in the [edit protocols rstp] Hierarchy Level . .	454
Unsupported Statements in the [edit protocols rstp] Hierarchy	
Level . . . . .	454
[edit protocols rsvp] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	455
Supported Statements in the [edit protocols rsvp] Hierarchy Level . .	455
Unsupported Statements in the [edit protocols rsvp] Hierarchy	
Level . . . . .	456
[edit protocols sflow] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	459
Supported Statements in the [edit protocols sflow] Hierarchy Level . .	459
Unsupported Statements in the [edit sflow] Hierarchy Level . . . . .	460
[edit protocols stp] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	460
Supported Statements in the [edit protocols stp] Hierarchy Level . .	460
Unsupported Statements in the [edit protocols stp] Hierarchy	
Level . . . . .	461
[edit protocols uplink-failure-detection] Configuration Statement Hierarchy	
on EX Series Switches . . . . .	461
Supported Statements in the [edit protocols uplink-failure-detection]	
Hierarchy Level . . . . .	462
Unsupported Statements in the [edit protocols uplink-failure-detection]	
Hierarchy Level . . . . .	462
[edit protocols vrrp] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	462
Supported Statements in the [edit protocols vrrp] Hierarchy Level . .	463
Unsupported Statements in the [edit protocols vrrp] Hierarchy	
Level . . . . .	463

[edit protocols vstp] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	463
Supported Statements in the [edit protocols vstp] Hierarchy Level . . .	464
Unsupported Statements in the [edit protocols vstp] Hierarchy	
Level . . . . .	464
[edit redundant-power-system] Configuration Statement Hierarchy on EX	
Series Switches . . . . .	465
Supported Statements in the [edit redundant-power-system] Hierarchy	
Level . . . . .	465
Unsupported Statements in the [edit redundant-power-system]	
Hierarchy Level . . . . .	465
[edit routing-options] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	465
Supported Statements in the [edit routing-options] Hierarchy Level . .	466
Unsupported Statements in the [edit routing-options] Hierarchy	
Level . . . . .	482
[edit security] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	482
Supported Statements in the [edit security] Hierarchy Level . . . . .	483
Unsupported Statements in the [edit security] Hierarchy Level . . . . .	485
[edit services] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	485
Supported Statements in the [edit services] Hierarchy Level . . . . .	486
Unsupported Statements in the [edit services] Hierarchy Level . . . . .	488
[edit snmp] Configuration Statement Hierarchy on EX Series Switches . .	488
Supported Statements in the [edit snmp] Hierarchy Level . . . . .	489
Unsupported Statements in the [edit snmp] Hierarchy Level . . . . .	493
[edit system] Configuration Statement Hierarchy on EX Series Switches . .	495
Supported Statements in the [edit system] Hierarchy Level . . . . .	495
Unsupported Statements in the [edit system] Hierarchy Level . . . . .	507
[edit virtual-chassis] Configuration Statement Hierarchy . . . . .	508
Supported Statements in the [edit virtual-chassis] Hierarchy Level . .	508
Unsupported Statements in the [edit virtual-chassis] Hierarchy	
Level . . . . .	509
[edit vlans] Configuration Statement Hierarchy on EX Series Switches . .	509
Supported Statements in the [edit vlans] Hierarchy Level . . . . .	510
Unsupported Statements in the [edit vlans] Hierarchy Level . . . . .	511
<b>Chapter 8</b>	
<b>Administration . . . . .</b>	<b>513</b>
Operational Commands . . . . .	513
set cli directory . . . . .	514
set cli idle-timeout . . . . .	515
set cli prompt . . . . .	516
set cli restart-on-upgrade . . . . .	517
set cli screen-length . . . . .	518
set cli screen-width . . . . .	519
set cli timestamp . . . . .	520
start shell . . . . .	521

<b>Part 4</b>	<b>Access and User Management</b>	
<b>Chapter 9</b>	<b>Overview</b>	<b>525</b>
	Software Overview	525
	Understanding Software Infrastructure and Processes	525
	Routing Engine and Packet Forwarding Engine	525
	Junos OS Processes	526
<b>Chapter 10</b>	<b>Configuration</b>	<b>529</b>
	Configuration Tasks	529
	Configuring Management Access for the EX Series Switch (J-Web Procedure)	529
	Generating SSL Certificates to Be Used for Secure Web Access	532
	Configuring MS-CHAPv2 to Provide Password-Change Support (CLI Procedure)	533
	Configuration Statements	533
	allow-commands	535
	allow-configuration	536
	announcement	537
	archive-sites	537
	authentication (Login)	538
	authentication-order	539
	change-type	540
	class (Assigning a Class to an Individual User)	540
	class (Defining Login Classes)	541
	class-usage-profile	542
	counters	543
	deny-commands	543
	deny-configuration	544
	destination-classes	545
	fields (for Interface Profiles)	546
	file (Associating with a Profile)	547
	file (Configuring a Log File)	548
	files	549
	filter-profile	550
	format	551
	full-name	552
	idle-timeout (System-Login)	552
	interface-profile	553
	interval (Accounting Options)	554
	login	555
	login-alarms	556
	login-tip	556
	maximum-length	557
	message	558
	mib-profile	559
	minimum-changes	560
	minimum-length	561
	object-names	562
	operation	562

	password (Login) . . . . .	563
	permissions . . . . .	564
	radius-options (edit system) . . . . .	565
	retry-options . . . . .	566
	root-authentication . . . . .	567
	root-login . . . . .	568
	routing-engine-profile . . . . .	569
	size . . . . .	570
	source-classes . . . . .	570
	start-time (Log File Transfer) . . . . .	571
	tacplus-options . . . . .	572
	tacplus-server . . . . .	573
	traceoptions (Address-Assignment Pool) . . . . .	574
	transfer-interval . . . . .	575
	uid . . . . .	576
	user (Access) . . . . .	577
<b>Chapter 11</b>	<b>Administration . . . . .</b>	<b>579</b>
	Routine Monitoring . . . . .	579
	Managing Users (J-Web Procedure) . . . . .	579
	Operational Commands . . . . .	581
	request message . . . . .	582
	show subscribers . . . . .	583
<b>Chapter 12</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>601</b>
	Troubleshooting Loss of the Root Password . . . . .	601
<b>Part 5</b>	<b>Configuration and File Management</b>	
<b>Chapter 13</b>	<b>Overview . . . . .</b>	<b>607</b>
	Software Overview . . . . .	607
	Understanding Software Infrastructure and Processes . . . . .	607
	Routing Engine and Packet Forwarding Engine . . . . .	607
	Junos OS Processes . . . . .	608
	Configuration Files Overview . . . . .	609
	Understanding Configuration Files for EX Series Switches . . . . .	609
	Using the CLI Viewer in the J-Web Interface to View Configuration Text . . .	610
	Using the CLI Editor in the J-Web Interface to Edit Configuration Text . . .	610
	Using the Point and Click CLI Tool in the J-Web Interface to Edit Configuration Text . . . . .	611
	Understanding Automatic Refreshing of Scripts on EX Series Switches . . .	613
	Understanding Autoinstallation of Configuration Files . . . . .	613
	Typical Uses for Autoinstallation . . . . .	613
	Autoinstallation Configuration Files and IP Addresses . . . . .	614
	Typical Autoinstallation Process on a New Switch . . . . .	614
	Configuration Files Terms . . . . .	615

<b>Chapter 14</b>	<b>Configuration</b>	<b>617</b>
	Configuration Tasks	617
	Uploading a Configuration File (CLI Procedure)	617
	Uploading a Configuration File (J-Web Procedure)	619
	Loading a Previous Configuration File (CLI Procedure)	619
	Reverting to the Default Factory Configuration for the EX Series Switch	620
	Reverting to the Factory-Default Configuration by Using the LCD Panel	621
	Reverting to the Factory-Default Configuration by Using the request system zeroize Command	622
	Reverting to the Factory-Default Configuration by Using the load factory-default Command	622
	Reverting to the Rescue Configuration for the EX Series Switch	624
	Setting or Deleting the Rescue Configuration (CLI Procedure)	625
	Setting or Deleting the Rescue Configuration (J-Web Procedure)	625
	Configuring Autoinstallation of Configuration Files (CLI Procedure)	626
	Using the Commit Options to Commit Configuration Changes (J-Web Procedure)	628
	Configuration Statements	629
	archival	630
	archive-sites (Configuration File)	631
	autoinstallation	633
	synchronize (Commit configuration)	634
	configuration	635
	configuration-servers	636
	interfaces	637
	transfer-interval (Configuration)	638
	transfer-on-commit	639
<b>Chapter 15</b>	<b>Administration</b>	<b>641</b>
	Routine Monitoring	641
	Managing Configuration Files Through the Configuration History (J-Web Procedure)	641
	Displaying Configuration History	641
	Displaying Users Editing the Configuration	642
	Comparing Configuration Files with the J-Web Interface	643
	Downloading a Configuration File with the J-Web Interface	643
	Loading a Previous Configuration File with the J-Web Interface	643
	Verifying Autoinstallation Status	644
	Operational Commands	645
	clear log	646
	clear system commit	647
	file archive	648
	file checksum md5	650
	file checksum sha1	651
	file checksum sha-256	652
	file compare	653
	file copy	656
	file delete	658

	file list . . . . .	659
	file rename . . . . .	661
	file show . . . . .	663
	request system configuration rescue delete . . . . .	665
	request system configuration rescue save . . . . .	666
	request system scripts refresh-from commit . . . . .	667
	request system scripts refresh-from event . . . . .	668
	request system scripts refresh-from op . . . . .	669
	request system zeroize . . . . .	670
	show system commit . . . . .	675
	show system configuration archival . . . . .	678
	show system configuration rescue . . . . .	679
	show system rollback . . . . .	681
	test configuration . . . . .	683
<b>Chapter 16</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>685</b>
	Troubleshooting Loss of the Root Password . . . . .	685
<b>Part 6</b>	<b>Software Installation</b>	
<b>Chapter 17</b>	<b>Overview . . . . .</b>	<b>691</b>
	Installation Overview . . . . .	691
	Understanding Software Installation on EX Series Switches . . . . .	691
	Overview of the Software Installation Process . . . . .	691
	Software Package Security . . . . .	692
	Installing Software on a Virtual Chassis . . . . .	692
	Installing Software on Switches with Redundant Routing Engines . . . . .	692
	Installing Software Using Automatic Software Download . . . . .	693
	Autoinstalling a Configuration File on an EX2200 or EX3300 Switch from a Disk-on-Key USB Memory Stick . . . . .	693
	Troubleshooting Software Installation . . . . .	693
	Understanding System Snapshot on EX Series Switches . . . . .	694
	Understanding Resilient Dual-Root Partitions on Switches . . . . .	695
	Resilient Dual-Root Partition Scheme (Junos OS Release 10.4R3 and Later) . . . . .	695
	Automatic Fixing of Corrupted Primary Root Partition with the Automatic Snapshot Feature . . . . .	696
	Earlier Partition Scheme (Junos OS Release 10.4R2 and Earlier) . . . . .	697
	Understanding Upgrading or Downgrading Between Resilient Dual-Root Partition Releases and Earlier Releases . . . . .	697
	Junos OS Package Names . . . . .	698
	Software Overview . . . . .	700
	Understanding Software Infrastructure and Processes . . . . .	700
	Routing Engine and Packet Forwarding Engine . . . . .	700
	Junos OS Processes . . . . .	700
	Licenses Overview . . . . .	701
	Understanding Software Licenses for EX Series Switches . . . . .	702
	Purchasing a Software Feature License . . . . .	702
	Features Requiring a License on EX2200 Switches . . . . .	703
	Features Requiring a License on EX3300 Switches . . . . .	704

	Features Requiring a License on EX4300 Switches . . . . .	705
	Features Requiring a License on EX4600 Switches . . . . .	706
	Features Requiring a License on EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 Switches . . . . .	707
	License Warning Messages . . . . .	708
<b>Chapter 18</b>	<b>Configuration . . . . .</b>	<b>711</b>
	Registering the Switch . . . . .	711
	Registering the EX Series Switch with the J-Web Interface . . . . .	711
	Booting the Switch . . . . .	711
	Booting an EX Series Switch Using a Software Package Stored on a USB Flash Drive . . . . .	712
	Creating a Snapshot and Using It to Boot an EX Series Switch . . . . .	713
	Creating a Snapshot on a USB Flash Drive and Using It to Boot the Switch . . . . .	713
	Creating a Snapshot on an Internal Flash Drive and Using it to Boot the Switch . . . . .	714
	Creating a Snapshot on the Alternate Slice of the Boot Media . . . . .	714
	Managing Licenses . . . . .	714
	Managing Licenses for the EX Series Switch (CLI Procedure) . . . . .	714
	Adding New Licenses . . . . .	715
	Deleting Licenses . . . . .	715
	Saving License Keys . . . . .	715
	Managing Licenses for the EX Series Switch (J-Web Procedure) . . . . .	716
	Adding New Licenses . . . . .	716
	Deleting Licenses . . . . .	716
	Displaying License Keys . . . . .	717
	Downloading Licenses . . . . .	717
<b>Chapter 19</b>	<b>Administration . . . . .</b>	<b>719</b>
	Routine Monitoring . . . . .	719
	Verifying That a System Snapshot Was Created on an EX Series Switch . . . . .	719
	Verifying Junos OS and Boot Loader Software Versions on an EX Series Switch . . . . .	720
	Verifying the Number of Partitions and File System Mountings . . . . .	720
	Verifying the Loader Software Version . . . . .	721
	Verifying Which Root Partition Is Active . . . . .	722
	Verifying the Junos OS Version in Each Root Partition . . . . .	722
	Monitoring Licenses . . . . .	723
	License Key Components for the EX Series Switch . . . . .	723
	Monitoring Licenses for the EX Series Switch . . . . .	724
	Displaying Installed Licenses and License Usage Details . . . . .	724
	Displaying Installed License Keys . . . . .	725
	Operational Commands . . . . .	725
	request system license add . . . . .	727
	request system license delete . . . . .	728
	request system license save . . . . .	729
	request system reboot . . . . .	730
	request system reboot . . . . .	735
	request system snapshot . . . . .	737

	request system software add . . . . .	743
	request system software delete . . . . .	752
	request system software rollback . . . . .	756
	request system software validate . . . . .	760
	show system auto-snapshot . . . . .	763
	show system boot-messages . . . . .	765
	show system license . . . . .	772
	show system snapshot . . . . .	779
	show system storage partitions (EX Series Switches Only) . . . . .	782
<b>Chapter 20</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>785</b>
	Troubleshooting Software Installation . . . . .	785
	Recovering from a Failed Software Upgrade on an EX Series Switch . . . .	785
	Rebooting from the Inactive Partition . . . . .	786
	Freeing Disk Space for Software Installation . . . . .	787
	Installation from the Boot Loader Generates 'cannot open package' Error . . . . .	787
	Troubleshooting a Switch That Has Booted from the Backup Junos OS Image . . . . .	788
	Resilient Dual-Root Partitions Frequently Asked Questions . . . . .	789
	What Happens to My Files If the System Detects a File System Corruption and Automatic Snapshot is Enabled? . . . . .	789
	What Happens to My Files If the System Detects a File System Corruption and Automatic Snapshot is Not Enabled? . . . . .	790
	How Will I Be Informed If My Switch Boots from the Alternate Slice Because of Corruption in the Root File System? . . . . .	791
<b>Part 7</b>	<b>System Monitoring</b>	
<b>Chapter 21</b>	<b>Overview . . . . .</b>	<b>795</b>
	Software Overview . . . . .	795
	Understanding Software Infrastructure and Processes . . . . .	795
	Routing Engine and Packet Forwarding Engine . . . . .	795
	Junos OS Processes . . . . .	796
	Alarms Overview . . . . .	797
	Understanding Alarm Types and Severity Levels on EX Series Switches . . .	797
	Dashboard Overview . . . . .	798
	Dashboard for EX Series Switches . . . . .	799
	Graphical Chassis Viewer . . . . .	799
	System Information Panel . . . . .	801
	Health Status Panel . . . . .	803
	Capacity Utilization Panel . . . . .	805
	Alarms Panel . . . . .	805
	File System Usage . . . . .	806
	Chassis Viewer . . . . .	806
	Hardware/CLI Terminology Mapping Overview . . . . .	817
	EX4300 Switch Hardware and CLI Terminology Mapping . . . . .	817



<b>Chapter 22</b>	<b>Configuration</b>	<b>821</b>
	Configuration Statements	821
	facility-override	822
	file (System Logging)	823
	files	824
	host (System)	825
	interface (Accounting or Sampling)	827
	log-prefix (System)	827
	match	828
	size (System)	829
	structured-data	830
	syslog (System)	831
	time-format	833
	time-zone	834
	user (System Logging)	836
	world-readable (System)	837
<b>Chapter 23</b>	<b>Administration</b>	<b>839</b>
	Routine Monitoring	839
	Monitoring System Log Messages	839
	Checking Active Alarms with the J-Web Interface	842
	Monitoring Switch Control Traffic	843
	Monitoring System Properties	846
	Monitoring Chassis Information	847
	Monitoring System Process Information	850
	Managing Log, Temporary, and Crash Files on the Switch (J-Web Procedure)	850
	Cleaning Up Files	851
	Downloading Files	851
	Deleting Files	852
	Operational Commands	853
	clear log	855
	file archive	856
	file checksum md5	858
	file checksum sha1	859
	file checksum sha-256	860
	file compare	861
	file copy	864
	file delete	866
	file list	867
	file rename	869
	file show	871
	monitor list	873
	monitor start	874
	monitor stop	876
	request chassis cb	877
	request chassis fabric plane	880
	request chassis fpc	882
	request system configuration rescue delete	886

request system configuration rescue save . . . . .	887
request system scripts refresh-from commit . . . . .	888
request system scripts refresh-from event . . . . .	889
request system scripts refresh-from op . . . . .	890
show chassis alarms . . . . .	891
show chassis environment . . . . .	906
show chassis environment cb . . . . .	970
show chassis environment fpc . . . . .	988
show chassis environment routing-engine . . . . .	1014
show chassis ethernet-switch . . . . .	1019
show chassis fabric fpcs . . . . .	1063
show chassis fabric map . . . . .	1101
show chassis fabric plane . . . . .	1108
show chassis fabric plane-location . . . . .	1150
show chassis fabric summary . . . . .	1155
show chassis fpc . . . . .	1160
show chassis led . . . . .	1197
show chassis location . . . . .	1207
show chassis pic . . . . .	1211
show chassis routing-engine . . . . .	1227
show log . . . . .	1249
show pfe next-hop . . . . .	1252
show pfe route . . . . .	1257
show pfe terse . . . . .	1266
show system alarms . . . . .	1268
show system audit . . . . .	1271
show system buffers . . . . .	1279
show system connections . . . . .	1286
show system core-dumps . . . . .	1305
show system directory-usage . . . . .	1319
show system processes . . . . .	1323

## Part 8

## System Services

### Chapter 24

<b>Overview . . . . .</b>	<b>1353</b>
Software Overview . . . . .	1353
Understanding Software Infrastructure and Processes . . . . .	1353
Routing Engine and Packet Forwarding Engine . . . . .	1353
Junos OS Processes . . . . .	1354
DHCP Local Server . . . . .	1355
Extended DHCP Local Server Overview . . . . .	1356
Interaction Among the DHCP Client, Extended DHCP Local Server, and Address-Assignment Pools . . . . .	1358
Providing DHCP Client Configuration Information . . . . .	1358
Minimal Configuration for Clients . . . . .	1360
DHCP Local Server and Address-Assignment Pools . . . . .	1360
DHCPv6 Local Server Overview . . . . .	1361
DHCP Local Server Handling of Client Information Request Messages . . .	1363
Configuring Group-Specific DHCP Local Server Options . . . . .	1364

Understanding Dynamic Reconfiguration of Extended DHCP Local Server	
Clients	1364
Default Client/Server Interaction	1364
Dynamic Client/Server Interaction for DHCPv4	1365
Dynamic Client/Server Interaction for DHCPv6	1365
Manually Forcing the Local Server to Initiate the Reconfiguration	
Process	1366
Action Taken for Events That Occur During a Reconfiguration	1366
DHCP Snooping Support	1367
DHCP Auto Logout Overview	1368
Auto Logout Overview	1368
How DHCP Identifies and Releases Clients	1368
Option 60 and Option 82 Requirements	1369
Address-Assignment Pools Overview	1370
Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6	
IA_NA Option	1371
Multiple Address Assignment for DHCPv6 Clients	1371
Multiple Address Assignment Using Local Address Pools or RADIUS	1371
Junos OS Predefined Variable for Multiple DHCPv6 Address	
Assignment	1372
Centrally Configured Opaque DHCP Options	1373
Data Flow for RADIUS-Sourced DHCP Options	1375
Multiple VSA 26-55 Instances Configuration	1376
DHCP Options That Cannot Be Centrally Configured	1376
Port Number Requirements for DHCP Firewall Filters	1377
DHCP Relay Agent	1378
Extended DHCP Relay Agent Overview	1379
Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP	
Servers	1380
DHCP Liveness Detection	1381
DHCP Relay Proxy Overview	1382
Interaction Among DHCP Relay Proxy, DHCP Client, and DHCP	
Servers	1382
DHCPv6 Relay Agent Overview	1384
Configuring Group-Specific DHCP Relay Options	1384
DHCP Snooping Support	1385
DHCP Auto Logout Overview	1387
Auto Logout Overview	1387
How DHCP Identifies and Releases Clients	1387
Option 60 and Option 82 Requirements	1388
Public Key Cryptography Overview	1388
Understanding Public Key Cryptography on Switches	1389
Public Key Infrastructure (PKI) and Digital Certificates	1389
Self-Signed Certificates Overview	1390
Understanding Self-Signed Certificates on EX Series Switches	1390

<b>Chapter 25</b>	<b>Configuration</b>	<b>1393</b>
	DHCP Local Server Examples	1393
	Example: Minimum Extended DHCP Local Server Configuration	1393
	Example: Extended DHCP Local Server Configuration with Optional Pool Matching	1394
	Example: Configuring Group Liveness Detection for DHCP Local Server Clients	1395
	DHCP Relay Agent Examples	1398
	Example: Minimum DHCP Relay Agent Configuration	1398
	Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings	1399
	Example: Configuring DHCP Snooping Support for DHCP Relay Agent	1403
	Configuration Tasks	1404
	Configuring DHCP Services (J-Web Procedure)	1405
	Configuring DHCP Services (J-Web Procedure) on EX Series Switches	1405
	Configuring DHCP Services on EX4300 Switches (J-Web Procedure)	1408
	Configuring a DHCP SIP Server (CLI Procedure)	1412
	Configuring a DHCP Client (CLI Procedure)	1413
	Configuring a DHCP Server on Switches (CLI Procedure)	1414
	Configuring an Extended DHCP Server on a Switch	1415
	Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)	1416
	Enabling HTTPS and XNM-SSL Services on Switches Using Self-Signed Certificates (CLI Procedure)	1417
	Manually Generating Self-Signed Certificates on Switches (CLI Procedure)	1417
	Generating a Public-Private Key Pair on Switches	1418
	Generating Self-Signed Certificates on Switches	1418
	Deleting Self-Signed Certificates (CLI Procedure)	1419
	Configuration Tasks for DHCP Local Server	1419
	Using External AAA Authentication Services with DHCP	1420
	Grouping Interfaces with Common DHCP Configurations	1421
	Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces	1422
	Overriding Default DHCP Local Server Configuration Settings	1423
	Specifying the Maximum Number of DHCP Clients Per Interface	1424
	Automatically Logging Out DHCP Clients	1426
	Enabling Processing of Client Information Requests	1427
	Specifying the Delegated Address Pool for IPv6 Prefix Assignment	1428
	Enabling DHCPv6 Rapid Commit Support	1428
	Deleting DHCP Local Server and DHCP Relay Override Settings	1429
	Configuring Dynamic Client Reconfiguration of Extended Local Server Clients	1429
	Configuring Dynamic Reconfiguration Attempts for DHCP Clients	1432
	Configuring Deletion of the Client When Dynamic Reconfiguration Fails	1433
	Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect	1433

Configuring a Token for DHCP Local Server Authentication . . . . .	1434
Preventing Binding of Clients That Do Not Support Reconfigure	
Messages . . . . .	1434
Requesting DHCP Local Server to Initiate Reconfiguration of Client	
Bindings . . . . .	1435
Configuring Detection of DHCP Local Server Client Connectivity . . . . .	1436
Configuring DHCP Snooped Packets Forwarding Support for DHCP Local	
Server . . . . .	1437
Configuring Passwords for Usernames . . . . .	1439
Creating Unique Usernames for DHCP Clients . . . . .	1439
Configuring How the Extended DHCP Local Server Determines Which	
Address-Assignment Pool to Use . . . . .	1442
Configuration Tasks for DHCP Relay Agent . . . . .	1443
Using External AAA Authentication Services with DHCP . . . . .	1444
Grouping Interfaces with Common DHCP Configurations . . . . .	1445
Guidelines for Configuring Interface Ranges for Groups of DHCP	
Interfaces . . . . .	1446
Overriding the Default DHCP Relay Configuration Settings . . . . .	1447
Changing the Gateway IP Address (giaddr) Field to the giaddr of the DHCP	
Relay Agent . . . . .	1449
Replacing the DHCP Relay Request and Release Packet Source Address . .	1450
Overriding Option 82 Information . . . . .	1450
Using Layer 2 Unicast Transmission for DHCP Packets . . . . .	1451
Trusting Option 82 Information . . . . .	1451
Specifying the Maximum Number of DHCP Clients Per Interface . . . . .	1451
Automatically Logging Out DHCP Clients . . . . .	1453
How DHCP Relay Agent Uses Option 82 for Auto Logout . . . . .	1454
Configuring DHCP Snooping for DHCP Relay Agent . . . . .	1455
Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay	
Agent . . . . .	1455
Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay	
Agent . . . . .	1460
Sending Release Messages When Clients Are Deleted . . . . .	1462
Disabling Automatic Binding of Stray DHCP Requests . . . . .	1463
Using DHCP Relay Agent Option 82 Information . . . . .	1464
Configuring Option 82 Information . . . . .	1465
Including a Prefix in DHCP Options . . . . .	1466
Including a Textual Description in DHCP Options . . . . .	1468
Configuring Server Groups . . . . .	1469
Configuring Active Server Groups . . . . .	1469
Enabling DHCP Relay Proxy Mode . . . . .	1470
Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets . .	1470
DHCP Liveness Detection Overview . . . . .	1472
Configuring Detection of DHCP Relay or DHCP Relay Proxy Client	
Connectivity . . . . .	1473
Disabling DHCP Relay . . . . .	1474
DHCP Local Server Configuration Statements . . . . .	1475
attempts (DHCP Local Server) . . . . .	1477
authentication (DHCP Local Server) . . . . .	1478

bfd	1479
circuit-type (DHCP Local Server)	1480
clear-on-abort (DHCP Local Server)	1481
client-discover-match (DHCP Local Server)	1482
client-id (DHCP Local Server)	1483
delegated-pool (DHCP Local Server)	1484
delimiter (DHCP Local Server)	1485
detection-time	1486
dhcp (DHCP Client)	1487
dhcp-local-server	1488
dhcpv6 (DHCP Local Server)	1493
domain-name (DHCP Local Server)	1496
dynamic-profile (DHCP Local Server)	1497
external-authority	1498
failure-action	1499
forward-snooped-clients (DHCP Local Server)	1500
group (DHCP Local Server)	1501
holddown-interval	1503
interface (DHCP Local Server)	1504
interface-client-limit (DHCP Local Server)	1506
interface-delete (Subscriber Management or DHCP Client Management)	1507
interface-name (DHCP Local Server)	1508
ip-address-first	1509
liveness-detection	1510
mac-address (DHCP Local Server)	1511
method	1512
minimum-interval	1513
minimum-receive-interval	1514
multiplier	1515
no-adaptation	1516
option-60 (DHCP Local Server)	1517
option-82 (DHCP Local Server Authentication)	1518
option-82 (DHCP Local Server Pool Matching)	1519
overrides (DHCP Local Server)	1520
password (DHCP Local Server)	1522
pool (DHCP Local Server Overrides)	1523
pool-match-order	1524
process-inform	1525
radius-disconnect (DHCP Local Server)	1527
rapid-commit (DHCPv6 Local Server)	1528
reconfigure (DHCP Local Server)	1529
relay-agent-interface-id (DHCP Local Server)	1530
relay-agent-remote-id (DHCP Local Server)	1531
routing-instance-name (DHCP Local Server)	1532
service-profile (DHCP Local Server)	1533
session-mode	1534
strict (DHCP Local Server)	1535
threshold (detection-time)	1536

threshold (transmit-interval) .....	1537
timeout (DHCP Local Server) .....	1538
token (DHCP Local Server) .....	1539
traceoptions (DHCP Server) .....	1540
transmit-interval .....	1543
trigger (DHCP Local Server) .....	1544
use-primary (DHCP Local Server) .....	1545
user-prefix (DHCP Local Server) .....	1546
username-include (DHCP Local Server) .....	1547
version (BFD) .....	1548
DHCP Relay Agent Configuration Statements .....	1548
[edit forwarding-options dhcp-relay] Configuration Statement Hierarchy	
for EX Series Switches .....	1550
Supported Statements in the [edit forwarding-options dhcp-relay]	
Hierarchy Level .....	1551
Unsupported Statements in the [edit forwarding-options dhcp-relay]	
Hierarchy Level .....	1555
access (Dynamic Access Routes) .....	1556
access-internal (Dynamic Access-Internal Routes) .....	1557
active-server-group .....	1558
allow-snooped-clients .....	1559
always-write-giaddr .....	1560
always-write-option-82 .....	1561
authentication (DHCP Relay Agent) .....	1562
bfd .....	1563
circuit-id (DHCP Relay Agent) .....	1564
circuit-type (DHCP Relay Agent) .....	1565
client-discover-match (DHCP Relay Agent) .....	1566
client-id (DHCP Relay Agent) .....	1567
delete-binding-on-renegotiation .....	1567
delimiter (DHCP Relay Agent) .....	1568
detection-time .....	1569
dhcp-relay .....	1570
dhcpv6 (DHCP Relay Agent) .....	1576
disable-relay .....	1579
domain-name (DHCP Relay Agent) .....	1580
drop (DHCP Relay Agent Option) .....	1581
dynamic-profile (DHCP Relay Agent) .....	1582
failure-action .....	1583
forward-snooped-clients (DHCP Relay Agent) .....	1584
group (DHCP Relay Agent) .....	1585
holddown-interval .....	1588
interface (DHCP Relay Agent) .....	1589
interface-client-limit (DHCP Relay Agent) .....	1591
interface-delete (Subscriber Management or DHCP Client	
Management) .....	1592
interface-name (DHCP Relay Agent) .....	1593
layer2-unicast-replies .....	1594
liveness-detection .....	1595

local-server-group (DHCP Relay Agent Option) .....	1596
mac-address (DHCP Relay Agent) .....	1597
method .....	1598
minimum-interval .....	1599
minimum-receive-interval .....	1600
multiplier .....	1601
next-hop (Dynamic Access-Internal Routes) .....	1602
no-adaptation .....	1603
no-allow-snooped-clients .....	1604
no-bind-on-request (DHCP Relay Agent) .....	1605
option-60 (DHCP Relay Agent) .....	1606
option-82 (DHCP Relay Agent) .....	1607
option-number (DHCP Relay Agent Option) .....	1608
overrides (DHCP Relay Agent) .....	1609
password (DHCP Relay Agent) .....	1611
preference (Subscriber Management) .....	1612
prefix (DHCP Relay Agent) .....	1613
proxy-mode .....	1614
relay-agent-interface-id (DHCPv6 Relay Agent) .....	1615
relay-agent-remote-id (DHCPv6 Relay Agent Username) .....	1616
relay-option (DHCP Relay Agent) .....	1617
relay-option-82 .....	1618
relay-server-group (DHCP Relay Agent Option) .....	1619
replace-ip-source-with .....	1620
routing-instance-name (DHCP Relay Agent) .....	1621
send-release-on-delete (DHCP Relay Agent) .....	1622
server-group .....	1623
service-profile (DHCP Relay Agent) .....	1624
session-mode .....	1625
threshold (detection-time) .....	1626
threshold (transmit-interval) .....	1627
trace (DHCP Relay Agent) .....	1628
transmit-interval .....	1629
trust-option-82 .....	1630
use-interface-description .....	1631
use-primary (DHCP Relay Agent) .....	1633
user-prefix (DHCP Relay Agent) .....	1634
username-include (DHCP Relay Agent) .....	1635
version (BFD) .....	1636
Other Configuration Statements .....	1637
cache-size .....	1639
cache-timeout-negative .....	1640
certificates .....	1641
certification-authority .....	1642
connection-limit .....	1643
crl (Encryption Interface) .....	1644
domain-search .....	1644
encoding .....	1645
enrollment-retry .....	1645



enrollment-url .....	1646
family (for EX Series switches) .....	1647
file .....	1650
ftp .....	1651
http .....	1652
https .....	1653
ldap-url .....	1654
lease-time .....	1655
load-key-file .....	1656
local .....	1657
local-certificate .....	1658
maximum-certificates .....	1658
maximum-hop-count .....	1659
maximum-lease-time (DHCP) .....	1659
minimum-wait-time .....	1660
name-server .....	1660
no-listen .....	1661
outbound-ssh .....	1662
path-length .....	1665
port (HTTP/HTTPS) .....	1665
port (SRC Server) .....	1666
process-inform .....	1666
protocol-version .....	1667
rate-limit .....	1668
reconfigure .....	1669
retransmission-attempt .....	1670
retransmission-interval .....	1671
server (DNS and TFTP Service) .....	1671
server-address .....	1672
server-identifier .....	1673
servers .....	1674
service-deployment .....	1674
services (System Services) .....	1675
session (Time-out) .....	1677
sip-server .....	1678
source-address (SRC Software) .....	1678
source-address-giaddr .....	1679
ssh .....	1680
static-binding .....	1681
system-generated-certificate .....	1682
telnet .....	1682
tftp .....	1683
traceoptions (DNS and TFTP Packet Forwarding) .....	1684
traceoptions .....	1686
update-server .....	1688
web-management .....	1689
wins-server (System) .....	1690

<b>Chapter 26</b>	<b>Administration</b>	<b>1691</b>
	Routine Monitoring	1691
	Monitoring DHCP Services	1691
	Verifying and Managing DHCP Local Server Configurations	1695
	Verifying and Managing DHCP Local Server Configuration	1696
	Verifying and Managing DHCPv6 Local Server Configuration	1696
	Verifying and Managing DHCP Relay Agent Configurations	1697
	Verifying and Managing DHCP Relay Configuration	1697
	Verifying and Managing DHCPv6 Relay Configuration	1697
	DHCP Local Server Monitoring Commands	1697
	clear dhcp server binding	1699
	clear dhcp server statistics	1702
	clear dhcpv6 server binding	1704
	clear dhcpv6 server statistics	1706
	request dhcp server reconfigure	1707
	request dhcpv6 server reconfigure	1709
	request system reboot	1711
	show dhcp server binding	1716
	show dhcp server statistics	1721
	show dhcpv6 server binding	1724
	show dhcpv6 server statistics	1730
	DHCP Relay Agent Monitoring Commands	1732
	clear dhcp relay binding	1733
	clear dhcp relay statistics	1735
	clear dhcpv6 relay binding	1738
	clear dhcpv6 relay statistics	1741
	show dhcp relay binding	1743
	show dhcp relay statistics	1748
	show dhcpv6 relay binding	1751
	show dhcpv6 relay statistics	1757
	show route extensive	1760
	show route protocol	1777
	Other Operational Commands	1788
	clear security pki local-certificate	1789
	clear system services dhcp binding	1790
	clear system services dhcp conflict	1791
	clear system services dhcp statistics	1792
	request ipsec switch	1793
	request security certificate (signed)	1794
	request security certificate (unsigned)	1796
	request security key-pair	1797
	request security pki generate-key-pair	1798
	request security pki local-certificate generate-self-signed	1799
	show security pki local-certificate	1800
	show system services dhcp binding	1803
	show system services dhcp conflict	1806
	show system services dhcp global	1807
	show system services dhcp pool	1809
	show system services dhcp statistics	1811

	show system services service-deployment . . . . .	1814
	ssh . . . . .	1815
	telnet . . . . .	1817
<b>Part 9</b>	<b>Access Control</b>	
<b>Chapter 27</b>	<b>Overview . . . . .</b>	<b>1821</b>
	Access Control Overview . . . . .	1821
	802.1X for EX Series Switches Overview . . . . .	1821
	How 802.1X Authentication Works . . . . .	1821
	802.1X Features Overview . . . . .	1822
	Supported Features Related to 802.1X Authentication . . . . .	1823
	Understanding Authentication on EX Series Switches . . . . .	1824
	Sample Basic Authentication Topology . . . . .	1824
	802.1X Authentication . . . . .	1825
	MAC RADIUS Authentication . . . . .	1827
	Captive Portal Authentication . . . . .	1827
	Static MAC Bypass of Authentication . . . . .	1828
	Fallback of Authentication Methods . . . . .	1828
	Understanding Guest VLANs for 802.1X on EX Series Switches . . . . .	1830
	Understanding 802.1X and RADIUS Accounting on EX Series Switches . .	1830
	Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches . .	1831
	Understanding 802.1X and VoIP on EX Series Switches . . . . .	1834
	Understanding 802.1X and VSAs on EX Series Switches . . . . .	1837
	Understanding Dynamic VLANs for 802.1X on EX Series Switches . . . . .	1837
	Understanding Server Fail Fallback and Authentication on EX Series Switches . . . . .	1838
	Authentication Process Flow for EX Series Switches . . . . .	1839
	Understanding Authentication Session Timeout . . . . .	1841
<b>Chapter 28</b>	<b>Configuration . . . . .</b>	<b>1843</b>
	Configuration Examples . . . . .	1843
	Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch . . . . .	1843
	Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an EX Series Switch . . . . .	1847
	Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch . . . . .	1852
	Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch . . . . .	1858
	Example: Configuring VoIP on an EX Series Switch Without Including 802.1X Authentication . . . . .	1866
	Example: Configuring Static MAC Bypass of Authentication on an EX Series Switch . . . . .	1872
	Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to an EX Series Switch . . . . .	1875
	Example: Configuring MAC RADIUS Authentication on an EX Series Switch . . . . .	1881
	Example: Applying a Firewall Filter to 802.1X-Authenticated Supplicants Using RADIUS Server Attributes on an EX Series Switch . . . . .	1887

Example: Applying Firewall Filters to Multiple Supplicants on Interfaces	
Enabled for 802.1X or MAC RADIUS Authentication . . . . .	1893
Example: Setting Up Captive Portal Authentication on an EX Series	
Switch . . . . .	1898
Example: Configuring Fallback Options on EX Series Switches for EAP-TTLS	
Authentication and Odyssey Access Clients . . . . .	1903
Configuration Tasks . . . . .	1907
Configuring 802.1X Interface Settings (CLI Procedure) . . . . .	1908
Configuring 802.1X RADIUS Accounting (CLI Procedure) . . . . .	1909
Filtering 802.1X Supplicants Using RADIUS Server Attributes . . . . .	1910
Configuring Match Statements on the RADIUS Server . . . . .	1911
Applying a Port Firewall Filter from the RADIUS Server . . . . .	1913
Configuring LLDP (CLI Procedure) . . . . .	1913
Enabling LLDP on Interfaces . . . . .	1914
Adjusting LLDP Advertisement Settings . . . . .	1914
Adjusting SNMP Notification Settings of LLDP Changes . . . . .	1915
Specifying a Management Address for the LLDP Management TLV . .	1916
Configuring LLDP Power Negotiation . . . . .	1916
Configuring LLDP-MED (CLI Procedure) . . . . .	1917
Enabling LLDP-MED on Interfaces . . . . .	1917
Configuring Location Information Advertised by the Switch . . . . .	1917
Configuring for Fast Start . . . . .	1918
VSA Match Conditions and Actions . . . . .	1918
Configuring Server Fail Fallback (CLI Procedure) . . . . .	1921
Configuring MAC RADIUS Authentication (CLI Procedure) . . . . .	1923
Configuring Static MAC Bypass of Authentication (CLI Procedure) . . . .	1924
Specifying RADIUS Server Connections on an EX Series Switch (CLI	
Procedure) . . . . .	1925
Configuring Captive Portal Authentication (CLI Procedure) . . . . .	1926
Configuring Secure Access for Captive Portal . . . . .	1926
Enabling an Interface for Captive Portal . . . . .	1927
Configuring Bypass of Captive Portal Authentication . . . . .	1927
Designing a Captive Portal Authentication Login Page on an EX Series	
Switch . . . . .	1928
Controlling Authentication Session Timeouts (CLI Procedure) . . . . .	1930
Configuration Statements . . . . .	1931
[edit access] Configuration Statement Hierarchy on EX Series Switches . .	1933
Supported Statements in the [edit access] Hierarchy Level . . . . .	1933
Unsupported Statements in the [edit access] Hierarchy Level . . . . .	1938
[edit protocols dot1x] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	1939
Supported Statements in the [edit protocols dot1x] Hierarchy Level . .	1940
Unsupported Statements in the [edit protocols dot1x] Hierarchy	
Level . . . . .	1940
accounting . . . . .	1942
accounting (Access Profile) . . . . .	1943
accounting-order . . . . .	1944
accounting-port . . . . .	1944
address-assignment (Address-Assignment Pools) . . . . .	1945

address-protection	1946
authorization-order	1947
authentication-order	1948
authentication-whitelist	1949
authenticator	1950
client-accounting-algorithm	1951
client-authentication-algorithm	1952
coa-dynamic-variable-validation	1952
destination (Accounting)	1953
destination-host (Gx-Plus)	1954
destination-realm (Gx-Plus)	1954
diameter-instance (Gx-Plus)	1955
domain (Domain Map)	1956
domain-name-server (Routing Instances and Access Profiles)	1957
domain-name-server-inet (Routing Instances and Access Profiles)	1958
domain-name-server-inet6 (Routing Instances and Access Profiles)	1959
ethernet-port-type-virtual	1959
global (Gx-Plus)	1960
gx-plus (Gx-Plus)	1960
ignore	1961
include-ipv6 (Gx-Plus)	1962
interface (Static MAC Bypass)	1963
interface (VoIP)	1964
interface-description-format	1965
juniper-dsl-attributes	1966
max-outstanding-requests (Gx-Plus)	1967
nas-identifier	1967
nas-port-extended-format (Access Profile)	1968
nas-port-id-delimiter (Subscriber Management)	1969
nas-port-id-format (Subscriber Management)	1970
nas-port-type (Subscriber Management)	1971
options (Access Profile)	1973
partition (Gx-Plus)	1974
port	1975
provisioning-order	1976
radius (Access Profile)	1977
radius (System)	1979
radius-options (Protocols 802.1X)	1980
radius-options (Access)	1981
radius-server (System)	1981
retry	1982
revert-interval	1983
routing-instance	1983
secret	1984
send-acct-status-on-config-change (Access Profile)	1984
server (RADIUS Accounting)	1985
service (Service Accounting)	1985
source-address	1986
timeout (RADIUS)	1987

	vlan (VoIP) .....	1988
	vlan-assignment .....	1989
	vlan-nas-port-stacked-format .....	1990
	voip .....	1990
	wait-for-acct-on-ack (Access Profile) .....	1991
<b>Chapter 29</b>	<b>Administration .....</b>	<b>1993</b>
	Routine Monitoring .....	1993
	Monitoring 802.1X Authentication .....	1993
	Verifying 802.1X Authentication .....	1994
	Operational Commands .....	1995
	clear captive-portal .....	1997
	clear dot1x .....	1999
	clear lldp neighbors .....	2001
	clear lldp statistics .....	2002
	show captive-portal authentication-failed-users .....	2003
	show captive-portal firewall .....	2004
	show captive-portal interface .....	2006
	show dot1x .....	2009
	show dot1x authentication-failed-users .....	2014
	show dot1x firewall .....	2015
	show dot1x static-mac-address .....	2016
	show lldp .....	2018
	show lldp local-information .....	2023
	show lldp neighbors .....	2025
	show lldp remote-global-statistics .....	2031
	show lldp statistics .....	2033
	show network-access aaa statistics accounting .....	2035
	show network-access aaa statistics authentication .....	2036
	show network-access aaa statistics dynamic-requests .....	2038
<b>Part 10</b>	<b>Class of Service</b>	
<b>Chapter 30</b>	<b>Overview .....</b>	<b>2041</b>
	CoS Overview .....	2041
	Junos OS CoS for EX Series Switches Overview .....	2042
	How Junos OS CoS Works .....	2042
	Default CoS Behavior on EX Series Switches .....	2043
	Understanding Junos OS CoS Components for EX Series Switches .....	2043
	Code-Point Aliases .....	2044
	Policers .....	2044
	Classifiers .....	2044
	Forwarding Classes .....	2044
	Tail Drop Profiles .....	2045
	Schedulers .....	2045
	Rewrite Rules .....	2045
	Understanding CoS Code-Point Aliases .....	2046
	Default Code-Point Aliases .....	2046

Understanding CoS Classifiers . . . . .	2049
Behavior Aggregate Classifiers . . . . .	2049
Multifield Classifiers . . . . .	2051
Understanding CoS Forwarding Classes . . . . .	2052
Default Forwarding Classes . . . . .	2053
Understanding CoS Congestion Management . . . . .	2055
Weighted Tail Drop Congestion Management . . . . .	2055
Weighted Random Early Detection Congestion Management . . . . .	2056
Understanding CoS Schedulers . . . . .	2060
Default Schedulers . . . . .	2060
Excess Rate . . . . .	2061
Transmission Rate . . . . .	2061
Scheduler Buffer Size . . . . .	2061
Priority Scheduling . . . . .	2062
Scheduler Drop-Profile Maps . . . . .	2063
Scheduler Maps . . . . .	2063
Understanding CoS Two-Color Marking . . . . .	2067
Understanding CoS Rewrite Rules . . . . .	2067
How Rewrite Rules Work . . . . .	2067
Default Rewrite Rule . . . . .	2068
Understanding Port Shaping and Queue Shaping for CoS on EX Series Switches . . . . .	2069
Port Shaping . . . . .	2069
Queue Shaping . . . . .	2070
Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches . . . . .	2070
Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports . . . . .	2071
Oversubscribed Ports on Line Cards . . . . .	2071
EX8200 Line Cards That Include Oversubscribed Ports . . . . .	2072
Ingress Queueing . . . . .	2072
Egress Queues . . . . .	2073
<b>Chapter 31 Configuration . . . . .</b>	<b>2075</b>
Configuration Examples . . . . .	2075
Example: Configuring CoS on EX Series Switches . . . . .	2075
Configuration Tasks . . . . .	2099
Configuring CoS (J-Web Procedure) . . . . .	2100
Defining CoS Code-Point Aliases (CLI Procedure) . . . . .	2101
Defining CoS Code-Point Aliases (J-Web Procedure) . . . . .	2101
Defining CoS Classifiers (CLI Procedure) . . . . .	2103
Defining CoS Classifiers (J-Web Procedure) . . . . .	2105
Defining CoS Forwarding Classes (CLI Procedure) . . . . .	2107
Defining CoS Forwarding Classes (J-Web Procedure) . . . . .	2107
Defining CoS Schedulers and Scheduler Maps (CLI Procedure) . . . . .	2109
Configuring a Scheduler and a Scheduler Map . . . . .	2109
Assigning a Scheduler Map to Interfaces . . . . .	2110
Assigning Scheduler Maps to Interfaces on EX8200 Line Cards That Include Oversubscribed Ports . . . . .	2110

Defining CoS Schedulers (J-Web Procedure) . . . . .	2111
Defining CoS Scheduler Maps (J-Web Procedure) . . . . .	2114
Configuring CoS Congestion Management (CLI Procedure) . . . . .	2115
Configuring a Weighted Tail Drop Profile . . . . .	2116
Configuring a Weighted Random Early Detection Drop Profile . . . . .	2116
Defining CoS Drop Profiles (J-Web Procedure) . . . . .	2118
Defining CoS Rewrite Rules (CLI Procedure) . . . . .	2120
Defining CoS Rewrite Rules (J-Web Procedure) . . . . .	2121
Assigning CoS Components to Interfaces (CLI Procedure) . . . . .	2123
Assigning CoS Components to Interfaces (J-Web Procedure) . . . . .	2123
Configuring Junos OS EZQoS for CoS (CLI Procedure) . . . . .	2125
Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) . . . . .	2126
Configuration Statements . . . . .	2127
[edit class-of-service] Configuration Statement Hierarchy on EX Series Switches . . . . .	2128
Supported Statements in the [edit class-of-service] Hierarchy Level . . . . .	2128
Unsupported Statements in the [edit class-of-service] Hierarchy Level . . . . .	2130
broadcast . . . . .	2131
buffer-size . . . . .	2132
classifiers . . . . .	2133
code-point-aliases . . . . .	2134
code-points . . . . .	2134
drop-profile-map . . . . .	2135
dscp . . . . .	2136
dscp-ipv6 . . . . .	2137
ethernet (CoS for Multidestination Traffic) . . . . .	2138
excess-rate (Schedulers) . . . . .	2138
family . . . . .	2139
forwarding-class . . . . .	2140
forwarding-classes . . . . .	2141
ieee-802.1 . . . . .	2142
import . . . . .	2143
inet-precedence . . . . .	2144
interfaces . . . . .	2145
loss-priority (Classifiers and Rewrite Rules) . . . . .	2146
policing . . . . .	2147
priority (Schedulers) . . . . .	2148
protocol (Drop Profiles) . . . . .	2148
rewrite-rules . . . . .	2149
scheduler-map . . . . .	2150
scheduler-maps . . . . .	2151
schedulers (CoS) . . . . .	2152
shaping-rate . . . . .	2153
transmit-rate (EX Series Switches) . . . . .	2154
unit . . . . .	2155



<b>Chapter 32</b>	<b>Administration</b> . . . . .	<b>2157</b>
	Routine Monitoring . . . . .	2157
	Monitoring CoS Classifiers . . . . .	2157
	Monitoring CoS Forwarding Classes . . . . .	2158
	Monitoring Interfaces That Have CoS Components . . . . .	2160
	Monitoring CoS Rewrite Rules . . . . .	2161
	Monitoring CoS Scheduler Maps . . . . .	2162
	Monitoring CoS Value Aliases . . . . .	2164
	Monitoring CoS Drop Profiles . . . . .	2164
	Operational Commands . . . . .	2165
	show class-of-service . . . . .	2167
	show class-of-service classifier . . . . .	2172
	show class-of-service code-point-aliases . . . . .	2174
	show class-of-service drop-profile . . . . .	2176
	show class-of-service forwarding-class . . . . .	2179
	show pfe statistics traffic cpu . . . . .	2181
	show pfe statistics traffic egress-queues . . . . .	2185
	show pfe statistics traffic multicast . . . . .	2187
<b>Chapter 33</b>	<b>Troubleshooting Procedures</b> . . . . .	<b>2191</b>
	Troubleshooting CoS Schedulers on a 40-port SFP+ Line Card in an EX8200 Switch . . . . .	2191
	Troubleshooting a CoS Classifier Configuration for a TCAM Space Error . . . . .	2192
<b>Part 11</b>	<b>Device Security</b>	
<b>Chapter 34</b>	<b>Overview</b> . . . . .	<b>2197</b>
	Storm Control Overview . . . . .	2197
	Understanding Storm Control on Switching Devices . . . . .	2197
	Unknown Unicast Forwarding Overview . . . . .	2199
	Understanding Unknown Unicast Forwarding . . . . .	2199
<b>Chapter 35</b>	<b>Configuration</b> . . . . .	<b>2201</b>
	Configuration Examples . . . . .	2201
	Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches . . . . .	2201
	Configuration Tasks . . . . .	2204
	Configuring Unknown Unicast Forwarding (CLI Procedure) . . . . .	2205
	Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) . . . . .	2206
	Configuring or Disabling Storm Control (CLI Procedure) . . . . .	2207
	Configuring Storm Control . . . . .	2208
	Disabling Storm Control on Broadcast Traffic . . . . .	2208
	Disabling Storm Control on All Multicast Traffic . . . . .	2209
	Disabling Storm Control on Registered Multicast Traffic . . . . .	2209
	Disabling Storm Control on Unregistered Multicast Traffic . . . . .	2209
	Disabling Storm Control on Unknown Unicast Traffic . . . . .	2210
	Disabling Storm Control on Multiple Types of Traffic . . . . .	2210

	Configuration Statements . . . . .	2211
	[edit forwarding-options storm-control-profiles] Configuration Statement	
	Hierarchy for EX Series Switches . . . . .	2211
	Supported Statements in the [edit forwarding-options	
	storm-control-profiles] Hierarchy Level . . . . .	2212
	Unsupported Statements in the [edit forwarding-options	
	storm-control-profiles] Hierarchy Level . . . . .	2212
	[edit switch-options] Configuration Statement Hierarchy on EX Series	
	Switches . . . . .	2212
	Supported Statements in the [edit switch-options] Hierarchy Level . .	2213
	Unsupported Statements in the [edit switch-options] Hierarchy	
	Level . . . . .	2213
	action-shutdown . . . . .	2215
	bandwidth-level . . . . .	2217
	bandwidth-percentage . . . . .	2218
	interface (Unknown Unicast Forwarding) . . . . .	2219
	no-broadcast . . . . .	2220
	no-multicast . . . . .	2222
	no-registered-multicast . . . . .	2224
	no-unknown-unicast . . . . .	2225
	no-unregistered-multicast . . . . .	2227
	recovery-timeout . . . . .	2228
	storm-control . . . . .	2229
	storm-control-profiles . . . . .	2230
	unknown-unicast-forwarding . . . . .	2231
<b>Chapter 36</b>	<b>Administration . . . . .</b>	<b>2233</b>
	Routine Monitoring . . . . .	2233
	Verifying That Unknown Unicast Packets Are Forwarded to a Single	
	Interface . . . . .	2233
	Operational Commands . . . . .	2234
	clear ethernet-switching recovery-timeout . . . . .	2235
	clear ethernet-switching table . . . . .	2236
	show ethernet-switching table . . . . .	2238
<b>Part 12</b>	<b>Ethernet Switching</b>	
<b>Chapter 37</b>	<b>Overview . . . . .</b>	<b>2245</b>
	Bridging and VLANs . . . . .	2245
	Understanding Bridging and VLANs on EX Series Switches . . . . .	2245
	History of VLANs . . . . .	2246
	How Bridging of VLAN Traffic Works . . . . .	2246
	Packets Are Either Tagged or Untagged . . . . .	2247
	Switch Interface Modes—Access, Trunk, or Tagged Access . . . . .	2248
	Additional Advantages of Using VLANs . . . . .	2250
	Maximum VLANs and VLAN Members Per Switch . . . . .	2250
	A Default VLAN Is Configured on Most Switches . . . . .	2251
	Assigning Traffic to VLANs . . . . .	2252
	Forwarding VLAN Traffic . . . . .	2252

VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces . . . . .	2252
Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches . . . . .	2254
When Should I Use an IRB Interface or RVI? . . . . .	2255
How Does an IRB Interface or RVI Work? . . . . .	2255
Creating an IRB Interface or RVI . . . . .	2256
Viewing IRB Interface and RVI Statistics . . . . .	2257
IRB Interfaces and RVI Functions and Other Technologies . . . . .	2257
Understanding Private VLANs . . . . .	2258
Typical Structure and Primary Application of PVLANS . . . . .	2259
Routing Between Isolated and Community VLANs . . . . .	2261
PVLANS Use 802.1Q Tags to Identify Packets . . . . .	2261
PVLANS Use IP Addresses Efficiently . . . . .	2262
PVLAN Port Types and Forwarding Rules . . . . .	2262
Guidelines and Restrictions for PVLANS . . . . .	2262
Understanding Virtual Routing Instances on EX Series Switches . . . . .	2263
Understanding Multiple VLAN Registration Protocol (MVRP) on EX Series Switches . . . . .	2264
How MVRP Updates, Creates, and Deletes VLANs on the Switches . . . . .	2264
MVRP Is Disabled by Default on the Switches . . . . .	2264
MRP Timers Control MVRP Updates . . . . .	2265
MVRP Uses MRP Messages to Transmit Switch and VLAN States . . . . .	2265
Compatibility Issues with Junos OS Releases of MVRP . . . . .	2266
Understanding MAC Address Aging . . . . .	2267
Q-in-Q Tunneling . . . . .	2269
Understanding Q-in-Q Tunneling on EX Series Switches . . . . .	2269
How Q-in-Q Tunneling Works . . . . .	2269
Sending and Receiving Untagged Packets . . . . .	2270
Disabling MAC Address Learning . . . . .	2270
Mapping C-VLANs to S-VLANs . . . . .	2271
Limitations for Q-in-Q Tunneling . . . . .	2272
Layer 2 Protocol Tunneling . . . . .	2272
Understanding Layer 2 Protocol Tunneling on EX Series Switches . . . . .	2272
Layer 2 Protocols Supported by L2PT on EX Series Switches . . . . .	2273
How L2PT Works . . . . .	2274
L2PT and Q-in-Q Tunneling on EX Series Switches . . . . .	2275
Redundant Trunk Groups . . . . .	2275
Understanding Redundant Trunk Links . . . . .	2276
Proxy ARP . . . . .	2278
Understanding Proxy ARP on EX Series Switches . . . . .	2278
What Is ARP? . . . . .	2278
Proxy ARP Overview . . . . .	2278
Best Practices for Proxy ARP on EX Series Switches . . . . .	2279
<b>Chapter 38 Configuration . . . . .</b>	<b>2281</b>
Configuration Examples . . . . .	2281
Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch . . . . .	2281
Example: Connecting Access Switches to a Distribution Switch . . . . .	2291

Example: Configuring a Private VLAN on a Single Switch . . . . .	2302
Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches . . . . .	2306
Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches . . . . .	2310
Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches . .	2322
Example: Configuring Redundant Trunk Links for Faster Recovery . . . . .	2326
Example: Configuring Proxy ARP on an EX Series Switch . . . . .	2331
Configuration Tasks . . . . .	2333
Configuring VLANs for EX Series Switches (J-Web Procedure) . . . . .	2334
Configuring VLANs for EX Series Switches (CLI Procedure) . . . . .	2337
Why Create a VLAN? . . . . .	2337
Creating a VLAN Using the Minimum Procedure . . . . .	2337
Creating a VLAN Using All of the Options . . . . .	2338
Configuration Guidelines for VLANs . . . . .	2339
Configuring Integrated Routing and Bridging Interfaces (CLI Procedure) . .	2340
Configuring MAC Table Aging (CLI Procedure) . . . . .	2341
Configuring the Native VLAN Identifier (CLI Procedure) . . . . .	2342
Creating a Private VLAN on a Single Switch (CLI Procedure) . . . . .	2343
Creating a Private VLAN Spanning Multiple Switches (CLI Procedure) . .	2345
Configuring Virtual Routing Instances (CLI Procedure) . . . . .	2347
Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure ) . . . . .	2348
Enabling MVRP . . . . .	2348
Disabling MVRP . . . . .	2348
Disabling Dynamic VLANs . . . . .	2348
Configuring Timer Values . . . . .	2349
Configuring MVRP Registration Mode . . . . .	2350
Using MVRP in a Mixed-Release Network . . . . .	2350
Configuring Q-in-Q Tunneling (CLI Procedure) . . . . .	2351
Configuring All-in-One Bundling . . . . .	2352
Configuring Many-to-Many Bundling . . . . .	2353
Configuring a Specific Interface Mapping with VLAN Rewrite Option . .	2355
Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure) . . . . .	2357
Configuring Redundant Trunk Groups (J-Web Procedure) . . . . .	2358
Configuring Proxy ARP (CLI Procedure) . . . . .	2360
Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure) . . . . .	2361
Configuring MAC Limiting (CLI Procedure) . . . . .	2361
Limiting the Number of MAC Addresses Learned by an Interface . . .	2362
Limiting the Number of MAC Addresses Learned by a VLAN . . . . .	2362
Configuring MAC Notification (CLI Procedure) . . . . .	2363
Enabling MAC Notification . . . . .	2363
Disabling MAC Notification . . . . .	2364
Setting the MAC Notification Interval . . . . .	2364

Configuration Statements .....	2364
[edit interfaces] Configuration Statement Hierarchy on EX Series	
Switches .....	2366
[edit switch-options] Configuration Statement Hierarchy on EX Series	
Switches .....	2367
Supported Statements in the [edit switch-options] Hierarchy Level ..	2367
Unsupported Statements in the [edit switch-options] Hierarchy	
Level .....	2368
[edit protocols] Configuration Statement Hierarchy on EX4300	
Switches .....	2368
[edit vlans] Configuration Statement Hierarchy on EX Series Switches ..	2370
Supported Statements in the [edit vlans] Hierarchy Level .....	2370
Unsupported Statements in the [edit vlans] Hierarchy Level .....	2372
address .....	2373
aggregated-ether-options .....	2374
community-vlan .....	2375
description (Interfaces) .....	2376
description (VLANs) .....	2377
encapsulation (Physical Interface) .....	2378
ether-options .....	2383
ethernet-switch-profile .....	2384
filter (VLANs) .....	2386
flexible-vlan-tagging .....	2387
global-mac-table-aging-time .....	2388
global-no-mac-learning .....	2388
input-vlan-map .....	2389
inter-switch-link .....	2390
interface (MVRP) .....	2391
interface (VLANs) .....	2392
interface-mac-limit .....	2393
interface-mode .....	2395
isolated-vlan .....	2396
join-timer (MVRP) .....	2397
l3-interface (VLANs) .....	2398
layer2-control .....	2399
leaveall-timer (MVRP) .....	2401
leave-timer (MVRP) .....	2402
mac (Static MAC-Based VLANs) .....	2403
mac-notification .....	2403
mac-rewrite .....	2404
mac-table-size .....	2405
members .....	2407
mvrp .....	2409
native-vlan-id .....	2410
no-attribute-length-in-pdu .....	2411
no-dynamic-vlan .....	2412
no-gratuitous-arp-request .....	2413
no-mac-learning .....	2414
output-vlan-map .....	2416

	notification-interval . . . . .	2417
	packet-action . . . . .	2418
	pop . . . . .	2420
	private-vlan . . . . .	2421
	proxy-arp . . . . .	2422
	push . . . . .	2423
	redundant-trunk-group . . . . .	2424
	registration . . . . .	2425
	swap . . . . .	2426
	tag-protocol-id (TPIDs Expected to Be Sent or Received) . . . . .	2427
	vlan (802.1Q Tagging) . . . . .	2428
	vlan-id (802.1Q Tagging) . . . . .	2429
	vlan-id (VLAN Tagging and Layer 3 Subinterfaces) . . . . .	2430
	vlan-id-list . . . . .	2431
	vlangs . . . . .	2433
<b>Chapter 39</b>	<b>Administration . . . . .</b>	<b>2435</b>
	Routine Monitoring . . . . .	2435
	Verifying That Virtual Routing Instances Are Working . . . . .	2435
	Verifying Integrated Routing and Bridging Interface Status and Statistics . . . . .	2436
	Verifying That MVRP Is Working Correctly . . . . .	2438
	Verifying That MAC Notification Is Working Properly . . . . .	2439
	Verifying That Proxy ARP Is Working Correctly . . . . .	2439
	Operational Commands . . . . .	2440
	show ethernet-switching interface . . . . .	2442
	show ethernet-switching mac-notification . . . . .	2445
	show ethernet-switching table . . . . .	2447
	show interfaces irb . . . . .	2454
	show mac-rewrite interface . . . . .	2460
	show mvrp . . . . .	2462
	show mvrp dynamic-vlan-memberships . . . . .	2464
	show mvrp statistics . . . . .	2465
	show redundant-trunk-group . . . . .	2467
	show route forwarding-table . . . . .	2469
	show system statistics arp . . . . .	2473
	show vlangs . . . . .	2474
<b>Part 13</b>	<b>High Availability</b>	
<b>Chapter 40</b>	<b>Overview . . . . .</b>	<b>2483</b>
	High Availability Features Overview . . . . .	2483
	High Availability Features for EX4300 Switches Overview . . . . .	2483
	Redundant Routing Engines . . . . .	2483
	Virtual Chassis . . . . .	2484
	VRRP . . . . .	2484
	Graceful Routing Engine Switchover . . . . .	2485
	Link Aggregation . . . . .	2485
	Nonstop Active Routing and Nonstop Bridging . . . . .	2485
	Nonstop Active Routing Overview . . . . .	2486
	Understanding Nonstop Active Routing on EX Series Switches . . . . .	2486

	Nonstop Bridging Overview . . . . .	2487
	Understanding Nonstop Bridging on EX Series Switches . . . . .	2487
	Nonstop Software Upgrade Overview . . . . .	2488
	Understanding Nonstop Software Upgrade on EX Series Switches . . . . .	2488
	Requirements for Performing an NSSU . . . . .	2489
	How an NSSU Works . . . . .	2490
	NSSU Limitations . . . . .	2493
	NSSU and Junos OS Release Support . . . . .	2493
	Overview of NSSU Configuration and Operation . . . . .	2494
	Power Management Overview . . . . .	2495
	Understanding Power Management on EX Series Switches . . . . .	2495
	Power Priority of Line Cards . . . . .	2496
	Power Supply Redundancy . . . . .	2499
	VRRP Overview . . . . .	2501
	Understanding VRRP on EX Series Switches . . . . .	2501
	Overview of VRRP on EX Series Switches . . . . .	2501
	Examples of VRRP Topologies . . . . .	2502
<b>Chapter 41</b>	<b>Configuration . . . . .</b>	<b>2505</b>
	Configuration Tasks . . . . .	2505
	Configuring Nonstop Active Routing on Switches . . . . .	2505
	Configuring Nonstop Bridging on Switches (CLI Procedure) . . . . .	2507
	Upgrading Software on an EX3300, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and Mixed Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure) . . . . .	2508
	Preparing the Switch for Software Installation . . . . .	2509
	Upgrading the Software Using NSSU . . . . .	2510
	Configuring Power Supply Redundancy (CLI Procedure) . . . . .	2512
	Configuring VRRP for IPv6 (CLI Procedure) . . . . .	2513
	Configuration Examples . . . . .	2514
	Example: Configuring Nonstop Active Routing on Switches . . . . .	2514
	Configuration Statements . . . . .	2517
	[edit chassis] Configuration Statement Hierarchy on EX Series Switches . . . . .	2517
	Supported Statements in the [edit chassis] Hierarchy Level . . . . .	2518
	[edit interfaces] Configuration Statement Hierarchy on EX Series Switches . . . . .	2519
	[edit protocols layer2-control] Hierarchy Level . . . . .	2520
	[edit protocols vrrp] Hierarchy Level . . . . .	2521
	chassis . . . . .	2522
	synchronize . . . . .	2524
	failover (Chassis) . . . . .	2525
	fpc . . . . .	2526
	graceful-switchover . . . . .	2527
	hold-time . . . . .	2528
	inet6-advertise-interval . . . . .	2528
	keepalive-time . . . . .	2529
	n-plus-n (Power Management) . . . . .	2530
	nonstop-bridging . . . . .	2530
	nonstop-routing . . . . .	2531

	preempt . . . . .	2532
	priority . . . . .	2533
	psu . . . . .	2534
	redundancy (Graceful Switchover) . . . . .	2535
	redundancy (Power Management) . . . . .	2536
	traceoptions (Routing Options) . . . . .	2537
	vcp-no-hold-time . . . . .	2540
	virtual-inet6-address . . . . .	2541
	virtual-link-local-address . . . . .	2542
	vrrp-inet6-group . . . . .	2543
<b>Chapter 42</b>	<b>Administration . . . . .</b>	<b>2545</b>
	Verification Tasks . . . . .	2545
	Verifying Power Configuration and Use . . . . .	2545
	Operational Commands . . . . .	2547
	request system software nonstop-upgrade . . . . .	2548
	show chassis power-budget-statistics . . . . .	2558
	show vrrp . . . . .	2562
<b>Chapter 43</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>2573</b>
	Tracing Nonstop Active Routing Synchronization Events . . . . .	2573
<b>Part 14</b>	<b>Interfaces</b>	
<b>Chapter 44</b>	<b>Overview . . . . .</b>	<b>2577</b>
	Interfaces Overview . . . . .	2577
	EX Series Switches Interfaces Overview . . . . .	2577
	Network Interfaces . . . . .	2578
	Special Interfaces . . . . .	2578
	Understanding Interface Naming Conventions on EX Series Switches . . . . .	2580
	Physical Part of an Interface Name . . . . .	2580
	Logical Part of an Interface Name . . . . .	2582
	Wildcard Characters in Interface Names . . . . .	2582
	Understanding Aggregated Ethernet Interfaces and LACP . . . . .	2582
	Link Aggregation Group (LAG) . . . . .	2583
	Link Aggregation Control Protocol (LACP) . . . . .	2584
	Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop	
	ECMP Traffic . . . . .	2585
	Understanding the Hashing Algorithm . . . . .	2585
	IP (IPv4 and IPv6) . . . . .	2586
	MPLS . . . . .	2587
	MAC-in-MAC Packet Hashing . . . . .	2588
	Layer 2 Header Hashing . . . . .	2589
	Understanding How Energy Efficient Ethernet Reduces Power Consumption	
	on Interfaces . . . . .	2590
	Understanding Local Link Bias . . . . .	2590
	Understanding Layer 3 Subinterfaces . . . . .	2591
	Understanding Unicast RPF . . . . .	2592
	Unicast RPF for Switches Overview . . . . .	2593
	Unicast RPF Implementation . . . . .	2593



	When to Enable Unicast RPF . . . . .	2594
	When Not to Enable Unicast RPF . . . . .	2595
	Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches . . . . .	2595
	Understanding IP Directed Broadcast for EX Series Switches . . . . .	2596
	IP Directed Broadcast for EX Series Switches Overview . . . . .	2596
	IP Directed Broadcast Implementation for EX Series Switches . . . . .	2597
	When to Enable IP Directed Broadcast . . . . .	2597
	When Not to Enable IP Directed Broadcast . . . . .	2597
	Understanding Interface Ranges on EX Series Switches . . . . .	2598
	Understanding Multichassis Link Aggregation . . . . .	2599
	MC-LAG Configuration Guidelines and Functional Behavior . . . . .	2600
	Active-Active and Active-Standby Mode . . . . .	2602
	ICCP and ICL . . . . .	2602
	LACP . . . . .	2603
	Data Traffic Forwarding Rules . . . . .	2603
	Multichassis Link Protection . . . . .	2604
	Failure Handling During a Split-Brain State . . . . .	2604
	Load-Balancing . . . . .	2606
	Layer 2 Unicast Features Supported . . . . .	2606
	Layer 2 Multicast Features Supported . . . . .	2607
	IGMP Snooping on an Active-Active MC-LAG . . . . .	2607
	Layer 3 Unicast Features Supported . . . . .	2608
	VRRP Active-Standby Support . . . . .	2608
	MAC Address Management . . . . .	2609
	MAC Address Synchronization . . . . .	2609
	Address Resolution Protocol Active-Active MC-LAG Support Methodology . . . . .	2610
	DHCP Relay with Option 82 . . . . .	2611
	Layer 3 Multicast . . . . .	2612
	MC-LAG Upgrade Guidelines . . . . .	2613
	802.1Q VLANs Overview . . . . .	2614
<b>Chapter 45</b>	<b>Configuration . . . . .</b>	<b>2615</b>
	Configuration Tasks . . . . .	2615
	Configuring Gigabit Ethernet Interfaces (CLI Procedure) . . . . .	2616
	Configuring VLAN Options and Interface Mode . . . . .	2617
	Configuring the Link Settings . . . . .	2617
	Configuring the IP Options . . . . .	2618
	Configuring Gigabit Ethernet Interfaces (J-Web Procedure) . . . . .	2619
	Port Role Configuration with the J-Web Interface (with CLI References) . . . . .	2625
	Adding a Logical Unit Description to the Configuration . . . . .	2629
	Disabling a Physical Interface . . . . .	2630
	Example: Disabling a Physical Interface . . . . .	2631
	Disabling a Logical Interface . . . . .	2631
	Configuring Flow Control . . . . .	2632
	Configuring the Interface Address . . . . .	2632
	Configuring Interface IPv4 Addresses . . . . .	2634
	Configuring Interface IPv6 Addresses . . . . .	2637

Configuring the Interface Bandwidth .....	2637
Configuring the Media MTU .....	2638
Media MTU Overview .....	2639
How to Configure the Media MTU .....	2640
Encapsulation Overhead by Encapsulation Type .....	2641
Media MTU Sizes by Interface Type for M5 and M7i Routers with CFEB, M10 and M10i Routers with CFEB, and M20 and M40 Routers ..	2642
Media MTU Sizes by Interface Type for M40e Routers .....	2642
Media MTU Sizes by Interface Type for M160 Routers .....	2644
Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers .....	2644
Media MTU Sizes by Interface Type for MX Series Routers .....	2645
Media MTU Sizes by Interface Type for T320 Routers .....	2646
Media MTU Sizes by Interface Type for T640 Platforms .....	2646
Media MTU Sizes by Interface Type for J2300 Platforms .....	2647
Media MTU Sizes by Interface Type for J4300 and J6300 Platforms ..	2647
Media MTU Sizes by Interface Type for J4350 and J6350 Platforms ..	2648
Media MTU Sizes by Interface Type for EX Series Switches and ACX Series Routers .....	2650
Media MTU Sizes by Interface Type for PTX Series Packet Transport Routers .....	2650
Setting the Protocol MTU .....	2651
Interface Ranges .....	2651
Configuring Interface Ranges .....	2652
Expanding Interface Range Member and Member Range Statements .....	2655
Configuration Inheritance for Member Interfaces .....	2656
Member Interfaces Inheriting Configuration from Configuration Groups .....	2657
Interfaces Inheriting Common Configuration .....	2658
Configuring Inheritance Range Priorities .....	2658
Configuration Expansion Where Interface Range Is Used .....	2659
Configuring Accounting for the Physical Interface .....	2660
Applying an Accounting Profile to the Physical Interface .....	2660
Configuring Accounting for the Logical Interface .....	2661
Applying an Accounting Profile to the Logical Interface .....	2661
Configuring Ethernet Loopback Capability .....	2662
Configuring Gratuitous ARP .....	2663
Configuring Static ARP Table Entries .....	2664
Example: Configuring Static ARP Table Entries .....	2665
Disabling the Transmission of Redirect Messages on an Interface .....	2665
Configuring Restricted and Unrestricted Proxy ARP .....	2665
Enabling or Disabling SNMP Notifications on Logical Interfaces .....	2666
Configuring Aggregated Ethernet Links (CLI Procedure) .....	2667
Configuring Aggregated Ethernet Interfaces (J-Web Procedure) .....	2668
Configuring Aggregated Ethernet LACP (CLI Procedure) .....	2671

Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) . . . . .	2672
Configuring LACP Link Protection for a Single Link at the Global Level . . . . .	2674
Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level . . . . .	2674
Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface . . . . .	2675
Configuring Aggregated Ethernet Link Protection . . . . .	2676
Configuring Link Protection for Aggregated Ethernet Interfaces . . . . .	2677
Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces . . . . .	2677
Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link . . . . .	2677
Disabling Link Protection for Aggregated Ethernet Interfaces . . . . .	2677
Configuring Aggregated Ethernet Link Speed . . . . .	2678
Configuring Aggregated Ethernet Minimum Links . . . . .	2679
Configuring Multichassis Link Aggregation . . . . .	2681
Configuring Energy Efficient Ethernet on Interfaces (CLI Procedure) . . . . .	2684
Enabling EEE on an EEE-Capable Base-T Copper Ethernet Port . . . . .	2685
Disabling EEE on a Base-T Copper Ethernet Port . . . . .	2685
Configuring Local Link Bias (CLI Procedure) . . . . .	2686
Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) . . . . .	2686
Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing . . . . .	2687
Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing . . . . .	2687
Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing . . . . .	2687
Configuring Tagged Aggregated Ethernet Interfaces . . . . .	2688
Configuring a Layer 3 Subinterface (CLI Procedure) . . . . .	2689
Configuring Unicast RPF (CLI Procedure) . . . . .	2689
Disabling Unicast RPF (CLI Procedure) . . . . .	2691
Configuring IP Directed Broadcast (CLI Procedure) . . . . .	2691
Tracing Operations of an Individual Router or Switch Interface . . . . .	2692
Tracing Operations of the Interface Process . . . . .	2693
Configuration Statements . . . . .	2694
[edit chassis] Configuration Statement Hierarchy on EX Series Switches . . . . .	2696
Supported Statements in the [edit chassis] Hierarchy Level . . . . .	2696
[edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches . . . . .	2698
Supported Subhierarchies in the [edit forwarding-options] Hierarchy Level . . . . .	2698
Unsupported Subhierarchies in the [edit forwarding-options] Hierarchy Level . . . . .	2698
[edit interfaces] Configuration Statement Hierarchy on EX Series Switches . . . . .	2699

[edit interfaces ae] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	2700
Supported Statements in the [edit interfaces ae] Hierarchy Level . . .	2700
Unsupported Statements in the [edit interfaces ae] Hierarchy Level . .	2704
[edit interfaces et] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	2705
Supported Statements in the [edit interfaces et] Hierarchy Level . . .	2706
Unsupported Statements in the [edit interfaces et] Hierarchy Level . .	2709
[edit interfaces ge] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	2711
Supported Statements in the [edit interfaces ge] Hierarchy Level . . .	2711
Unsupported Statements in the [edit interfaces ge] Hierarchy Level . .	2715
[edit interfaces interface-range] Configuration Statement Hierarchy on EX	
Series Switches . . . . .	2716
Supported Statements in the [edit interfaces interface-range] Hierarchy	
Level . . . . .	2717
Unsupported Statements in the [edit interfaces interface-range]	
Hierarchy Level . . . . .	2720
[edit interfaces irb] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	2725
Supported Statements in the [edit interfaces irb] Hierarchy Level . . .	2725
Unsupported Statements in the [edit interfaces irb] Hierarchy Level . .	2728
[edit interfaces lo] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	2728
Supported Statements in the [edit interfaces lo] Hierarchy Level . . .	2729
Unsupported Statements in the [edit interfaces lo] Hierarchy Level . .	2731
[edit interfaces me] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	2731
Supported Statements in the [edit interfaces me] Hierarchy Level . . .	2732
Unsupported Statements in the [edit interfaces me] Hierarchy	
Level . . . . .	2734
[edit interfaces vme] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	2735
Supported Statements in the [edit interfaces vme] Hierarchy Level . .	2735
Unsupported Statements in the [edit interfaces vme] Hierarchy	
Level . . . . .	2738
[edit interfaces xe] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	2738
Supported Statements in the [edit interfaces xe] Hierarchy Level . . .	2738
Unsupported Statements in the [edit interfaces xe] Hierarchy Level . .	2742
[edit protocols lACP] Configuration Statement Hierarchy on EX Series	
Switches . . . . .	2743
Supported Statements in the [edit protocols lACP] Hierarchy Level . .	2744
Unsupported Statements in the [edit protocols lACP] Hierarchy	
Level . . . . .	2744
802.3ad . . . . .	2745
accounting-profile . . . . .	2746
address . . . . .	2747
aggregated-devices . . . . .	2749

aggregated-ether-options	2750
arp (Interfaces)	2752
auto-negotiation	2753
backup-liveness-detection	2754
backup-peer-ip	2755
bandwidth (Interfaces)	2756
broadcast	2757
chassis	2758
description (Interfaces)	2760
device-count	2761
disable (Interface)	2762
enhanced-hash-key	2764
ether-options	2766
ethernet (Aggregated Devices)	2767
eui-64	2767
family	2768
filter	2774
flow-control	2775
force-up	2776
gratuitous-arp-reply	2776
hash-mode	2777
hold-time (Physical Interface)	2779
iccp	2781
ieee-802-3az-eee	2782
inet (enhanced-hash-key)	2783
inet6 (enhanced-hash-key)	2785
interface (Multichassis Protection)	2786
interface-mode	2787
interface-range	2789
lACP (Aggregated Ethernet)	2791
lACP (802.3ad)	2793
layer2 (enhanced-hash-key)	2794
link-mode	2796
link-protection	2798
link-speed (Aggregated Ethernet)	2800
liveness-detection	2801
local-bias	2802
local-ip-addr (ICCP)	2802
loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)	2803
mc-ae	2804
mc-ae-id	2806
member (Interface Ranges)	2807
member-range	2808
members	2809
minimum-interval (Liveness Detection)	2811
minimum-receive-interval (Liveness Detection)	2811
mtu	2812
multi-chassis	2814
multi-chassis-protection	2815

	native-vlan-id .....	2816
	no-gratuitous-arp-request .....	2817
	no-redirects .....	2817
	peer (ICCP) .....	2818
	periodic .....	2819
	preferred .....	2820
	primary (Address on Interface) .....	2821
	proxy-arp .....	2822
	rpf-check .....	2823
	session-establishment-hold-time .....	2824
	speed (Ethernet) .....	2825
	traceoptions (Individual Interfaces) .....	2826
	traceoptions (Interface Process) .....	2828
	transmit-interval (Liveness Detection) .....	2829
	traps .....	2830
	unit .....	2831
	vlan (802.1Q Tagging) .....	2832
	vlan-id (VLAN Tagging and Layer 3 Subinterfaces) .....	2833
	vlan-tagging .....	2834
<b>Chapter 46</b>	<b>Administration .....</b>	<b>2835</b>
	Routine Monitoring .....	2835
	Monitoring Interface Status and Traffic .....	2835
	Verifying the Status of a LAG Interface .....	2837
	Verifying That EEE Is Saving Energy on Configured Ports .....	2837
	Verifying That LACP Is Configured Correctly and Bundle Members Are	
	Exchanging LACP Protocol Packets .....	2839
	Verifying the LACP Setup .....	2839
	Verifying That LACP Packets Are Being Exchanged .....	2840
	Verifying That Layer 3 Subinterfaces Are Working .....	2840
	Verifying Unicast RPF Status .....	2841
	Verifying IP Directed Broadcast Status .....	2843
	Operational Commands .....	2844
	monitor interface .....	2845
	request diagnostics tdr .....	2854
	show diagnostics tdr .....	2856
	show forwarding-options enhanced-hash-key .....	2861
	show interfaces diagnostics optics .....	2864
	show interfaces ge- .....	2878
	show interfaces irb .....	2890
	show interfaces mc-ae .....	2896
	show interfaces me0 .....	2898
	show interfaces queue .....	2905
	show interfaces xe- .....	2911
	show lacp interfaces .....	2925
	test interface restart-auto-negotiation .....	2930

<b>Chapter 47</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>2931</b>
	Troubleshooting an Aggregated Ethernet Interface . . . . .	2931
	Show Interfaces Command Shows the LAG is Down . . . . .	2931
	Logical Interface Statistics Do Not Reflect All Traffic . . . . .	2931
	IPv6 Interface Traffic Statistics Are Not Supported . . . . .	2932
	SNMP Counters ifHCInBroadcastPkts and ifInBroadcastPkts Are Always 0 . . . . .	2932
	Troubleshooting Interface Configuration and Cable Faults . . . . .	2932
	Interface Configuration or Connectivity Is Not Working . . . . .	2932
	Troubleshooting Unicast RPF . . . . .	2933
	Legitimate Packets Are Discarded . . . . .	2933
	Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure) . . . . .	2934
<b>Part 15</b>	<b>BGP</b>	
<b>Chapter 48</b>	<b>Overview . . . . .</b>	<b>2939</b>
	Layer 3 Protocols . . . . .	2939
	Layer 3 Protocols Supported on EX Series Switches . . . . .	2939
	Layer 3 Protocols Not Supported on EX Series Switches . . . . .	2940
<b>Chapter 49</b>	<b>Configuration . . . . .</b>	<b>2943</b>
	Configuration Tasks . . . . .	2943
	Configuring BGP Sessions (J-Web Procedure) . . . . .	2943
	Configuration Statements . . . . .	2947
	accept-remote-nexthop . . . . .	2949
	advertise-external . . . . .	2950
	advertise-inactive . . . . .	2952
	advertise-peer-as . . . . .	2953
	aggregate-label . . . . .	2954
	allow . . . . .	2955
	as-override . . . . .	2956
	authentication-algorithm . . . . .	2957
	authentication-key (Protocols BGP and BMP) . . . . .	2958
	authentication-key-chain (Protocols BGP and BMP) . . . . .	2959
	autonomous-system . . . . .	2960
	bfd-liveness-detection (Protocols BGP) . . . . .	2963
	bgp . . . . .	2967
	bgp-orf-cisco-mode . . . . .	2968
	cluster . . . . .	2970
	damping (Protocols BGP) . . . . .	2972
	description (Protocols BGP) . . . . .	2974
	disable (Protocols BGP) . . . . .	2975
	explicit-null (Protocols BGP) . . . . .	2976
	export (Protocols BGP) . . . . .	2978
	family (Protocols BGP) . . . . .	2979
	graceful-restart (Protocols BGP) . . . . .	2983
	group (Protocols BGP) . . . . .	2984
	hold-time (Protocols BGP) . . . . .	2987
	idle-after-switch-over . . . . .	2989

	import (Protocols BGP) .....	2990
	include-mp-next-hop .....	2992
	keep .....	2993
	labeled-unicast (Protocols BGP) .....	2995
	local-address (Protocols BGP) .....	2997
	local-as .....	2999
	local-interface (IPv6) .....	3001
	local-preference .....	3002
	log-updown (Protocols BGP) .....	3003
	metric-out (Protocols BGP) .....	3004
	mtu-discovery .....	3006
	multihop .....	3008
	multipath (Protocols BGP) .....	3010
	neighbor (Protocols BGP) .....	3011
	no-aggregator-id .....	3014
	no-client-reflect .....	3015
	no-validate .....	3016
	out-delay .....	3018
	outbound-route-filter .....	3020
	passive (Protocols BGP) .....	3021
	peer-as (Protocols BGP) .....	3022
	preference (Protocols BGP) .....	3024
	prefix-limit .....	3025
	remove-private .....	3027
	rib-group (Protocols BGP) .....	3029
	tcp-mss (Protocols BGP) .....	3030
	traceoptions (Protocols BGP) .....	3031
	type (Protocols BGP) .....	3034
<b>Chapter 50</b>	<b>Administration .....</b>	<b>3035</b>
	Routine Monitoring .....	3035
	Monitoring BGP Routing Information .....	3035
	Operational Commands .....	3037
	clear bgp damping .....	3038
	clear bgp neighbor .....	3039
	clear bgp table .....	3041
	show bgp bmp .....	3043
	show bgp group .....	3045
	show bgp neighbor .....	3052
	show bgp summary .....	3066
	show policy damping .....	3072
<b>Part 16</b>	<b>IS-IS</b>	
<b>Chapter 51</b>	<b>Overview .....</b>	<b>3077</b>
	Layer 3 Protocols .....	3077
	Layer 3 Protocols Supported on EX Series Switches .....	3077
	Layer 3 Protocols Not Supported on EX Series Switches .....	3078



<b>Chapter 52</b>	<b>Configuration</b>	<b>3081</b>
	Configuration Statements	3081
	authentication-key (Protocols IS-IS)	3083
	authentication-type (Protocols IS-IS)	3084
	bfd-liveness-detection (Protocols IS-IS)	3085
	checksum (Protocols IS-IS)	3087
	csnp-interval	3088
	disable (Protocols IS-IS)	3089
	export (Protocols IS-IS)	3090
	external-preference (Protocols IS-IS)	3091
	graceful-restart (Protocols IS-IS)	3092
	hello-authentication-key	3093
	hello-authentication-type	3094
	hello-interval (Protocols IS-IS)	3095
	hello-padding	3096
	hold-time (Protocols IS-IS)	3098
	ignore-attached-bit	3099
	interface (Protocols IS-IS)	3100
	ipv4-multicast	3102
	ipv4-multicast-metric	3103
	ipv6-multicast	3103
	ipv6-multicast-metric	3104
	ipv6-unicast	3105
	ipv6-unicast-metric	3106
	isis	3107
	level (Global IS-IS)	3108
	link-protection (Protocols IS-IS)	3109
	loose-authentication-check	3109
	lsp-interval	3110
	lsp-lifetime	3111
	max-areas	3112
	mesh-group (Protocols IS-IS)	3113
	metric (Protocols IS-IS)	3114
	no-adjacency-holddown	3115
	no-authentication-check	3116
	no-csnp-authentication	3116
	no-eligible-backup (Protocols IS-IS)	3117
	no-hello-authentication	3117
	no-ipv4-multicast	3118
	no-ipv4-routing	3119
	no-ipv6-multicast	3120
	no-ipv6-routing	3121
	no-ipv6-unicast	3122
	no-psnp-authentication	3122
	no-unicast-topology	3123
	node-link-protection (Protocols IS-IS)	3123
	overload (Protocols IS-IS)	3124
	passive (Protocols IS-IS)	3127
	point-to-point	3128

	preference (Protocols IS-IS) . . . . .	3129
	prefix-export-limit (Protocols IS-IS) . . . . .	3130
	priority (Protocols IS-IS) . . . . .	3131
	reference-bandwidth (Protocols IS-IS) . . . . .	3132
	rib-group (Protocols IS-IS) . . . . .	3133
	spf-options (Protocols IS-IS) . . . . .	3134
	topologies (Protocols IS-IS) . . . . .	3135
	traceoptions (Protocols IS-IS) . . . . .	3136
	wide-metrics-only . . . . .	3139
<b>Chapter 53</b>	<b>Administration . . . . .</b>	<b>3141</b>
	Operational Commands . . . . .	3141
	clear isis adjacency . . . . .	3142
	clear isis database . . . . .	3144
	clear isis overload . . . . .	3146
	clear isis statistics . . . . .	3148
	show isis adjacency . . . . .	3150
	show isis authentication . . . . .	3154
	show isis backup coverage . . . . .	3156
	show isis backup label-switched-path . . . . .	3158
	show isis backup spf results . . . . .	3160
	show isis database . . . . .	3164
	show isis hostname . . . . .	3176
	show isis interface . . . . .	3178
	show isis overview . . . . .	3182
	show isis route . . . . .	3185
	show isis spf . . . . .	3189
	show isis statistics . . . . .	3194
<b>Part 17</b>	<b>OSPF</b>	
<b>Chapter 54</b>	<b>Overview . . . . .</b>	<b>3199</b>
	Layer 3 Protocols . . . . .	3199
	Layer 3 Protocols Supported on EX Series Switches . . . . .	3199
	Layer 3 Protocols Not Supported on EX Series Switches . . . . .	3200
	OSPF Overview . . . . .	3202
	Understanding IPsec Authentication for OSPF Packets on EX Series Switches . . . . .	3202
	Authentication Algorithms . . . . .	3203
	Encryption Algorithms . . . . .	3203
	IPsec Protocols . . . . .	3204
	Security Associations . . . . .	3204
	IPsec Modes . . . . .	3204

<b>Chapter 55</b>	<b>Configuration</b>	<b>3207</b>
	Configuration Tasks	3207
	Configuring an OSPF Network (J-Web Procedure)	3207
	Using IPsec to Secure OSPFv3 Networks (CLI Procedure)	3211
	Configuring Security Associations	3211
	Securing OPSFv3 Networks	3212
	Configuration Statements	3212
	area	3214
	area-range	3216
	bandwidth-based-metrics	3218
	bfd-liveness-detection (Protocols OSPF)	3220
	dead-interval	3224
	default-lsa	3225
	default-metric	3226
	disable (OSPF)	3228
	domain-id	3229
	domain-vpn-tag	3230
	export (Protocols OSPF)	3231
	external-preference (Protocols OSPF)	3232
	graceful-restart (Protocols OSPF)	3233
	hello-interval (Protocols OSPF)	3235
	ignore-lsp-metrics	3236
	import (Protocols OSPF)	3237
	inter-area-prefix-export	3238
	inter-area-prefix-import	3239
	interface (Protocols OSPF)	3240
	interface-type (Protocols OSPF)	3242
	lsp-metric-into-summary	3243
	metric (Protocols OSPF Interface)	3244
	metric-type	3246
	no-nssa-abr	3247
	no-rfc-1583	3248
	nssa	3249
	ospf	3250
	ospf3	3251
	overload (Protocols OSPF)	3252
	passive (Protocols OSPF)	3254
	preference (Protocols OSPF)	3255
	prefix-export-limit (Protocols OSPF)	3256
	priority (Protocols OSPF)	3257
	realm	3258
	reference-bandwidth (Protocols OSPF)	3259
	retransmit-interval (OSPF)	3260
	rib-group (Protocols OSPF)	3261
	route-type-community	3262
	shortcuts (Protocols OSPF)	3263
	spf-options (Protocols OSPF)	3264
	stub	3266
	summaries	3267

	traceoptions (Protocols OSPF) . . . . .	3268
	traffic-engineering (OSPF) . . . . .	3271
	transit-delay (OSPF) . . . . .	3273
	type-7 . . . . .	3274
	virtual-link . . . . .	3275
<b>Chapter 56</b>	<b>Administration . . . . .</b>	<b>3277</b>
	Routine Monitoring . . . . .	3277
	Monitoring OSPF Routing Information . . . . .	3277
	Operational Commands . . . . .	3280
	clear (ospf   ospf3) database . . . . .	3281
	clear (ospf   ospf3) io-statistics . . . . .	3284
	clear (ospf   ospf3) neighbor . . . . .	3285
	clear (ospf   ospf3) statistics . . . . .	3287
	clear (ospf   ospf3) overload . . . . .	3289
	show (ospf   ospf3) interface . . . . .	3290
	show (ospf   ospf3) io-statistics . . . . .	3296
	show (ospf   ospf3) log . . . . .	3298
	show (ospf   ospf3) neighbor . . . . .	3301
	show (ospf   ospf3) overview . . . . .	3307
	show (ospf   ospf3) route . . . . .	3312
	show (ospf   ospf3) statistics . . . . .	3318
	show ospf database . . . . .	3322
	show ospf3 database . . . . .	3330
<b>Part 18</b>	<b>RIP</b>	
<b>Chapter 57</b>	<b>Overview . . . . .</b>	<b>3343</b>
	Layer 3 Protocols . . . . .	3343
	Layer 3 Protocols Supported on EX Series Switches . . . . .	3343
	Layer 3 Protocols Not Supported on EX Series Switches . . . . .	3344
<b>Chapter 58</b>	<b>Configuration . . . . .</b>	<b>3347</b>
	Configuration Tasks . . . . .	3347
	Configuring a RIP Network (J-Web Procedure) . . . . .	3347
	Configuration Statements: RIP . . . . .	3350
	any-sender . . . . .	3351
	authentication-key (Protocols RIP) . . . . .	3352
	authentication-type (Protocols RIP) . . . . .	3353
	bfd-liveness-detection (Protocols RIP) . . . . .	3354
	check-zero . . . . .	3357
	export (Protocols RIP) . . . . .	3358
	graceful-restart (Protocols RIP) . . . . .	3359
	group (Protocols RIP) . . . . .	3360
	holddown (Protocols RIP) . . . . .	3362
	import (Protocols RIP) . . . . .	3363
	message-size . . . . .	3364
	metric-in (Protocols RIP) . . . . .	3365
	metric-out (Protocols RIP) . . . . .	3366
	neighbor (Protocols RIP) . . . . .	3367

	preference (Protocols RIP) .....	3368
	receive (Protocols RIP) .....	3369
	rib-group (Protocols RIP) .....	3370
	rip .....	3370
	route-timeout (Protocols RIP) .....	3371
	send (Protocols RIP) .....	3372
	traceoptions (Protocols RIP) .....	3373
	update-interval (Protocols RIP) .....	3376
	Configuration Statements: RIPng .....	3376
	export (Protocols RIPng) .....	3377
	graceful-restart (Protocols RIPng) .....	3378
	group (Protocols RIPng) .....	3379
	holddown (Protocols RIPng) .....	3380
	import (Protocols RIPng) .....	3381
	metric-in (Protocols RIPng) .....	3382
	metric-out (Protocols RIPng) .....	3383
	neighbor (Protocols RIPng) .....	3384
	preference (Protocols RIPng) .....	3385
	receive (Protocols RIPng) .....	3386
	ripng .....	3387
	route-timeout (Protocols RIPng) .....	3387
	send (Protocols RIPng) .....	3388
	traceoptions (Protocols RIPng) .....	3389
	update-interval (Protocols RIPng) .....	3391
<b>Chapter 59</b>	<b>Administration .....</b>	<b>3393</b>
	Routine Monitoring .....	3393
	Monitoring RIP Routing Information .....	3393
	Operational Commands: RIP .....	3394
	clear rip general-statistics .....	3395
	clear rip statistics .....	3396
	show rip general-statistics .....	3397
	show rip neighbor .....	3399
	show rip statistics .....	3401
	Operational Commands: RIPng .....	3403
	clear ripng general-statistics .....	3404
	clear ripng statistics .....	3405
	show ripng general-statistics .....	3406
	show ripng neighbor .....	3408
	show ripng statistics .....	3410
<b>Part 19</b>	<b>Routing Options</b>	
<b>Chapter 60</b>	<b>Overview .....</b>	<b>3415</b>
	Layer 3 Protocols .....	3415
	Layer 3 Protocols Supported on EX Series Switches .....	3415
	Layer 3 Protocols Not Supported on EX Series Switches .....	3416
	Routing Options Overview .....	3418
	Understanding Distributed Periodic Packet Management on EX Series Switches .....	3418

<b>Chapter 61</b>	<b>Configuration</b>	<b>3421</b>
	Configuration Tasks	3421
	Configuring Static Routing (CLI Procedure)	3422
	Configuring Static Routing (J-Web Procedure)	3422
	Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)	3424
	Disabling or Enabling Distributed Periodic Packet Management Globally	3425
	Disabling or Enabling Distributed Periodic Packet Management for LACP Packets	3425
	Configuration Statements	3425
	active	3429
	aggregate (Routing)	3430
	as-path (Routing Options)	3432
	asm-override-ssm	3434
	authentication-key-chains	3435
	autonomous-system	3436
	backup-pe-group	3438
	backups	3439
	bandwidth (Multicast Flow Map)	3440
	bfd-liveness-detection (Routing Options Static Route)	3441
	bgp-orf-cisco-mode	3445
	bmp	3447
	brief	3449
	centralized	3450
	community (Routing Options)	3451
	confederation	3453
	disable (Routing Options)	3454
	discard	3455
	export (Routing Options)	3456
	export-rib	3457
	fate-sharing	3459
	flow	3460
	flow-map	3461
	forwarding-cache (Flow Maps)	3462
	forwarding-cache (Multicast)	3463
	forwarding-table	3464
	generate	3465
	import (Routing Options)	3466
	import-policy	3467
	import-rib	3468
	indirect-next-hop	3469
	install (Routing Options)	3470
	instance-export	3471
	instance-import	3471
	interface (Multicast Static Routes)	3472
	interface (Routing Options)	3473
	interface-routes	3474
	local-address (Routing Options)	3475

martians	3476
maximum-bandwidth (Routing Options)	3477
maximum-paths	3478
maximum-prefixes	3480
med-igp-update-interval	3481
metric (Aggregate, Generated, or Static Route)	3482
multicast (Routing Options)	3483
no-qos-adjust	3484
options (Routing Options)	3485
pim-to-igmp-proxy	3486
pim-to-mld-proxy	3487
policy (Aggregate and Generated Routes)	3488
policy (Flow Maps)	3489
policy (SSM Maps)	3490
ppm	3491
ppm (Ethernet Switching)	3492
preference (Routing Options)	3493
prefix	3494
qualified-next-hop (Static Routes)	3495
readvertise	3497
redundant-sources	3498
resolution	3499
resolution-ribs	3500
resolve	3501
restart-duration (Routing Options)	3502
retain	3503
reverse-oif-mapping	3504
rpf-check-policy (Routing Options RPF)	3505
rib (General)	3506
rib (Route Resolution)	3508
rib-group (Routing Options)	3509
rib-groups	3510
route-distinguisher-id	3512
route-record	3513
router-id	3514
routing-options	3515
scope	3516
scope-policy	3517
source (Source-Specific Multicast)	3518
source-routing	3519
ssm-groups	3520
ssm-map (Routing Options Multicast)	3521
static (Routing Options)	3522
subscriber-leave-timer	3528
tag (Routing Options)	3529
threshold (Multicast Forwarding Cache)	3530
timeout (Flow Maps)	3531
timeout (Multicast)	3532
traceoptions (Routing Options)	3533

	upstream-interface . . . . .	3536
<b>Chapter 62</b>	<b>Administration . . . . .</b>	<b>3537</b>
	Routine Monitoring . . . . .	3537
	Monitoring Routing Information . . . . .	3537
	Operational Commands . . . . .	3539
	clear ipv6 neighbors . . . . .	3541
	show as-path . . . . .	3542
	show as-path domain . . . . .	3546
	show as-path summary . . . . .	3549
	show ipv6 neighbors . . . . .	3551
	show route . . . . .	3553
	show route active-path . . . . .	3559
	show route all . . . . .	3564
	show route aspath-regex . . . . .	3566
	show route best . . . . .	3568
	show route brief . . . . .	3571
	show route community . . . . .	3573
	show route community-name . . . . .	3575
	show route damping . . . . .	3577
	show route detail . . . . .	3582
	show route exact . . . . .	3599
	show route export . . . . .	3601
	show route extensive . . . . .	3604
	show route flow validation . . . . .	3621
	show route inactive-path . . . . .	3623
	show route inactive-prefix . . . . .	3626
	show route instance . . . . .	3628
	show route label . . . . .	3636
	show route label-switched-path . . . . .	3639
	show route martians . . . . .	3641
	show route next-hop . . . . .	3643
	show route no-community . . . . .	3649
	show route protocol . . . . .	3652
	show route range . . . . .	3664
	show route receive-protocol . . . . .	3668
	show route resolution . . . . .	3676
	show route snooping . . . . .	3679
	show route source-gateway . . . . .	3687
	show route summary . . . . .	3693
	show route table . . . . .	3697
	show route terse . . . . .	3712
<b>Part 20</b>	<b>Multicast</b>	
<b>Chapter 63</b>	<b>Overview . . . . .</b>	<b>3717</b>
	IGMP Snooping Overview . . . . .	3717
	IGMP Snooping on EX Series Switches Overview . . . . .	3717
	How IGMP Snooping Works . . . . .	3718
	IGMP Message Types . . . . .	3718



	How Hosts Join and Leave Multicast Groups . . . . .	3719
	Support for IGMPv3 Multicast Sources . . . . .	3719
	IGMP Snooping and Forwarding Interfaces . . . . .	3720
	General Forwarding Rules . . . . .	3720
	Examples of IGMP Snooping Multicast Forwarding . . . . .	3721
<b>Chapter 64</b>	<b>Configuration . . . . .</b>	<b>3727</b>
	Configuration Examples . . . . .	3727
	Example: Configuring IGMP Snooping . . . . .	3727
	Configuration Tasks . . . . .	3730
	Configuring IGMP Snooping (CLI Procedure) . . . . .	3730
	Enabling IGMP Snooping on VLANs . . . . .	3731
	Enabling Immediate Leave . . . . .	3731
	Configuring an Interface as a Multicast-Router Interface . . . . .	3732
	Configuring Static Group Membership on an Interface . . . . .	3733
	Changing the Timer and Counter Values . . . . .	3734
	Configuring IGMP Snooping (J-Web Procedure) . . . . .	3735
	Configuring IGMP Snooping Tracing Operations (CLI Procedure) . . . . .	3737
	Configuring Tracing Operations . . . . .	3738
	Viewing, Stopping, and Restarting Tracing Operations . . . . .	3739
	Configuration Statements . . . . .	3739
	[edit protocols] Configuration Statement Hierarchy on EX Series Switches . . . . .	3741
	[edit protocols igmp-snooping] Configuration Statement Hierarchy . . . . .	3743
	Supported Statements in the [edit protocols igmp-snooping] Hierarchy Level . . . . .	3744
	Unsupported Statements in the [edit protocols igmp-snooping] Hierarchy Level . . . . .	3744
	accounting (Protocols IGMP Interface) . . . . .	3745
	accounting (Protocols IGMP) . . . . .	3745
	address (Anycast RPs) . . . . .	3746
	address (Local RPs) . . . . .	3746
	anycast-pim . . . . .	3747
	assert-timeout . . . . .	3748
	auto-rp . . . . .	3749
	bootstrap . . . . .	3750
	bootstrap-export . . . . .	3751
	bootstrap-import . . . . .	3752
	bootstrap-priority . . . . .	3753
	data-forwarding . . . . .	3754
	dense-groups . . . . .	3755
	disable (IGMP Snooping) . . . . .	3755
	disable (Protocols IGMP) . . . . .	3756
	disable (PIM) . . . . .	3757
	dr-election-on-p2p . . . . .	3758
	dr-register-policy . . . . .	3758
	embedded-rp . . . . .	3759
	export (Protocols PIM Bootstrap) . . . . .	3760
	family (Bootstrap) . . . . .	3761

family (Local RP) .....	3762
graceful-restart (Protocols PIM) .....	3763
group (IGMP Snooping) .....	3763
group (Protocols IGMP) .....	3764
group-ranges .....	3765
hello-interval (Protocols PIM) .....	3766
hold-time (Protocols PIM) .....	3767
igmp-snooping .....	3768
immediate-leave (Protocols IGMP) .....	3769
immediate-leave (IGMP Snooping) .....	3770
import (Protocols PIM Bootstrap) .....	3771
import (Protocols PIM) .....	3772
infinity .....	3773
interface (IGMP Snooping) .....	3774
interface (Protocols PIM) .....	3775
interface (Protocols IGMP) .....	3777
join-load-balance .....	3778
local .....	3779
local-address (Protocols PIM) .....	3780
mapping-agent-election .....	3781
maximum-rps .....	3782
mode (Protocols PIM) .....	3783
multicast-router-interface (IGMP Snooping) .....	3784
neighbor-policy .....	3785
pim .....	3786
priority (PIM Interfaces) .....	3790
priority (Bootstrap) .....	3791
priority (PIM RPs) .....	3792
query-interval (Protocols IGMP) .....	3793
query-last-member-interval (Protocols IGMP) .....	3794
query-response-interval (Protocols IGMP) .....	3795
receiver .....	3796
restart-duration (Protocols PIM) .....	3797
rib-group (Protocols PIM) .....	3798
robust-count (IGMP Snooping) .....	3799
robust-count (Protocols IGMP) .....	3799
rp .....	3800
rp-register-policy .....	3802
rp-set .....	3803
source (Protocols IGMP) .....	3804
source-vlans .....	3805
spt-threshold .....	3806
ssm-map (Protocols IGMP) .....	3807
static (IGMP Snooping) .....	3807
static (Protocols PIM) .....	3808
static (Protocols IGMP) .....	3809
traceoptions (Protocols PIM) .....	3810
traceoptions (Protocols IGMP) .....	3813
traceoptions (IGMP Snooping) .....	3816

	version (Protocols IGMP) . . . . .	3818
	version (IGMP Snooping) . . . . .	3818
	version (PIM) . . . . .	3819
	vlan (IGMP Snooping) . . . . .	3820
<b>Chapter 65</b>	<b>Administration . . . . .</b>	<b>3823</b>
	Routine Monitoring . . . . .	3823
	Monitoring IGMP Snooping . . . . .	3823
	Verifying IGMP Snooping (CLI Procedure) . . . . .	3826
	Verifying IGMP Snooping Memberships . . . . .	3826
	Viewing IGMP Snooping Statistics . . . . .	3827
	Viewing IGMP Snooping Routing Information . . . . .	3828
	Operational Commands . . . . .	3828
	clear igmp membership . . . . .	3830
	clear igmp statistics . . . . .	3833
	clear igmp-snooping membership . . . . .	3835
	clear igmp-snooping statistics . . . . .	3836
	clear multicast bandwidth-admission . . . . .	3837
	clear multicast scope . . . . .	3839
	clear multicast sessions . . . . .	3840
	clear multicast statistics . . . . .	3841
	clear pim join . . . . .	3842
	clear pim register . . . . .	3844
	clear pim statistics . . . . .	3846
	mtrace . . . . .	3849
	mtrace from-source . . . . .	3852
	mtrace monitor . . . . .	3855
	mtrace to-gateway . . . . .	3857
	show igmp group . . . . .	3860
	show igmp interface . . . . .	3864
	show igmp statistics . . . . .	3868
	show igmp-snooping membership . . . . .	3871
	show igmp-snooping route . . . . .	3874
	show igmp-snooping statistics . . . . .	3876
	show multicast flow-map . . . . .	3878
	show multicast interface . . . . .	3880
	show multicast minfo . . . . .	3882
	show multicast next-hops . . . . .	3884
	show multicast pim-to-igmp-proxy . . . . .	3887
	show multicast pim-to-mld-proxy . . . . .	3889
	show multicast route . . . . .	3891
	show multicast rpf . . . . .	3898
	show multicast scope . . . . .	3902
	show multicast sessions . . . . .	3904
	show multicast usage . . . . .	3907
	show pim bootstrap . . . . .	3910
	show pim interfaces . . . . .	3912
	show pim join . . . . .	3915
	show pim neighbors . . . . .	3929

show pim rps . . . . .	3933
show pim source . . . . .	3940
show pim statistics . . . . .	3943

## Part 21

### Chapter 66

## Network Management and Monitoring

<b>Overview . . . . .</b>	<b>3959</b>
Real-Time Performance Monitoring . . . . .	3959
Understanding Real-Time Performance Monitoring on EX Series	
Switches . . . . .	3960
RPM Packet Collection . . . . .	3960
Tests and Probe Types . . . . .	3960
Hardware Timestamps . . . . .	3961
Limitations of RPM on EX Series Switches . . . . .	3963
Analyzers and Port Mirroring . . . . .	3963
Understanding Port Mirroring and Analyzers on EX4300 Switches . . . . .	3964
Port Mirroring Overview . . . . .	3964
Analyzer Overview . . . . .	3965
Port Mirroring and Analyzer Terminologies . . . . .	3965
Configuration Guidelines for Port Mirroring and Analyzers on EX4300	
Switches . . . . .	3966
sFlow Monitoring Technology . . . . .	3969
Understanding How to Use sFlow Technology for Network Monitoring on	
an EX Series Switch . . . . .	3969
Sampling Mechanism and Architecture of sFlow Technology on EX	
Series Switches . . . . .	3969
Adaptive Sampling . . . . .	3970
sFlow Agent Address Assignment . . . . .	3971
Ethernet OAM Link Fault Management and Connectivity Fault Management . . . . .	3971
Understanding Ethernet OAM Link Fault Management for an EX Series	
Switch . . . . .	3972
Understanding Ethernet OAM Connectivity Fault Management for an EX	
Series Switch . . . . .	3973
Understanding Ethernet Frame Delay Measurements on Switches . . . . .	3974
Ethernet Frame Delay Measurements . . . . .	3975
Types of Ethernet Frame Delay Measurements . . . . .	3975
Limitations . . . . .	3976
Uplink Failure Detection . . . . .	3976
Understanding Uplink Failure Detection . . . . .	3977
Uplink Failure Detection Overview . . . . .	3977
Failure Detection Pair . . . . .	3978
Network Analytics . . . . .	3978
Network Analytics Overview . . . . .	3978
Analytics Feature Overview . . . . .	3979
Network Analytics Enhancements Overview . . . . .	3980
Summary of CLI Changes . . . . .	3981
Understanding Enhanced Network Analytics Streaming Data . . . . .	3985
Google Protocol Buffer (GPB) . . . . .	3985
JavaScript Object Notation (JSON) . . . . .	3988

	Comma-separated Values (CSV) . . . . .	3988
	Tab-separated Values (TSV) . . . . .	3988
	Queue Statistics Output for JSON, CSV, and TSV . . . . .	3989
	Traffic Statistics Output for JSON, CSV, and TSV . . . . .	3989
	Understanding Enhanced Analytics Local File Output . . . . .	3991
	Prototype File for the Google Protocol Buffer Stream Format . . . . .	3993
<b>Chapter 67</b>	<b>Configuration . . . . .</b>	<b>3995</b>
	Configuration Examples . . . . .	3995
	Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches . . . . .	3995
	Example: Configuring Mirroring for Local Monitoring of Employee Resource Use on EX4300 Switches . . . . .	4000
	Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches . . . . .	4006
	Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX4300 Switches . . . . .	4016
	Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches . . . . .	4023
	Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches . . . . .	4027
	Example: Configuring Enhanced Network Analytics Features . . . . .	4029
	Configuration Tasks . . . . .	4039
	Configuring SNMP (J-Web Procedure) . . . . .	4040
	Configuring Mirroring on EX4300 Switches to Analyze Traffic (CLI Procedure) . . . . .	4043
	Configuring an Analyzer for Local Traffic Analysis . . . . .	4044
	Configuring an Analyzer for Remote Traffic Analysis . . . . .	4044
	Configuring Port Mirroring . . . . .	4045
	Configuring Port Mirroring to Analyze Traffic (J-Web Procedure) . . . . .	4047
	Configuring sFlow Technology for Network Monitoring (CLI Procedure) . . . . .	4049
	Configuring Real-Time Performance Monitoring (J-Web Procedure) . . . . .	4050
	Configuring the Interface for RPM Timestamping for Client/Server on an EX Series Switch (CLI Procedure) . . . . .	4057
	Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) . . . . .	4058
	Creating the Maintenance Domain . . . . .	4058
	Configuring the Maintenance Domain MIP Half Function . . . . .	4059
	Creating a Maintenance Association . . . . .	4059
	Configuring the Continuity Check Protocol . . . . .	4060
	Configuring a Maintenance Association End Point . . . . .	4060
	Configuring a Connectivity Fault Management Action Profile . . . . .	4061
	Configuring the Linktrace Protocol . . . . .	4061
	Configuring Ethernet OAM Link Fault Management (CLI Procedure) . . . . .	4062
	Configuring Interfaces for Uplink Failure Detection (CLI Procedure) . . . . .	4064
	Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements (CLI Procedure) . . . . .	4065
	Configuring One-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure) . . . . .	4066

Configuring an Iterator Profile on a Switch (CLI Procedure) . . . . .	4067
Triggering an Ethernet Frame Delay Measurement Session on a Switch . .	4068
Configuring Two-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure) . . . . .	4069
Configuring Queue Monitoring . . . . .	4070
Configuring Traffic Monitoring . . . . .	4072
Configuring a Local File for Network Analytics Data . . . . .	4073
Configuring a Remote Collector for Streaming Analytics Data . . . . .	4074
Configuration Statements: SNMP . . . . .	4075
[edit snmp] Configuration Statement Hierarchy on EX Series Switches . .	4078
Supported Statements in the [edit snmp] Hierarchy Level . . . . .	4078
Unsupported Statements in the [edit snmp] Hierarchy Level . . . . .	4083
address (SNMP) . . . . .	4085
address-mask . . . . .	4086
agent-address . . . . .	4086
alarm (SNMP RMON) . . . . .	4087
authorization . . . . .	4088
bucket-size . . . . .	4088
categories . . . . .	4089
client-list . . . . .	4089
client-list-name . . . . .	4090
clients . . . . .	4090
commit-delay . . . . .	4091
community (SNMP) . . . . .	4092
community (SNMP RMON) . . . . .	4093
community-name (SNMP) . . . . .	4094
contact (SNMP) . . . . .	4095
description (SNMP) . . . . .	4096
description (SNMP RMON) . . . . .	4097
destination-port . . . . .	4097
engine-id (SNMP) . . . . .	4098
event (SNMP) . . . . .	4099
falling-event-index . . . . .	4099
falling-threshold . . . . .	4100
falling-threshold . . . . .	4101
falling-threshold-interval . . . . .	4102
filter-duplicates . . . . .	4102
filter-interfaces . . . . .	4103
group (Configuring Access Privileges) . . . . .	4104
group (Associating a Security Name) . . . . .	4105
health-monitor . . . . .	4105
history . . . . .	4106
interface (SNMP RMON History) . . . . .	4107
interface (SNMP) . . . . .	4107
interval . . . . .	4108
interval (SNMP Health Monitor) . . . . .	4108
interval (SNMP RMON) . . . . .	4109
location (SNMP) . . . . .	4109
logical-system (SNMP) . . . . .	4110

message-processing-model	4111
name	4111
nonvolatile	4112
notify	4113
notify-filter (Applying to the Management Target)	4114
notify-filter (Configuring the Profile Name)	4114
notify-view	4115
oid (SNMP)	4115
oid	4116
owner	4116
parameters	4117
port (SNMP)	4117
read-view	4118
request-type	4119
rising-event-index	4120
rising-threshold (SNMP Health Monitor)	4121
rising-threshold (SNMP RMON)	4122
rmon	4122
rmon	4123
routing-instance (SNMP)	4124
routing-instance (SNMPv3)	4125
sample-type	4125
security-level (Defining Access Privileges)	4126
security-level (Generating SNMP Notifications)	4127
security-model (Access Privileges)	4128
security-model (Group)	4129
security-model (SNMP Notifications)	4130
security-name (Security Group)	4131
security-name (Community String)	4132
security-name (SNMP Notifications)	4133
security-to-group	4134
snmp	4134
snmp	4135
snmp-community	4136
source-address (SNMP)	4137
startup-alarm	4138
syslog-subtag	4138
tag (SNMPv3)	4139
tag-list	4139
target-address	4140
target-parameters	4141
targets	4142
traceoptions (SNMP)	4143
trap-group	4145
trap-options	4146
type (SNMP RMON)	4147
type (SNMPv3)	4148
v3	4149
vacm	4151

variable .....	4152
version (SNMP) .....	4152
view (Associating a MIB View with a Community) .....	4153
view (Configuring a MIB View) .....	4154
write-view .....	4155
Configuration Statements: Analyzers and Port Mirroring .....	4155
[edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches .....	4155
Supported Subhierarchies in the [edit forwarding-options] Hierarchy Level .....	4156
Unsupported Subhierarchies in the [edit forwarding-options] Hierarchy Level .....	4156
[edit forwarding-options port-mirroring] Configuration Statement Hierarchy .....	4157
[edit forwarding-options analyzer] Configuration Statement Hierarchy ..	4158
egress .....	4158
egress (Analyzer) .....	4159
ingress (vlans) .....	4159
ingress (Analyzer) .....	4160
input (Analyzer) .....	4161
interface .....	4162
no-tag .....	4163
output (Mirroring) .....	4164
vlan (Mirroring) .....	4165
Configuration Statements for Network Analytics .....	4165
address (Analytics Collector) .....	4166
analytics .....	4167
collector (Analytics) .....	4171
depth-threshold .....	4172
export-profiles .....	4173
file (Analytics) .....	4175
interface (Export Profiles) .....	4176
interfaces (Analytics Resource) .....	4177
latency-threshold .....	4178
local (Analytics Collector) .....	4179
resource (Analytics) .....	4180
resource-profiles (Analytics) .....	4181
system (Analytics Resource) .....	4182
system (Export Profiles) .....	4183
traceoptions (Analytics) .....	4184
Configuration Statements: sFlow Technology .....	4184
[edit protocols] Configuration Statement Hierarchy on EX4300 Switches .....	4185
[edit protocols sflow] Configuration Statement Hierarchy on EX Series Switches .....	4186
Supported Statements in the [edit protocols sflow] Hierarchy Level ..	4187
Unsupported Statements in the [edit sflow] Hierarchy Level .....	4187
agent-id .....	4188
collector .....	4188



disable (sFlow Monitoring Technology) . . . . .	4189
interfaces (sFlow Monitoring Technology) . . . . .	4190
polling-interval . . . . .	4191
sample-rate . . . . .	4192
sflow . . . . .	4193
source-ip . . . . .	4194
udp-port . . . . .	4194
Configuration Statements: Ethernet OAM Connectivity Fault Management . . .	4194
action-profile (Applying to OAM CFM, for EX Series Switch Only) . . . . .	4196
age (EX Series Switch Only) . . . . .	4197
auto-discovery (EX Series Switch Only) . . . . .	4197
calculation-weight . . . . .	4198
connectivity-fault-management (EX Series Switch Only) . . . . .	4199
continuity-check (EX Series Switch Only) . . . . .	4200
cycle-time . . . . .	4201
delay . . . . .	4202
delay-variation . . . . .	4203
direction (EX Series Switch Only) . . . . .	4204
hold-interval (OAM CFM, for EX Series Switch Only) . . . . .	4204
interface (OAM CFM, for EX Series Switch Only) . . . . .	4205
interval (EX Series Switch Only) . . . . .	4206
iteration-period . . . . .	4207
level (EX Series Switch Only) . . . . .	4207
linktrace (EX Series Switch Only) . . . . .	4208
loss-threshold (EX Series Switch Only) . . . . .	4208
maintenance-association (EX Series Switch Only) . . . . .	4209
maintenance-domain (EX Series Switch Only) . . . . .	4210
measurement-type (OAM LFM) . . . . .	4211
mep (EX Series Switch Only) . . . . .	4212
mip-half-function (EX Series Switch Only) . . . . .	4213
name-format (EX Series Switch Only) . . . . .	4214
path-database-size (EX Series Switch Only) . . . . .	4214
performance-monitoring (OAM LFM) . . . . .	4215
remote-mep (EX Series Switch Only) . . . . .	4216
sla-iterator-profiles (OAM LFM) . . . . .	4217
Configuration Statements: Ethernet OAM Link Fault Management . . . . .	4217
action (OAM LFM) . . . . .	4218
action-profile . . . . .	4219
allow-remote-loopback . . . . .	4220
ethernet (OAM LFM) . . . . .	4221
event (OAM LFM) . . . . .	4224
event-thresholds . . . . .	4224
frame-error . . . . .	4225
frame-period . . . . .	4225
frame-period-summary . . . . .	4226
interface (OAM LFM) . . . . .	4227
link-adjacency-loss . . . . .	4228
link-discovery . . . . .	4228
link-down . . . . .	4229

link-event-rate	4229
link-fault-management	4230
negotiation-options	4231
no-allow-link-events	4231
oam	4232
pdu-interval	4234
pdu-threshold	4235
remote-loopback	4235
symbol-period	4236
syslog (OAM LFM)	4236
Configuration Statements: RPM	4236
data-fill	4238
data-size	4239
destination-port	4240
dscp-code-point	4241
hardware-timestamp	4242
history-size	4243
moving-average-size	4243
one-way-hardware-timestamp	4244
port (RPM)	4244
probe	4245
probe-count	4246
probe-interval	4246
probe-limit	4247
probe-server	4248
probe-type	4249
routing-instance	4250
routing-instances	4250
rpm (Interfaces)	4251
source-address (Services)	4251
tcp	4252
test	4253
test-interval	4254
thresholds	4255
traps	4256
udp	4257
Configuration Statements: Uplink Failure Detection	4257
action (Uplink Failure Detection)	4258
group (Uplink Failure Detection)	4258
link-to-disable	4259
link-to-monitor	4259
traceoptions (Uplink Failure Detection)	4260
uplink-failure-detection	4261
<b>Chapter 68 Administration</b>	<b>4263</b>
Routine Monitoring	4263
Monitoring Hosts Using the J-Web Ping Host Tool	4263
Monitoring Network Traffic Using Traceroute	4265

Verifying Input and Output for Port Mirroring Analyzers on EX Series Switches . . . . .	4267
Viewing Real-Time Performance Monitoring Information . . . . .	4268
Verifying That Uplink Failure Detection Is Working Correctly . . . . .	4268
Operational Commands: General . . . . .	4269
monitor traffic . . . . .	4270
ping . . . . .	4280
show pfe statistics bridge . . . . .	4284
traceroute . . . . .	4289
Operational Commands: RPM . . . . .	4292
show services rpm active-servers . . . . .	4293
show services rpm history-results . . . . .	4294
show services rpm probe-results . . . . .	4297
Operational Commands: SNMP . . . . .	4302
clear snmp rmon history . . . . .	4303
clear snmp statistics . . . . .	4304
request snmp spoof-trap . . . . .	4306
show snmp health-monitor . . . . .	4312
show snmp inform-statistics . . . . .	4319
show snmp mib . . . . .	4321
show snmp rmon . . . . .	4324
show snmp rmon history . . . . .	4328
show snmp statistics . . . . .	4332
show snmp v3 . . . . .	4336
Operational Commands: Analyzers and Port Mirroring . . . . .	4338
show forwarding-options port-mirroring . . . . .	4339
show forwarding-options analyzer . . . . .	4341
Commands for Network Analytics . . . . .	4342
monitor start (Analytics) . . . . .	4343
show analytics collector . . . . .	4346
show analytics configuration . . . . .	4348
show analytics queue-statistics . . . . .	4352
show analytics status . . . . .	4354
show analytics traffic-statistics . . . . .	4358
Operational Commands: sFlow . . . . .	4360
show sflow . . . . .	4361
show sflow interface . . . . .	4363
show sflow collector . . . . .	4365
clear sflow collectors statistics . . . . .	4366
Operational Commands: Ethernet OAM Connectivity Fault Management . . . . .	4366
clear oam ethernet connectivity-fault-management delay-statistics . . . . .	4367
clear oam ethernet connectivity-fault-management sla-iterator-statistics . . . . .	4369
clear oam ethernet connectivity-fault-management statistics . . . . .	4370
monitor ethernet delay-measurement . . . . .	4371
show oam ethernet connectivity-fault-management delay-statistics . . . . .	4376
show oam ethernet connectivity-fault-management forwarding-state . . . . .	4380
show oam ethernet connectivity-fault-management interfaces . . . . .	4384
show oam ethernet connectivity-fault-management path-database . . . . .	4390

	show oam ethernet connectivity-fault-management mep-database . . .	4393
	show oam ethernet connectivity-fault-management mip . . . . .	4399
	show oam ethernet connectivity-fault-management	
	sla-iterator-statistics . . . . .	4400
	Operational Commands: Ethernet OAM Link Fault Management . . . . .	4404
	show oam ethernet link-fault-management . . . . .	4405
	Operational Commands: Uplink Failure Detection . . . . .	4409
	show uplink-failure-detection . . . . .	4410
<b>Chapter 69</b>	<b>Troubleshooting . . . . .</b>	<b>4413</b>
	Troubleshooting Procedures . . . . .	4413
	Troubleshooting Port Mirroring Configuration Error Messages . . . . .	4413
	An Analyzer Configuration Returns a “Multiple interfaces cannot be configured as a member of Analyzer output VLAN” Error Message . . . . .	4413
	Operational Commands . . . . .	4414
	show pfe statistics bridge . . . . .	4415
<b>Part 22</b>	<b>PoE</b>	
<b>Chapter 70</b>	<b>Overview . . . . .</b>	<b>4423</b>
	Power over Ethernet Overview . . . . .	4423
	Understanding PoE on EX Series Switches . . . . .	4423
	PoE, PoE+, and Enhanced PoE . . . . .	4423
	PoE Power Allocation . . . . .	4425
<b>Chapter 71</b>	<b>Configuration . . . . .</b>	<b>4433</b>
	Configuration Examples . . . . .	4433
	Example: Configuring PoE Interfaces on an EX Series Switch . . . . .	4433
	Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch . . . . .	4435
	Configuration Tasks . . . . .	4440
	Configuring PoE on EX Series Switches (CLI Procedure) . . . . .	4440
	PoE Configurable Options . . . . .	4441
	Configuring the PoE Controller on EX2200, EX3200, EX3300, EX4200, and EX4300 Switches . . . . .	4443
	Configuring the PoE Controllers on EX6200 and EX8200 Switches . .	4445
	Configuring PoE Interfaces . . . . .	4446
	Configuring PoE (J-Web Procedure) . . . . .	4447
	Configuring PoE on EX2200, EX2200-C, EX3200, EX3300, EX4200, and EX4300 Switches . . . . .	4447
	Configuring PoE on EX6200 Switches . . . . .	4448
	Configuration Statements . . . . .	4450
	[edit poe] Configuration Statement Hierarchy on EX Series Switches . .	4450
	Supported Statements in the [edit poe] Hierarchy Level . . . . .	4451
	af-mode . . . . .	4452
	disable (Power over Ethernet) . . . . .	4453
	duration . . . . .	4454
	fpc (Notification Control) . . . . .	4455
	guard-band . . . . .	4456

	interface (Power over Ethernet) . . . . .	4457
	interval (Power over Ethernet) . . . . .	4458
	management . . . . .	4459
	maximum-power (Interface) . . . . .	4460
	notification-control . . . . .	4462
	poe . . . . .	4463
	priority (Power over Ethernet) . . . . .	4464
	telemetries . . . . .	4465
<b>Chapter 72</b>	<b>Administration . . . . .</b>	<b>4467</b>
	Upgrading . . . . .	4467
	Upgrading the PoE Controller Software . . . . .	4467
	Determining Whether the PoE Controller Software Needs	
	Upgrading . . . . .	4468
	Upgrading the PoE Controller Software . . . . .	4468
	Monitoring the Upgrade Progress . . . . .	4469
	Routine Monitoring . . . . .	4470
	Monitoring PoE . . . . .	4470
	Monitoring PoE Power Consumption (CLI Procedure) . . . . .	4471
	PoE Power Consumption on a Switch . . . . .	4471
	Current Power Consumption for PoE Interfaces . . . . .	4471
	Power Consumption for PoE Interfaces over Time . . . . .	4473
	Verifying PoE Configuration and Status (CLI Procedure) . . . . .	4474
	PoE Controller Configuration and Status . . . . .	4474
	PoE Interface Configuration and Status . . . . .	4475
	PoE SNMP Trap Generation Status . . . . .	4476
	PoE Line Card Configuration and Status . . . . .	4477
	Operational Commands . . . . .	4478
	request system firmware upgrade poe . . . . .	4479
	show poe controller . . . . .	4481
	show poe interface . . . . .	4484
	show poe notification-control . . . . .	4488
	show poe telemetries . . . . .	4490
<b>Chapter 73</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>4493</b>
	Troubleshooting PoE Interfaces . . . . .	4493
<b>Part 23</b>	<b>Port Security</b>	
<b>Chapter 74</b>	<b>Overview . . . . .</b>	<b>4497</b>
	Port Security Overview . . . . .	4497
	Understanding Port Security . . . . .	4497
	Understanding DAI for Port Security . . . . .	4500
	Address Resolution Protocol . . . . .	4500
	ARP Spoofing . . . . .	4500
	Dynamic ARP Inspection . . . . .	4501
	Prioritizing Inspected Packets . . . . .	4502
	Understanding DHCP Option 82 for Port Security on Switching Devices . . . . .	4503
	DHCP Option 82 Processing . . . . .	4503
	Suboption Components of Option 82 . . . . .	4504

Switching Device Configurations That Support Option 82 . . . . .	4505
DHCPv6 Options . . . . .	4506
Understanding DHCP Snooping for Port Security . . . . .	4507
DHCP Snooping Basics . . . . .	4507
Enabling DHCP Snooping . . . . .	4508
DHCP Snooping Process . . . . .	4508
DHCPv6 Snooping . . . . .	4509
Rapid Commit for DHCPv6 . . . . .	4509
DHCP Server Access . . . . .	4510
Static IP Address Additions to the DHCP Snooping Database . . . . .	4513
Snooping DHCP Packets That Have Invalid IP Addresses . . . . .	4513
Understanding IP Source Guard for Port Security on EX Series Switches . .	4514
IP Address Spoofing . . . . .	4514
How IP Source Guard Works . . . . .	4514
IPv6 Source Guard . . . . .	4515
The DHCP Snooping Table . . . . .	4515
Typical Uses of Other Junos OS Features with IP Source Guard . . . .	4516
Understanding IPv6 Neighbor Discovery Inspection . . . . .	4516
Understanding MAC Limiting and MAC Move Limiting for Port Security on	
EX Series Switches . . . . .	4518
MAC Limiting for Port Security by Limiting the Number of MAC Addresses	
That Can be Learned on Interfaces . . . . .	4518
MAC Limiting for Port Security by Specifying MAC Addresses That Are	
Allowed to Access Interfaces . . . . .	4519
MAC Move Limiting for Port Security by Monitoring MAC Address Moves	
within VLANs . . . . .	4519
Understanding Media Access Control Security (MACsec) . . . . .	4520
How MACsec Works . . . . .	4520
Understanding Connectivity Associations and Secure Channels . . . .	4521
Understanding MACsec Security Modes . . . . .	4521
Understanding the Requirements to Enable MACsec on a	
Switch-to-Host Link . . . . .	4523
Understanding MACsec Hardware Requirements for EX Series and QFX	
Series Switches . . . . .	4524
Understanding MACsec Software Requirements for EX Series and QFX	
Series Switches . . . . .	4524
Understanding the MACsec Feature License Requirement . . . . .	4525
MACsec Limitations . . . . .	4525
Understanding Persistent MAC Learning (Sticky MAC) . . . . .	4526
Understanding Trusted DHCP Servers for Port Security . . . . .	4527

<b>Chapter 75</b>	<b>Configuration</b>	<b>4529</b>
	Configuration Examples	4529
	Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing	4529
	Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect the Switch from IPv6 Address Spoofing	4534
	Configuration Tasks	4539
	Configuring Port Security (CLI Procedure)	4539
	Configuring Port Security (J-Web Procedure)	4541
	Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)	4545
	Configuring IP Source Guard (CLI Procedure)	4546
	Enabling IPv6 Neighbor Discovery Inspection	4547
	Configuring Media Access Control Security (MACsec)	4548
	Acquiring and Downloading the Junos OS Software	4548
	Acquiring and Downloading the MACsec Feature License	4549
	Configuring the PIC Mode of the MACsec-capable Interfaces (EX4200 switches only)	4550
	Configuring MACsec Using Static Connectivity Association Key Security Mode (Recommended for Enabling MACsec on Switch-to-Switch Links)	4552
	Configuring MACsec on the Switch Using Dynamic Secure Association Key Security Mode to Secure a Switch-to-Host Link	4556
	Configuring MACsec Using Static Secure Association Key Security Mode to Secure a Switch-to-Switch Link	4560
	Configuring MAC Limiting (CLI Procedure)	4564
	Limiting the Number of MAC Addresses Learned by an Interface	4565
	Limiting the Number of MAC Addresses Learned by a VLAN	4565
	Configuring MAC Limiting (J-Web Procedure)	4566
	Configuring MAC Move Limiting (CLI Procedure)	4568
	Configuring Persistent MAC Learning (CLI Procedure)	4569
	Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports (CLI Procedure)	4571
	Enabling Dynamic ARP Inspection (CLI Procedure)	4572
	Enabling Dynamic ARP Inspection (J-Web Procedure)	4572
	Enabling a Trusted DHCP Server (CLI Procedure)	4573
	Enabling a Trusted DHCP Server (J-Web Procedure)	4574
	Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Performance (CLI Procedure)	4575
	Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)	4577
	Enabling DHCPv6 Rapid Commit Support	4579
	Configuration Statements	4579
	[edit vlans] Configuration Statement Hierarchy on EX Series Switches	4581
	Supported Statements in the [edit vlans] Hierarchy Level	4582
	Unsupported Statements in the [edit vlans] Hierarchy Level	4584
	arp-inspection	4585
	cak	4586
	circuit-id	4587

ckn	4588
connectivity-association	4589
connectivity-association (MACsec Interfaces)	4590
direction	4591
dhcp-security	4592
dhcp-service	4594
dhcp-snooping-file	4595
dhcpv6-options	4596
dhcpv6-snooping-file	4597
encryption	4598
exclude-protocol	4599
group (DHCP Security)	4600
host-name	4601
id	4602
interface (DHCP Security)	4603
interface-mac-limit	4604
interfaces (MACsec)	4606
ip-source-guard	4607
ipv6-source-guard	4609
key	4610
key-server-priority	4611
mac	4612
mac-address (MACsec)	4613
mac-move-limit	4614
macsec	4616
mka	4617
must-secure	4618
neighbor-discovery-inspection	4619
no-dhcpv6-options	4620
no-dhcpv6-snooping	4620
no-encryption	4621
no-option16	4622
no-option18	4622
no-option37	4623
offset	4624
option-16 (DHCPv6 Snooping)	4625
option-18 (DHCPv6 Snooping)	4626
option-37 (DHCPv6 Snooping)	4627
overrides (DHCP Security)	4628
packet-action	4629
persistent-learning	4631
port-id	4632
pre-shared-key	4633
prefix (DHCPv6 Options)	4634
recovery-timeout	4635
remote-id	4637
replay-protect	4638
replay-window-size	4639
routing-instance-name (circuit-id)	4640



	secure-channel . . . . .	4641
	security-association . . . . .	4642
	security-mode . . . . .	4643
	static-ip . . . . .	4644
	static-ipv6 . . . . .	4645
	traceoptions (DHCP) . . . . .	4646
	transmit-interval (MACsec) . . . . .	4648
	trusted . . . . .	4649
	untrusted . . . . .	4649
	use-interface-description . . . . .	4650
	use-interface-index . . . . .	4651
	use-interface-name . . . . .	4652
	use-string . . . . .	4653
	use-vlan-id . . . . .	4654
	vendor-id . . . . .	4655
	write-interval . . . . .	4656
<b>Chapter 76</b>	<b>Administration . . . . .</b>	<b>4657</b>
	Operational Commands . . . . .	4657
	clear arp . . . . .	4658
	clear dhcp-security binding . . . . .	4660
	clear dhcp-security ipv6 binding . . . . .	4661
	clear dot1x . . . . .	4662
	clear ethernet-switching recovery-timeout . . . . .	4664
	clear security mka statistics . . . . .	4665
	show dhcp-security arp inspection statistics . . . . .	4666
	show dhcp-security binding . . . . .	4668
	show dhcp-security binding ip-source-guard . . . . .	4671
	show dhcp-security ipv6 binding . . . . .	4673
	show dhcp-security ipv6 statistics . . . . .	4675
	show dhcp-security nd-inspection statistics . . . . .	4678
	show security macsec connections . . . . .	4680
	show security macsec statistics . . . . .	4682
	show security mka sessions . . . . .	4686
	show security mka statistics . . . . .	4688
<b>Part 24</b>	<b>Routing Policy and Packet Filtering</b>	
<b>Chapter 77</b>	<b>Overview . . . . .</b>	<b>4693</b>
	Security Features Overview . . . . .	4693
	Security Features for EX Series Switches Overview . . . . .	4693
	Firewall Filters Overview . . . . .	4695
	Firewall Filters for EX Series Switches Overview . . . . .	4696
	Firewall Filter Types . . . . .	4696
	Firewall Filter Components . . . . .	4697
	Firewall Filter Processing . . . . .	4698
	Understanding Planning of Firewall Filters . . . . .	4699
	Understanding Firewall Filter Processing Points for Bridged and Routed Packets on EX Series Switches . . . . .	4702

Understanding How Firewall Filters Control Packet Flows . . . . .	4703
Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches . . . . .	4704
Firewall Filter Elements . . . . .	4705
Match Conditions Supported on Switches . . . . .	4705
Actions for Firewall Filters . . . . .	4712
Action Modifiers for Firewall Filters . . . . .	4713
Support for Match Conditions and Actions for Loopback Firewall Filters on Switches . . . . .	4715
Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches . . . . .	4718
Firewall Filter Types and Their Bind Points . . . . .	4719
Support for IPv4 and IPv6 Firewall Filters on Switches . . . . .	4720
Platform Support for Match Conditions for IPv4 Traffic . . . . .	4720
Platform Support for Match Conditions for IPv6 Traffic . . . . .	4733
Platform Support for Match Conditions for Non-IP Traffic . . . . .	4742
Platform Support for Actions for IPv4 Traffic . . . . .	4742
Platform Support for Actions for IPv6 Traffic . . . . .	4745
Platform Support for Action Modifiers for IPv4 Traffic . . . . .	4748
Platform Support for Action Modifiers for IPv6 Traffic . . . . .	4754
Understanding How Firewall Filters Are Evaluated . . . . .	4760
Understanding Firewall Filter Match Conditions . . . . .	4762
Filter Match Conditions . . . . .	4762
Numeric Filter Match Conditions . . . . .	4763
Interface Filter Match Conditions . . . . .	4763
IP Address Filter Match Conditions . . . . .	4764
MAC Address Filter Match Conditions . . . . .	4764
Bit-Field Filter Match Conditions . . . . .	4765
Understanding How Firewall Filters Test a Packet's Protocol . . . . .	4766
Understanding the Use of Policers in Firewall Filters . . . . .	4767
Policers Overview . . . . .	4767
Policer Types . . . . .	4767
Policer Actions . . . . .	4768
Policer Levels . . . . .	4769
Color Modes . . . . .	4769
Naming Conventions for Policers . . . . .	4770
Understanding Filter-Based Forwarding for EX Series Switches . . . . .	4770
Understanding Tricolor Marking Architecture . . . . .	4770
<b>Chapter 78 Configuration . . . . .</b>	<b>4773</b>
Configuration Examples . . . . .	4773
Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches . . . . .	4773
Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch . . . . .	4791
Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches . . . . .	4795

Example: Applying Firewall Filters to Multiple Supplicants on Interfaces	
Enabled for 802.1X or MAC RADIUS Authentication . . . . .	4798
Configuration Tasks . . . . .	4803
Configuring Firewall Filters (CLI Procedure) . . . . .	4804
Configuring a Firewall Filter . . . . .	4804
Configuring a Term Specifically for IPv4 or IPv6 Traffic . . . . .	4808
Applying a Firewall Filter to a Port on a Switch . . . . .	4809
Applying a Firewall Filter to a Management Interface on a Switch . . . . .	4810
Applying a Firewall Filter to a VLAN on a Network . . . . .	4811
Applying a Firewall Filter to a Layer 3 (Routed) Interface . . . . .	4812
Configuring Firewall Filters (J-Web Procedure) . . . . .	4813
Configuring Policers to Control Traffic Rates (CLI Procedure) . . . . .	4818
Configuring Policers . . . . .	4819
Specifying Policers in a Firewall Filter Configuration . . . . .	4820
Applying a Firewall Filter That Is Configured with a Policer . . . . .	4820
Assigning Multifield Classifiers in Firewall Filters to Specify	
Packet-Forwarding Behavior (CLI Procedure) . . . . .	4821
Configuring Routing Policies (J-Web Procedure) . . . . .	4822
Configuring Tricolor Marking Policers . . . . .	4827
Configuring a Tricolor Marking Policer . . . . .	4828
Applying Tricolor Marking Policers to Firewall Filters . . . . .	4828
Configuration Statements . . . . .	4829
[edit firewall] Configuration Statement Hierarchy on EX Series Switches . . . . .	4830
Supported Statements in the [edit firewall] Hierarchy Level . . . . .	4831
Unsupported Statements in the [edit firewall] Hierarchy Level . . . . .	4832
Firewall Filter Configuration Statements Supported by Junos OS for EX	
Series Switches . . . . .	4832
action (TCM Policers) . . . . .	4835
apply-path . . . . .	4835
as-path (Policy Options) . . . . .	4836
as-path-group . . . . .	4837
bandwidth-limit . . . . .	4838
burst-size-limit . . . . .	4839
color-aware . . . . .	4840
color-blind . . . . .	4840
committed-burst-size . . . . .	4841
committed-information-rate . . . . .	4842
community (Policy Options) . . . . .	4843
condition . . . . .	4846
damping (Policy Options) . . . . .	4847
dynamic-db . . . . .	4848
excess-burst-size . . . . .	4849
family (Firewall Filter) . . . . .	4850
filter (Firewall Filters) . . . . .	4851
filter (VLANs) . . . . .	4852
filter-specific . . . . .	4853
firewall . . . . .	4854
from . . . . .	4855
if-exceeding . . . . .	4856

	interface-specific . . . . .	4857
	loss-priority high then discard (Three-Color Policer) . . . . .	4857
	peak-burst-size . . . . .	4858
	policer . . . . .	4859
	policy-options . . . . .	4860
	policy-statement . . . . .	4861
	prefix-list . . . . .	4865
	routing-instance . . . . .	4866
	single-rate . . . . .	4867
	term . . . . .	4868
	then (Firewall Filters) . . . . .	4869
	then (Policer Action) . . . . .	4870
	three-color-policer (Configuring) . . . . .	4871
	two-rate . . . . .	4872
<b>Chapter 79</b>	<b>Administration . . . . .</b>	<b>4873</b>
	Routine Monitoring . . . . .	4873
	Verifying That Firewall Filters Are Operational . . . . .	4873
	Verifying That Policers Are Operational . . . . .	4874
	Monitoring Firewall Filter Traffic . . . . .	4875
	Monitoring Traffic for All Firewall Filters and Policers That Are Configured on the Switch . . . . .	4875
	Monitoring Traffic for a Specific Firewall Filter . . . . .	4875
	Monitoring Traffic for a Specific Policer . . . . .	4876
	Operational Commands . . . . .	4876
	clear firewall . . . . .	4877
	show firewall . . . . .	4879
	show firewall log . . . . .	4886
	show policer . . . . .	4889
	show policy . . . . .	4891
	show policy conditions . . . . .	4893
	test policy . . . . .	4895
<b>Chapter 80</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>4897</b>
	Troubleshooting Firewall Filters . . . . .	4897
	A Firewall Filter Configuration Returns a “No Space Available in TCAM” Message . . . . .	4897
<b>Part 25</b>	<b>Spanning-Tree Protocols</b>	
<b>Chapter 81</b>	<b>Overview . . . . .</b>	<b>4903</b>
	Spanning Trees Overview . . . . .	4903
	Understanding MSTP for EX Series Switches . . . . .	4904
	MSTP Maps Multiple VLANs . . . . .	4904
	Configuring MSTP Regions . . . . .	4905
	Selecting a Spanning-Tree Protocol . . . . .	4905
	Understanding RSTP for EX Series Switches . . . . .	4906
	Spanning-Tree Protocols Help Prevent Broadcast Storms . . . . .	4906
	RSTP is an Enhancement of the Original STP . . . . .	4907
	Port Roles Determine Participation in The Spanning-Tree . . . . .	4907

Port States Determine How a Port Processes a Frame . . . . .	4907
Edge Ports Connect to Devices That Cannot Be Part of a Spanning-Tree . . . . .	4908
BPDUs Maintain the Spanning-Tree . . . . .	4908
When an RSTP Root Bridge Fails . . . . .	4909
Switches Must Relearn MAC Addresses After a Link Failure . . . . .	4909
Selecting a Spanning-Tree Protocol . . . . .	4909
Understanding STP for EX Series Switches . . . . .	4910
Selecting a Spanning-Tree Protocol . . . . .	4911
RSTP or VSTP Forced to Run as IEEE 802.1D STP . . . . .	4912
Understanding VSTP for EX Series Switches and QFX Series Switches . . . . .	4912
Selecting a Spanning-Tree Protocol . . . . .	4913
Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches . . . . .	4914
Different Kinds of BPDUs . . . . .	4914
Protecting Switches From Incompatible BPDUs . . . . .	4914
Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches . . . . .	4916
Understanding Root Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches . . . . .	4917
<b>Chapter 82 Configuration . . . . .</b>	<b>4919</b>
Configuration Examples . . . . .	4919
Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches . . . . .	4919
Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches . . . . .	4936
Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches . . . . .	4956
Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches . . . . .	4961
Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches . . . . .	4967
Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on EX Series Switches . . . . .	4971
Configuration Tasks . . . . .	4976
Forcing RSTP or VSTP to Run as IEEE 802.1D STP (CLI Procedure) . . . . .	4976
Configuring Spanning-Tree Protocols (J-Web Procedure) . . . . .	4977
Configuring VSTP (CLI Procedure) . . . . .	4982
Unblocking an Interface That Receives BPDUs in Error (CLI Procedure) . . . . .	4984
Operational Tasks . . . . .	4984
Reverting to RSTP or VSTP from Forced IEEE 802.1D STP . . . . .	4985
Configuration Statements . . . . .	4985
[edit protocols] Configuration Statement Hierarchy on EX Series Switches . . . . .	4986
block (Spanning Trees) . . . . .	4989
bpdu-block . . . . .	4990
bpdu-block-on-edge . . . . .	4991

bpdu-timeout-action (Spanning Trees) .....	4992
bridge-priority (Spanning Trees) .....	4993
configuration-name (Spanning Trees) .....	4994
cost (Spanning Trees) .....	4995
disable (Spanning Trees) .....	4996
disable-timeout (Spanning Trees) .....	4997
edge (Spanning Trees) .....	4998
force-version (IEEE 802.1D STP) .....	4999
forward-delay (Spanning Trees) .....	5000
hello-time (Spanning Trees) .....	5001
interface (BPDU Block) .....	5002
interface (Spanning Trees) .....	5004
log (Spanning Trees) .....	5005
max-age (Spanning Trees) .....	5006
max-hops (Spanning Trees) .....	5007
mode (Spanning Trees) .....	5008
msti (Spanning Trees) .....	5009
mstp (Spanning Trees) .....	5010
no-root-port (Spanning Trees) .....	5011
priority (Spanning Trees) .....	5012
revision-level (Spanning Trees) .....	5013
rstp (Spanning Trees) .....	5014
stp (Spanning Trees) .....	5016
traceoptions (Spanning Trees) .....	5018
vlan (Spanning Trees) .....	5021
vlan (VSTP) .....	5023
vstp .....	5025

## Chapter 83

<b>Administration .....</b>	<b>5027</b>
Routine Monitoring .....	5027
Monitoring Spanning-Tree Protocols .....	5027
Operational Commands .....	5029
clear ethernet-switching bpdu-error .....	5031
clear spanning-tree protocol-migration .....	5032
clear spanning-tree statistics .....	5033
clear spanning-tree statistics .....	5034
show spanning-tree bridge .....	5035
show spanning-tree bridge .....	5040
show spanning-tree interface .....	5045
show spanning-tree interface .....	5049
show spanning-tree mstp configuration .....	5055
show spanning-tree mstp configuration .....	5057
show spanning-tree statistics .....	5059
show spanning-tree statistics .....	5061

## Part 26

## Chapter 84

## Virtual Chassis

<b>Overview</b>	<b>5065</b>
Virtual Chassis Overview	5065
Understanding EX4300 Virtual Chassis	5065
Understanding EX Series Virtual Chassis Components	5067
Maximum Number of Switches per Virtual Chassis	5067
Virtual Chassis Ports (VCPs)	5068
Master Role	5070
Backup Role	5070
Linecard Role	5071
Member Switch and Member ID	5071
Mastership Priority	5072
Mixed Virtual Chassis	5072
Virtual Chassis Identifier (VCID)	5073
Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric	5073
Virtual Chassis Fabric Summary	5074
Understanding Mixed Virtual Chassis Fabric	5075
Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches	5075
Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches	5076
Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis	5077
Understanding Mixed EX4300 and EX4600 Virtual Chassis	5077
Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis	5077
Understanding How the Master in a Virtual Chassis Is Elected	5078
Understanding Software Upgrade in an EX Series Virtual Chassis	5079
Understanding Global Management of a Virtual Chassis	5080
Understanding Nonvolatile Storage in a Virtual Chassis	5083
Nonvolatile Memory Features	5083
Understanding EX Series Virtual Chassis Port Link Aggregation	5083
Understanding EX Series Virtual Chassis Configuration	5086
Understanding Split and Merge in a Virtual Chassis	5087
What Happens When a Virtual Chassis Configuration Splits	5088
Merging Virtual Chassis Configurations	5089
Understanding Automatic Software Update on Virtual Chassis Member Switches	5090
Automatic Software Update Basics	5091
Automatic Software Update Restrictions	5091
Understanding MAC Address Assignment on a Virtual Chassis	5093
Understanding High Availability on an EX Series Virtual Chassis	5094

<b>Chapter 85</b>	<b>Configuration</b>	<b>5097</b>
	Configuration Tasks	5097
	Configuring an EX4300 Virtual Chassis (CLI Procedure)	5097
	Configuring an EX4300 Virtual Chassis with a Nonprovisioned Configuration File	5099
	Configuring an EX4300 Virtual Chassis with a Preprovisioned Configuration File	5101
	Adding a New Switch to an Existing EX4300 Virtual Chassis (CLI Procedure)	5103
	Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure)	5104
	Remove, Repair, and Reinstall the Same Switch	5105
	Remove a Member Switch, Replace It with a Different Switch, and Reapply the Old Configuration	5105
	Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch	5106
	Configuring Mastership of a Virtual Chassis (CLI Procedure)	5107
	Configuring Mastership Using a Preprovisioned Configuration File	5107
	Configuring Mastership Using a Configuration File That Is Not Preprovisioned	5108
	Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)	5109
	Setting an Uplink VCP Between the Member Switches	5111
	Setting an Uplink VCP on a Standalone Switch	5111
	Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis (CLI Procedure)	5113
	Disabling Split and Merge in a Virtual Chassis (CLI Procedure)	5114
	Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure)	5114
	Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure)	5116
	Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)	5116
	Configuration Statements	5117
	[edit virtual-chassis] Configuration Statement Hierarchy	5117
	Supported Statements in the [edit virtual-chassis] Hierarchy Level	5118
	Unsupported Statements in the [edit virtual-chassis] Hierarchy Level	5119
	aliases (Virtual Chassis)	5120
	alias-name (Virtual Chassis aliases)	5121
	auto-sw-update	5122
	graceful-restart (Enabling Globally)	5124
	graceful-switchover	5125
	id	5126
	location (Virtual Chassis)	5127
	mac-persistence-timer	5128
	mastership-priority	5129
	member	5131
	no-management-vlan	5132



	no-split-detection . . . . .	5133
	package-name . . . . .	5134
	preprovisioned . . . . .	5135
	redundancy (Graceful Switchover) . . . . .	5136
	role . . . . .	5137
	serial-number . . . . .	5140
	serial-number (Virtual Chassis aliases) . . . . .	5141
	traceoptions (Virtual Chassis) . . . . .	5142
	vcp-no-hold-time . . . . .	5145
	virtual-chassis . . . . .	5147
<b>Chapter 86</b>	<b>Administration . . . . .</b>	<b>5149</b>
	Routine Monitoring . . . . .	5149
	Command Forwarding Usage with an EX Series Virtual Chassis . . . . .	5149
	Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member . . . . .	5154
	Verifying That Virtual Chassis Ports Are Operational . . . . .	5155
	Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis . . . . .	5157
	Verifying That Graceful Routing Engine Switchover Is Working in the Virtual Chassis . . . . .	5159
	Operational Commands . . . . .	5160
	clear virtual-chassis vc-port statistics . . . . .	5161
	request session member . . . . .	5163
	request virtual-chassis recycle . . . . .	5164
	request virtual-chassis renumber . . . . .	5165
	request virtual-chassis vc-port . . . . .	5166
	request virtual-chassis vc-port diagnostics optics . . . . .	5168
	show system uptime . . . . .	5169
	show virtual-chassis active-topology . . . . .	5174
	show virtual-chassis device-topology . . . . .	5179
	show virtual-chassis login . . . . .	5185
	show virtual-chassis protocol adjacency . . . . .	5186
	show virtual-chassis protocol database . . . . .	5190
	show virtual-chassis protocol interface . . . . .	5194
	show virtual-chassis protocol route . . . . .	5197
	show virtual-chassis protocol statistics . . . . .	5200
	show virtual-chassis . . . . .	5203
	show virtual-chassis vc-path . . . . .	5207
	show virtual-chassis vc-port . . . . .	5209
	show virtual-chassis vc-port diagnostics optics . . . . .	5213
	show virtual-chassis vc-port statistics . . . . .	5227
<b>Chapter 87</b>	<b>Troubleshooting . . . . .</b>	<b>5235</b>
	Troubleshooting Procedures . . . . .	5235
	Troubleshooting an EX Series Virtual Chassis . . . . .	5235
	A Disconnected Member Switch's ID Is Not Available for Reassignment . . . . .	5235
	Load Factory Default Does Not Commit on a Multimember Virtual Chassis . . . . .	5235

The Member ID Persists When a Member Switch Is Disconnected From a Virtual Chassis . . . . .	5236
A Member Switch Is Not Participating in a Mixed Virtual Chassis . . . .	5236

## Part 27

### Chapter 88

## Virtual Chassis Fabric

<b>Overview . . . . .</b>	<b>5241</b>
Virtual Chassis Fabric Overview . . . . .	5241
Virtual Chassis Fabric Overview . . . . .	5241
Understanding Virtual Chassis Fabric Components . . . . .	5243
Spine-and-Leaf Topology . . . . .	5243
Spine Devices . . . . .	5244
Leaf Devices . . . . .	5244
Routing Engine Role . . . . .	5245
Linecard Role . . . . .	5246
Master Routing Engine Election Process . . . . .	5246
Virtual Chassis Ports (VCPs) . . . . .	5247
Automatic Virtual Chassis Port (VCP) Conversion . . . . .	5247
VCF Configuration Options . . . . .	5248
Fabric Mode . . . . .	5248
Mixed Mode . . . . .	5249
Virtual Management Ethernet Interface . . . . .	5249
Virtual Chassis Fabric Port Link Aggregation Group Bundles . . . . .	5249
Virtual Chassis Fabric License Requirements . . . . .	5250
Hardware Requirements for a Virtual Chassis Fabric . . . . .	5250
Software Requirements in a Virtual Chassis Fabric . . . . .	5250
Understanding Virtual Chassis Fabric Configuration . . . . .	5251
Virtual Chassis Fabric Setup . . . . .	5251
Configuration File Management in a VCF . . . . .	5253
Logging into a Virtual Chassis Fabric . . . . .	5253
Understanding Interface Numbering . . . . .	5253
Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric . . . . .	5254
Virtual Chassis Fabric Summary . . . . .	5254
Understanding Mixed Virtual Chassis Fabric . . . . .	5255
Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches . . . . .	5255
Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches . . . . .	5256
Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis . . . . .	5257
Understanding Mixed EX4300 and EX4600 Virtual Chassis . . . . .	5257
Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis . . . . .	5257
Understanding Traffic Flow Through a Virtual Chassis Fabric . . . . .	5258
Understanding Software Upgrades in a Virtual Chassis Fabric . . . . .	5258
Virtual Chassis Fabric Software Basics . . . . .	5259
Nonstop Software Upgrade (NSSU) . . . . .	5259

	Automatic Software Update . . . . .	5259
	Traditional Software Upgrade . . . . .	5259
<b>Chapter 89</b>	<b>Configuration . . . . .</b>	<b>5261</b>
	Configuration Tasks . . . . .	5261
	Autoprovisioning a Virtual Chassis Fabric . . . . .	5261
	Preprovisioning a Virtual Chassis Fabric . . . . .	5264
	Configuring a Nonprovisioned Virtual Chassis Fabric . . . . .	5267
	Adding a Device to a Virtual Chassis Fabric . . . . .	5270
	Adding a Leaf Device to an Autoprovisioned Virtual Chassis Fabric . . . . .	5271
	Adding a Spine Device to an Autoprovisioned Virtual Chassis Fabric . . . . .	5272
	Adding a Spine or Leaf Device to a Preprovisioned Virtual Chassis Fabric . . . . .	5273
	Adding a Spine or Leaf Device to a Nonprovisioned Virtual Chassis Fabric . . . . .	5275
	Removing a Device From a Virtual Chassis Fabric . . . . .	5277
	Upgrading Software for a Virtual Chassis Fabric . . . . .	5278
	NSSU . . . . .	5279
	Automatic Software Update . . . . .	5279
	Standard Upgrade . . . . .	5279
	Configuration Statements . . . . .	5280
	[edit virtual-chassis] Configuration Statement Hierarchy . . . . .	5280
	Supported Statements in the [edit virtual-chassis] Hierarchy Level . . . . .	5281
	Unsupported Statements in the [edit virtual-chassis] Hierarchy Level . . . . .	5282
	aliases (Virtual Chassis) . . . . .	5283
	alias-name (Virtual Chassis aliases) . . . . .	5284
	auto-provisioned . . . . .	5285
	auto-sw-update . . . . .	5286
	enhanced-hash-key . . . . .	5288
	fabric-load-balance . . . . .	5290
	id . . . . .	5291
	inactivity-interval (Fabric Load Balance) . . . . .	5292
	location (Virtual Chassis) . . . . .	5293
	mac-persistence-timer . . . . .	5294
	mastership-priority . . . . .	5295
	member . . . . .	5297
	no-management-vlan . . . . .	5298
	no-split-detection . . . . .	5299
	package-name . . . . .	5300
	preprovisioned . . . . .	5301
	role . . . . .	5302
	serial-number . . . . .	5305
	serial-number (Virtual Chassis aliases) . . . . .	5306
	traceoptions (Virtual Chassis) . . . . .	5307
	virtual-chassis . . . . .	5310

<b>Chapter 90</b>	<b>Administration . . . . .</b>	<b>5313</b>
	Routine Monitoring . . . . .	5313
	Verifying the Member ID, Role, Status, and Neighbor Member Connections of a Virtual Chassis Fabric Member Device . . . . .	5313
	Verifying Virtual Chassis Port Connections in a Virtual Chassis Fabric . . . .	5314
	Verifying the Virtual Chassis Fabric Mode Settings . . . . .	5315
	Operational Commands . . . . .	5315
	clear virtual-chassis vc-port statistics . . . . .	5317
	request session member . . . . .	5319
	request virtual-chassis mode . . . . .	5320
	request virtual-chassis reactivate . . . . .	5323
	request virtual-chassis vc-port . . . . .	5324
	request virtual-chassis vc-port diagnostics optics . . . . .	5326
	show forwarding-options enhanced-hash-key . . . . .	5327
	show virtual-chassis active-topology . . . . .	5330
	show virtual-chassis device-topology . . . . .	5335
	show virtual-chassis login . . . . .	5341
	show virtual-chassis mode . . . . .	5342
	show virtual-chassis protocol adjacency . . . . .	5345
	show virtual-chassis protocol database . . . . .	5349
	show virtual-chassis protocol interface . . . . .	5353
	show virtual-chassis protocol route . . . . .	5356
	show virtual-chassis protocol statistics . . . . .	5359
	show virtual-chassis . . . . .	5362
	show virtual-chassis vc-port . . . . .	5366
	show virtual-chassis vc-port diagnostics optics . . . . .	5370
	show virtual-chassis vc-port statistics . . . . .	5384
<b>Chapter 91</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>5391</b>
	Troubleshooting Virtual Chassis Fabric . . . . .	5391
	Virtual Chassis Port Link Does Not Form . . . . .	5391
	QFX5100 Leaf Device Assumes Routing Engine Role . . . . .	5392

# List of Figures

<b>Part 2</b>	<b>System Setup</b>	
<b>Chapter 5</b>	<b>Troubleshooting Procedures</b> .....	<b>305</b>
	Figure 1: Connecting to the Console Port on the EX Series Switch .....	306
<b>Part 4</b>	<b>Access and User Management</b>	
<b>Chapter 12</b>	<b>Troubleshooting Procedures</b> .....	<b>601</b>
	Figure 2: Connecting to the Console Port on the EX Series Switch .....	602
<b>Part 5</b>	<b>Configuration and File Management</b>	
<b>Chapter 14</b>	<b>Configuration</b> .....	<b>617</b>
	Figure 3: EX Series Switch LCD Panel .....	621
<b>Chapter 16</b>	<b>Troubleshooting Procedures</b> .....	<b>685</b>
	Figure 4: Connecting to the Console Port on the EX Series Switch .....	686
<b>Part 8</b>	<b>System Services</b>	
<b>Chapter 24</b>	<b>Overview</b> .....	<b>1353</b>
	Figure 5: DHCP Options Data Flow .....	1375
<b>Part 9</b>	<b>Access Control</b>	
<b>Chapter 27</b>	<b>Overview</b> .....	<b>1821</b>
	Figure 6: Example Authentication Topology .....	1825
	Figure 7: VoIP Multiple Supplicant Topology .....	1835
	Figure 8: VoIP Single Supplicant Topology .....	1836
	Figure 9: Authentication Process Flow for an EX Series Switch .....	1840
<b>Chapter 28</b>	<b>Configuration</b> .....	<b>1843</b>
	Figure 10: Topology for Configuration .....	1845
	Figure 11: Topology for Guest VLAN Example .....	1849
	Figure 12: Topology for Configuring Supplicant Modes .....	1854
	Figure 13: VoIP Topology .....	1860
	Figure 14: Topology for Static MAC Authentication Configuration .....	1873
	Figure 15: Topology for Configuration .....	1877
	Figure 16: Topology for MAC RADIUS Authentication Configuration .....	1883
	Figure 17: Topology for Firewall Filter and RADIUS Server Attributes Configuration .....	1889
	Figure 18: Conceptual Model: Dynamic Filter Updated for Each New User . . .	1895

	Figure 19: Multiple Supplicants on an 802.1X-Enabled Interface Connecting to a File Server . . . . .	1896
	Figure 20: EX Series Switch Connecting OAC to RADIUS Server Using EAP-TTLS Authentication . . . . .	1905
	Figure 21: Example of a Captive Portal Login Page . . . . .	1928
<b>Part 10</b>	<b>Class of Service</b>	
<b>Chapter 30</b>	<b>Overview . . . . .</b>	<b>2041</b>
	Figure 22: Packet Flow Across the Network . . . . .	2043
	Figure 23: Graphical Representation of a Segmented Drop Profile . . . . .	2057
	Figure 24: Graphical Representation of an Interpolated Drop Profile on EX Series Switches Except EX4300 Switches . . . . .	2058
	Figure 25: Tail-Drop Profile Packet Drop on EX4300 Switches . . . . .	2059
<b>Chapter 31</b>	<b>Configuration . . . . .</b>	<b>2075</b>
	Figure 26: Topology for Configuring CoS . . . . .	2076
<b>Part 12</b>	<b>Ethernet Switching</b>	
<b>Chapter 37</b>	<b>Overview . . . . .</b>	<b>2245</b>
	Figure 27: An IRB Interface or RVI on a Switch Providing Routing Between Two Access Switches . . . . .	2255
	Figure 28: Creating an IRB Interface or RVI . . . . .	2256
	Figure 29: Private VLAN on a Single Switch . . . . .	2260
	Figure 30: PVLAN Spanning Multiple Switches . . . . .	2261
	Figure 31: L2PT Example . . . . .	2274
	Figure 32: Redundant Trunk Group, Link 1 Active . . . . .	2277
	Figure 33: Redundant Trunk Group, Link 2 Active . . . . .	2277
<b>Chapter 38</b>	<b>Configuration . . . . .</b>	<b>2281</b>
	Figure 34: Sample Access Switch-Distribution Switch Topology . . . . .	2293
	Figure 35: Topology of a Private VLAN on a Single EX Series Switch . . . . .	2303
	Figure 36: MVRP Configured on Two Access Switches and One Distribution Switch for Automatic VLAN Administration . . . . .	2312
	Figure 37: L2PT Topology . . . . .	2324
	Figure 38: Topology for Configuring the Redundant Trunk Links . . . . .	2328
<b>Part 13</b>	<b>High Availability</b>	
<b>Chapter 40</b>	<b>Overview . . . . .</b>	<b>2483</b>
	Figure 39: Basic VRRP on EX Series Switches . . . . .	2502
	Figure 40: VRRP on Virtual Chassis Switches . . . . .	2503
<b>Part 14</b>	<b>Interfaces</b>	
<b>Chapter 44</b>	<b>Overview . . . . .</b>	<b>2577</b>
	Figure 41: Egress Traffic Flow with Local Link Bias . . . . .	2590
	Figure 42: Egress Traffic Flow without Local Link Bias . . . . .	2591
	Figure 43: Symmetrically Routed Interfaces . . . . .	2594
	Figure 44: Asymmetrically Routed Interfaces . . . . .	2595

<b>Part 20</b>	<b>Multicast</b>	
<b>Chapter 63</b>	<b>Overview</b>	<b>3717</b>
	Figure 45: Multicast Traffic Flow with IGMP Snooping Enabled	3718
	Figure 46: Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts	3722
	Figure 47: Scenario 2: Switch Forwarding Multicast Traffic to Another Switch	3723
	Figure 48: Scenario 3: Switch Connected to Hosts Only (No IGMP Querier)	3724
	Figure 49: Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs	3725
<b>Chapter 64</b>	<b>Configuration</b>	<b>3727</b>
	Figure 50: IGMP Snooping Topology Sample Topology	3728
<b>Part 21</b>	<b>Network Management and Monitoring</b>	
<b>Chapter 66</b>	<b>Overview</b>	<b>3959</b>
	Figure 51: RPM Timestamps	3962
	Figure 52: Relationship Among MEPs, MIPs, and Maintenance Domain Levels	3974
	Figure 53: Uplink Failure Detection Configuration on Switches	3977
<b>Chapter 67</b>	<b>Configuration</b>	<b>3995</b>
	Figure 54: sFlow Technology Monitoring System	3997
	Figure 55: Network Topology for Local Mirroring Example	4002
	Figure 56: Remote Mirroring Network Topology Example	4008
	Figure 57: Remote Mirroring Through a Transit Switch Network—Sample Topology	4017
<b>Part 23</b>	<b>Port Security</b>	
<b>Chapter 74</b>	<b>Overview</b>	<b>4497</b>
	Figure 58: DHCP Clients, Switching Device, and the DHCP Server Are All on the Same VLAN or Bridge Domain	4505
	Figure 59: Switching Device Acting as an Extended Relay Server	4506
	Figure 60: DHCP Server Connected Directly to Switch	4511
	Figure 61: DHCP Server Connected Directly to Switch 2, with Switch 2 Connected to Switch 1 Through a Trusted Trunk Port	4511
	Figure 62: Switch Is the DHCP Server	4512
	Figure 63: Switch Acting as Relay Agent Through Router to DHCP Server	4513
<b>Chapter 75</b>	<b>Configuration</b>	<b>4529</b>
	Figure 64: Network Topology for Basic Port Security	4531
	Figure 65: Network Topology for Basic Port Security	4536
<b>Part 24</b>	<b>Routing Policy and Packet Filtering</b>	
<b>Chapter 77</b>	<b>Overview</b>	<b>4693</b>
	Figure 66: Firewall Filter Processing Points in the Packet Forwarding Path	4702
	Figure 67: Application of Firewall Filters to Control Packet Flow	4704
	Figure 68: Evaluation of Terms Within a Firewall Filter	4761
<b>Chapter 78</b>	<b>Configuration</b>	<b>4773</b>

	Figure 69: Application of Port, VLAN, and Layer 3 Routed Firewall Filters . . . .	4775
	Figure 70: SSH Connection From a Management PC to an EX Series Switch . .	4792
	Figure 71: Conceptual Model: Dynamic Filter Updated for Each New User . . .	4800
	Figure 72: Multiple Supplicants on an 802.1X-Enabled Interface Connecting to a File Server . . . . .	4801
<b>Part 25</b>	<b>Spanning-Tree Protocols</b>	
<b>Chapter 82</b>	<b>Configuration . . . . .</b>	<b>4919</b>
	Figure 73: Network Topology for RSTP . . . . .	4921
	Figure 74: Network Topology for MSTP . . . . .	4938
	Figure 75: BPDU Protection Topology . . . . .	4958
	Figure 76: BPDU Protection Topology . . . . .	4963
	Figure 77: Network Topology for Loop Protection . . . . .	4968
	Figure 78: Network Topology for Root Protection . . . . .	4973
<b>Part 26</b>	<b>Virtual Chassis</b>	
<b>Chapter 84</b>	<b>Overview . . . . .</b>	<b>5065</b>
	Figure 79: Console Session Redirection (EX4200 Virtual Chassis Pictured) . . .	5081
	Figure 80: Management Ethernet Port Redirection to the VME Interface . . . .	5082
<b>Part 27</b>	<b>Virtual Chassis Fabric</b>	
<b>Chapter 88</b>	<b>Overview . . . . .</b>	<b>5241</b>
	Figure 81: VCF Spine-and-Leaf Architecture . . . . .	5242
	Figure 82: VCF Spine-and-Leaf Architecture . . . . .	5244



# List of Tables

	<b>About the Documentation</b> . . . . .	<b>cix</b>
	Table 1: Notice Icons . . . . .	cxi
	Table 2: Text and Syntax Conventions . . . . .	cxi
<b>Part 1</b>	<b>Getting Started with ELS</b>	
<b>Chapter 1</b>	<b>Getting Started with ELS</b> . . . . .	<b>3</b>
	Table 3: ELS Support . . . . .	3
	Table 4: Enhanced Layer 2 CLI Changes . . . . .	9
	Table 5: Enhanced Layer 2 CLI Changes . . . . .	17
<b>Part 2</b>	<b>System Setup</b>	
<b>Chapter 2</b>	<b>Overview</b> . . . . .	<b>29</b>
	Table 6: Junos OS Processes . . . . .	30
<b>Chapter 3</b>	<b>Configuration</b> . . . . .	<b>33</b>
	Table 7: Date and Time Settings . . . . .	35
	Table 8: System Identity Configuration Summary . . . . .	36
	Table 9: Menu Options of the LCD Menus Supported on the Switches . . . . .	60
<b>Chapter 4</b>	<b>Administration</b> . . . . .	<b>79</b>
	Table 10: request system storage cleanup Output Fields . . . . .	127
	Table 11: show chassis fan Output Fields . . . . .	152
	Table 12: show chassis firmware Output Fields . . . . .	166
	Table 13: show chassis lcd Output Fields . . . . .	176
	Table 14: show ntp associations Output Fields . . . . .	191
	Table 15: show ntp status Output Fields . . . . .	193
	Table 16: show system firmware Output Fields . . . . .	196
	Table 17: show system storage Output Fields . . . . .	212
	Table 18: show system switchover Output Fields . . . . .	219
	Table 19: show system uptime Output Fields . . . . .	223
	Table 20: show system users Output Fields . . . . .	228
	Table 21: show system virtual-memory Output Fields . . . . .	234
	Table 22: show task replication Output Fields . . . . .	289
<b>Part 3</b>	<b>User Interfaces</b>	
<b>Chapter 6</b>	<b>Overview</b> . . . . .	<b>311</b>
	Table 23: Junos OS Processes . . . . .	312
<b>Chapter 7</b>	<b>Configuration</b> . . . . .	<b>317</b>

Table 24: Unsupported [edit access] Configuration Statements on EX Series Switches .....	324
Table 25: Unsupported [edit forwarding-options] Configuration Statements on EX Series Switches .....	340
Table 26: Unsupported [edit interfaces ae] Configuration Statements on EX Series Switches .....	346
Table 27: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches .....	358
Table 28: Unsupported [edit interfaces lo] Configuration Statements for EX Series Switches .....	363
Table 29: Unsupported [edit interfaces me] Configuration Statements for EX Series Switches .....	366
Table 30: Unsupported [edit protocols bgp] Configuration Statements on EX Series Switches .....	400
Table 31: Unsupported [edit protocols connections] Configuration Statements on EX Series Switches .....	402
Table 32: Unsupported [edit protocols dcbx] Configuration Statements on EX Series Switches .....	404
Table 33: Unsupported [edit protocols isis] Configuration Statements on EX Series Switches .....	411
Table 34: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches .....	420
Table 35: Unsupported [edit protocols mvrp] Configuration Statements on EX Series Switches .....	434
Table 36: Unsupported [edit protocols ospf] Configuration Statements on EX Series Switches .....	440
Table 37: Unsupported [edit protocols ospf 3] Configuration Statements on EX Series Switches .....	444
Table 38: Unsupported [edit protocols-rip] Configuration Statements on EX Series Switches .....	450
Table 39: Unsupported [edit protocols rsvp] Configuration Statements on EX Series Switches .....	456
Table 40: Unsupported [edit routing-options] Configuration Statements on EX Series Switches .....	482
Table 41: Unsupported [edit security] Configuration Statements on EX Series Switches .....	485
Table 42: Unsupported [edit services] Configuration Statements on EX Series Switches .....	488
Table 43: Unsupported [edit snmp] Configuration Statements on EX Series Switches .....	493
Table 44: Unsupported [edit system] Configuration Statements on EX Series Switches .....	508
Table 45: Unsupported [edit vlans] Configuration Statements on EX Series Switches .....	511

## Part 4

### Chapter 9

## Access and User Management

Overview .....	525
Table 46: Junos OS Processes .....	526

<b>Chapter 10</b>	<b>Configuration . . . . .</b>	<b>529</b>
	Table 47: Secure Management Access Configuration Summary . . . . .	530
<b>Chapter 11</b>	<b>Administration . . . . .</b>	<b>579</b>
	Table 48: User Management Configuration Page Summary . . . . .	580
	Table 49: Add an Authentication Server . . . . .	581
	Table 50: show subscribers Output Fields . . . . .	586
<b>Part 5</b>	<b>Configuration and File Management</b>	
<b>Chapter 13</b>	<b>Overview . . . . .</b>	<b>607</b>
	Table 51: Junos OS Processes . . . . .	608
	Table 52: J-Web Edit Point & Click Configuration Links . . . . .	611
	Table 53: J-Web Edit Point & Click Configuration Icons . . . . .	612
	Table 54: J-Web Edit Point & Click Configuration Buttons . . . . .	612
	Table 55: Configuration File Terms . . . . .	615
<b>Chapter 14</b>	<b>Configuration . . . . .</b>	<b>617</b>
	Table 56: Options for the load command . . . . .	618
	Table 57: Commit Options . . . . .	629
	Table 58: Commit Preference Options . . . . .	629
<b>Chapter 15</b>	<b>Administration . . . . .</b>	<b>641</b>
	Table 59: J-Web Configuration History Summary . . . . .	642
	Table 60: J-Web Configuration Database Information Summary . . . . .	642
	Table 61: show system commit Output Fields . . . . .	675
<b>Part 6</b>	<b>Software Installation</b>	
<b>Chapter 17</b>	<b>Overview . . . . .</b>	<b>691</b>
	Table 62: Resilient Dual-Root Partition Scheme . . . . .	695
	Table 63: Earlier Partition Scheme . . . . .	697
	Table 64: Junos OS Processes . . . . .	701
	Table 65: Junos OS EFL Part Number on EX2200 Switches . . . . .	703
	Table 66: Junos OS EFL Part Number on EX3300 Switches . . . . .	704
	Table 67: Junos OS AFL Part Number on EX3300 Switches . . . . .	705
	Table 68: Junos OS EFL Part Number on EX4300 Switches . . . . .	705
	Table 69: Junos OS AFL Part Number on EX4300 Switches . . . . .	706
	Table 70: Junos OS AFL Part Number on EX4600 Switches . . . . .	707
	Table 71: Junos OS AFL Part Number on EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 Switches . . . . .	707
<b>Chapter 19</b>	<b>Administration . . . . .</b>	<b>719</b>
	Table 72: show system auto-snapshot status Output Fields . . . . .	763
	Table 73: show system license Output Fields . . . . .	772
	Table 74: show system snapshot Output Fields . . . . .	780
	Table 75: show system storage partitions Output Fields . . . . .	782
<b>Chapter 20</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>785</b>
	Table 76: Actions If Corrupt Files Are Found and Automatic Snapshot is Enabled . . . . .	790

	Table 77: Actions If Corrupt Files Are Found . . . . .	790
<b>Part 7</b>	<b>System Monitoring</b>	
<b>Chapter 21</b>	<b>Overview . . . . .</b>	<b>795</b>
	Table 78: Junos OS Processes . . . . .	796
	Table 79: Alarm Terms . . . . .	797
	Table 80: Details of a Virtual Chassis Member Switch . . . . .	800
	Table 81: Status of a Member Switch in a Virtual Chassis . . . . .	800
	Table 82: System Information . . . . .	801
	Table 83: Health Status . . . . .	803
	Table 84: Capacity Utilization . . . . .	805
	Table 85: Chassis Viewer for EX2200 Switches . . . . .	806
	Table 86: Chassis Viewer for EX2200-C Switches . . . . .	807
	Table 87: Chassis Viewer for EX3200, EX3300, and EX4200 Switches . . . . .	807
	Table 88: Chassis Viewer for EX4300 Switches . . . . .	809
	Table 89: Chassis Viewer for EX4500 Switches . . . . .	810
	Table 90: Chassis Viewer for EX4550 Switches . . . . .	811
	Table 91: Chassis Viewer for EX6210 Switches . . . . .	813
	Table 92: Chassis Viewer for EX8208 Switches . . . . .	813
	Table 93: Chassis Viewer for EX8216 Switches . . . . .	815
	Table 94: Chassis Viewer for XRE200 External Routing Engines . . . . .	815
	Table 95: CLI Equivalents of Terms Used in the Documentation for EX4300 Switches . . . . .	817
<b>Chapter 23</b>	<b>Administration . . . . .</b>	<b>839</b>
	Table 96: Filtering System Log Messages . . . . .	840
	Table 97: Viewing System Log Messages . . . . .	841
	Table 98: Summary of Key Alarm Output Fields . . . . .	843
	Table 99: Packet Capture Field Summary . . . . .	844
	Table 100: Summary of Key System Properties Output Fields . . . . .	846
	Table 101: Summary of the Key Output Fields for Chassis Information . . . . .	848
	Table 102: Summary of System Process Information Output Fields . . . . .	850
	Table 103: monitor list Output Fields . . . . .	873
	Table 104: monitor start Output Fields . . . . .	874
	Table 105: show chassis alarms Output Fields . . . . .	897
	Table 106: show chassis environment Output Fields . . . . .	913
	Table 107: show chassis environment cb Output Fields . . . . .	972
	Table 108: show chassis environment fpc Output Fields . . . . .	991
	Table 109: show chassis environment routing-engine Output Fields . . . . .	1016
	Table 110: show chassis ethernet-switch Output Fields . . . . .	1022
	Table 111: show chassis fabric fpcs Output Fields . . . . .	1066
	Table 112: show chassis fabric map Output Fields . . . . .	1102
	Table 113: show chassis fabric plane Output Fields . . . . .	1110
	Table 114: show chassis fabric plane-location Output Fields . . . . .	1151
	Table 115: show chassis fabric summary Output Fields . . . . .	1155
	Table 116: show chassis fpc Output Fields . . . . .	1167
	Table 117: show chassis led Output Fields . . . . .	1198
	Table 118: show chassis location Output Fields . . . . .	1209
	Table 119: show chassis pic Output Fields . . . . .	1215

	Table 120: show chassis routing-engine Output Fields . . . . .	1230
	Table 121: show pfe next-hop Output Fields . . . . .	1254
	Table 122: show pfe route Output Fields . . . . .	1260
	Table 123: QFX Series show pfe route Hardware Table Output Fields . . . . .	1260
	Table 124: show system alarms Output Fields . . . . .	1268
	Table 125: show system buffers Output Fields . . . . .	1282
	Table 126: show system connections Output Fields . . . . .	1288
	Table 127: show system core-dumps Output Fields . . . . .	1308
	Table 128: show system directory-usage Output Fields . . . . .	1321
	Table 129: show system processes Output Fields . . . . .	1330
<b>Part 8</b>	<b>System Services</b>	
<b>Chapter 24</b>	<b>Overview . . . . .</b>	<b>1353</b>
	Table 130: Junos OS Processes . . . . .	1354
	Table 131: Comparing the Extended DHCP Local Server to the Traditional DHCP Local Server . . . . .	1357
	Table 132: Information in Authentication Grant . . . . .	1359
	Table 133: RADIUS Attributes and VSAs for DHCPv6 Local Server . . . . .	1361
	Table 134: Action Taken for Events That Occur During a Reconfiguration . . . . .	1366
	Table 135: Unsupported Opaque DHCP Options . . . . .	1377
<b>Chapter 25</b>	<b>Configuration . . . . .</b>	<b>1393</b>
	Table 136: DHCP Service Configuration Pages Summary . . . . .	1406
	Table 137: DHCP Service Configuration Pages Summary for EX4300 Switches . . . . .	1408
	Table 138: DHCP Client Settings . . . . .	1413
	Table 139: Actions for DHCP Local Server Snooped Packets . . . . .	1438
	Table 140: DHCP Relay Agent Option 82 Value for Auto Logout . . . . .	1454
	Table 141: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Enabled . . . . .	1460
	Table 142: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Disabled . . . . .	1461
	Table 143: Actions for Snooped BOOTREPLY Packets . . . . .	1461
	Table 144: Unsupported [edit forwarding-options] Configuration Statements on EX Series Switches . . . . .	1555
	Table 145: Protocol Families and Supported Interface Types . . . . .	1649
<b>Chapter 26</b>	<b>Administration . . . . .</b>	<b>1691</b>
	Table 146: Summary of DHCP Output Fields . . . . .	1692
	Table 147: Summary of DHCP Output Fields for EX4300 Switches . . . . .	1694
	Table 148: Summary of the DHCP Statistics Information Output for EX4300 switches . . . . .	1695
	Table 149: show dhcp server binding Output Fields . . . . .	1717
	Table 150: show dhcp server statistics Output Fields . . . . .	1722
	Table 151: show dhcpv6 server binding Output Fields . . . . .	1725
	Table 152: show dhcpv6 server statistics Output Fields . . . . .	1731
	Table 153: clear dhcp relay statistics Output Fields . . . . .	1736
	Table 154: show dhcp relay binding Output Fields . . . . .	1744
	Table 155: show dhcp relay statistics Output Fields . . . . .	1749

Table 156: show dhcpv6 relay binding Output Fields . . . . .	1752
Table 157: show dhcpv6 relay statistics Output Fields . . . . .	1757
Table 158: show route extensive Output Fields . . . . .	1760
Table 159: show security pki local-certificate Output Fields . . . . .	1800
Table 160: show system services dhcp binding Output Fields . . . . .	1803
Table 161: show system services dhcp conflict Output Fields . . . . .	1806
Table 162: show system services dhcp global Output Fields . . . . .	1807
Table 163: show system services dhcp pool Output Fields . . . . .	1809
Table 164: show system services dhcp statistics Output Fields . . . . .	1811
Table 165: show system services service-deployment Output Fields . . . . .	1814

## Part 9

### Chapter 28

## Access Control

### Configuration . . . . . 1843

Table 166: Components of the Topology . . . . .	1845
Table 167: Components of the Guest VLAN Topology . . . . .	1849
Table 168: Components of the Supplicant Mode Configuration Topology . . . . .	1854
Table 169: Components of the VoIP Configuration Topology . . . . .	1860
Table 170: Components of the Static MAC Authentication Configuration Topology . . . . .	1873
Table 171: Components of the Topology . . . . .	1877
Table 172: Components of the MAC RADIUS Authentication Configuration Topology . . . . .	1883
Table 173: Components of the Firewall Filter and RADIUS Server Attributes Topology . . . . .	1889
Table 174: Components of the OAC Deployment . . . . .	1905
Table 175: Match Conditions . . . . .	1919
Table 176: Actions for VSAs . . . . .	1920
Table 177: Configurable Elements of a Captive Portal Login Page . . . . .	1928
Table 178: Unsupported [edit access] Configuration Statements on EX Series Switches . . . . .	1938

### Chapter 29

### Administration . . . . . 1993

Table 179: clear captive-portal interface Output Fields . . . . .	1997
Table 180: show captive-portal authentication-failed-users Output Fields . . . . .	2003
Table 181: show captive-portal interface Output Fields . . . . .	2006
Table 182: show dot1x Output Fields . . . . .	2009
Table 183: show dot1x authentication-failed-users Output Fields . . . . .	2014
Table 184: show dot1x static-mac-address Output Fields . . . . .	2016
Table 185: show lldp Output Fields . . . . .	2018
Table 186: show lldp local-information Output Fields . . . . .	2023
Table 187: show lldp neighbors Output Fields . . . . .	2025
Table 188: show lldp remote-global-statistics Output Fields . . . . .	2031
Table 189: show lldp statistics Output Fields . . . . .	2033
Table 190: show network-access aaa statistics accounting Output Fields . . . . .	2035
Table 191: show network-access aaa statistics authentication Output Fields . . . . .	2036
Table 192: show network-access aaa statistics dynamic-requests Output Fields . . . . .	2038

<b>Part 10</b>	<b>Class of Service</b>	
<b>Chapter 30</b>	<b>Overview</b>	<b>2041</b>
	Table 193: Default Code-Point Aliases	2046
	Table 194: Default BA Classification	2050
	Table 195: Allowed BA Classification	2050
	Table 196: Default Forwarding Classes for Unicast Traffic	2053
	Table 197: Default Forwarding Classes for Multicast Traffic	2053
	Table 198: Support for Scheduler Maps on Switches and Line Cards	2063
	Table 199: Default Packet Header Rewrite Mappings	2069
	Table 200: EX8200 Line Cards That Include Oversubscribed Ports	2072
<b>Chapter 31</b>	<b>Configuration</b>	<b>2075</b>
	Table 201: Configuration Components: VLANs	2077
	Table 202: Configuration Components: Switch Ports on a 48-Port All-PoE Switch	2077
	Table 203: CoS Value Aliases Configuration Fields	2102
	Table 204: BA-classifier Loss Priority Assignments	2103
	Table 205: Classifiers Configuration Fields	2105
	Table 206: Forwarding Classes Configuration Fields	2108
	Table 207: Schedulers Configuration Page	2112
	Table 208: Scheduler Maps Configuration Fields	2115
	Table 209: Drop Profiles Configuration parameters	2118
	Table 210: Rewrite Rules Configuration Page Summary	2122
	Table 211: Assigning CoS Components to Logical Interfaces	2124
<b>Chapter 32</b>	<b>Administration</b>	<b>2157</b>
	Table 212: Summary of Key CoS Classifier Output Fields	2158
	Table 213: Summary of Key CoS Forwarding Class Output Fields	2159
	Table 214: Summary of Key CoS Interfaces Output Fields	2160
	Table 215: Summary of Key CoS Rewrite Rules Output Fields	2161
	Table 216: Summary of Key CoS Scheduler Maps Output Fields	2162
	Table 217: Summary of Key CoS Value Alias Output Fields	2164
	Table 218: Summary of the Key Output Fields for CoS Red Drop Profiles	2165
	Table 219: show class-of-service Output Fields	2167
	Table 220: show class-of-service classifier Output Fields	2172
	Table 221: show class-of-service code-point-aliases Output Fields	2174
	Table 222: show class-of-service drop-profile Output Fields	2176
	Table 223: show class-of-service forwarding-class Output Fields	2179
	Table 224: show pfe statistics traffic cpu Output Fields	2181
	Table 225: show pfe statistics traffic egress-queues Output Fields	2185
	Table 226: show pfe statistics traffic multicast Output Fields	2188
<b>Part 11</b>	<b>Device Security</b>	
<b>Chapter 35</b>	<b>Configuration</b>	<b>2201</b>
	Table 227: Unsupported [edit switch-options] Configuration Statements on EX Series Switches	2213
<b>Chapter 36</b>	<b>Administration</b>	<b>2233</b>
	Table 228: show ethernet-switching table Output Fields	2239

<b>Part 12</b>	<b>Ethernet Switching</b>	
<b>Chapter 37</b>	<b>Overview</b>	<b>2245</b>
	Table 229: Tracking IRB Interface and RVI Usage	2257
	Table 230: PVLAN Ports and Layer 2 Forwarding	2262
	Table 231: Junos OS MVRP Versions and Inclusion of Extra Byte in PDU	2266
	Table 232: MVRP Environments and Description of Required Actions	2266
	Table 233: Protocol Destination MAC Addresses	2275
<b>Chapter 38</b>	<b>Configuration</b>	<b>2281</b>
	Table 234: Components of the Basic Bridging Configuration Topology	2283
	Table 235: Components of the Topology for Connecting an Access Switch to a Distribution Switch	2293
	Table 236: Interfaces of the Topology for Configuring a PVLAN	2302
	Table 237: VLAN IDs in the Topology for Configuring a PVLAN	2303
	Table 238: Components of the Network Topology	2312
	Table 239: Components of the Redundant Trunk Link Topology	2329
	Table 240: VLAN Configuration Details	2335
	Table 241: RTG Configuration Fields	2359
	Table 242: Unsupported [edit switch-options] Configuration Statements on EX Series Switches	2368
	Table 243: Unsupported [edit vlans] Configuration Statements on EX Series Switches	2372
<b>Chapter 39</b>	<b>Administration</b>	<b>2435</b>
	Table 244: show ethernet-switching interface Output Fields	2442
	Table 245: show ethernet-switching mac-notification Output Fields	2445
	Table 246: show ethernet-switching table Output fields	2448
	Table 247: show interfaces irb Output Fields	2454
	Table 248: show mac-rewrite interface Output Fields	2460
	Table 249: show mvrp Output Fields	2462
	Table 250: show mvrp dynamic-vlan-memberships Output Fields	2464
	Table 251: show mvrp statistics Output Fields	2465
	Table 252: show redundant-trunk-group Output Fields	2467
	Table 253: show route forwarding-table Output Fields	2470
<b>Part 13</b>	<b>High Availability</b>	
<b>Chapter 40</b>	<b>Overview</b>	<b>2483</b>
	Table 254: Platform and Release Support for NSSU	2494
	Table 255: Available Operating Power in N+1 and N+N Redundancy Configurations	2500
<b>Chapter 42</b>	<b>Administration</b>	<b>2545</b>
	Table 256: show chassis power-budget-statistics Output Fields	2558
	Table 257: show vrrp Output Fields	2562
<b>Part 14</b>	<b>Interfaces</b>	
<b>Chapter 44</b>	<b>Overview</b>	<b>2577</b>
	Table 258: Network Interface Types and Purposes	2578



**Chapter 45**

Table 259: Special Interface Types and Purposes . . . . .	2578
Table 260: Maximum Interfaces per LAG and Maximum LAGs per Switch . . . .	2583
Table 261: IPv4 and IPv6 Hashing Fields . . . . .	2587
Table 262: MPLS Hashing Fields . . . . .	2588
Table 263: MAC-in-MAC Hashing Fields . . . . .	2589
Table 264: Layer 2 Header Hashing Fields . . . . .	2589
Table 265: MC-LAG Configuration Guidelines . . . . .	2601
Table 266: MC-LAG Functional Behavior . . . . .	2602
Table 267: MC-LAG Failure Scenarios . . . . .	2604
<b>Configuration . . . . .</b>	<b>2615</b>
Table 268: Factory Default Configuration Link Settings for EX Series Switches . . . . .	2617
Table 269: Port Edit Options . . . . .	2621
Table 270: Recommended CoS Settings for Port Roles . . . . .	2624
Table 271: Port Role Configuration Summary . . . . .	2625
Table 272: Recommended CoS Settings for Port Roles . . . . .	2628
Table 273: Effect of set interfaces disable <interface_name> on T series PICs . .	2630
Table 274: Encapsulation Overhead by Encapsulation Type . . . . .	2641
Table 275: Media MTU Sizes by Interface Type for M5 and M7i Routers with CFEB, M10 and M10i Routers with CFEB, and M20 and M40 Routers . . . . .	2642
Table 276: Media MTU Sizes by Interface Type for M40e Routers . . . . .	2642
Table 277: Media MTU Sizes by Interface Type for M160 Routers . . . . .	2644
Table 278: Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers . . . . .	2644
Table 279: Media MTU Sizes by Interface Type for MX Series Routers . . . . .	2645
Table 280: Media MTU Sizes by Interface Type for T320 Routers . . . . .	2646
Table 281: Media MTU Sizes by Interface Type for T640 Platforms . . . . .	2646
Table 282: Media MTU Sizes by Interface Type for J2300 Platforms . . . . .	2647
Table 283: Media MTU Sizes by Interface Type for J4300 and J6300 Platforms . . . . .	2647
Table 284: Media MTU Sizes by Interface Type for J4350 and J6350 Platforms . . . . .	2648
Table 285: Media MTU Sizes by Interface Type for EX Series Switches and ACX Series Routers . . . . .	2650
Table 286: Media MTU Sizes by Interface Type for PTX Series Packet Transport Routers . . . . .	2650
Table 287: Aggregated Ethernet Interface Options . . . . .	2669
Table 288: VLAN Options . . . . .	2670
Table 289: IP Options . . . . .	2670
Table 290: Unsupported [edit forwarding-options] Subhierarchies on EX Series Switches . . . . .	2699
Table 291: Unsupported [edit interfaces ae] Configuration Statements for EX Series Switches . . . . .	2704
Table 292: Unsupported [edit interfaces et] Configuration Statements for EX Series Switches . . . . .	2709
Table 293: Unsupported [edit interfaces ge] Configuration Statements for EX Series Switches . . . . .	2715

	Table 294: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches . . . . .	2720
	Table 295: Unsupported [edit interfaces irb] Configuration Statements for EX Series Switches . . . . .	2728
	Table 296: Unsupported [edit interfaces lo] Configuration Statements for EX Series Switches . . . . .	2731
	Table 297: Unsupported [edit interfaces me] Configuration Statements for EX Series Switches . . . . .	2734
	Table 298: Unsupported [edit interfaces xe] Configuration Statements for EX Series Switches . . . . .	2742
	Table 299: Protocol Families and Supported Interface Types . . . . .	2772
<b>Chapter 46</b>	<b>Administration . . . . .</b>	<b>2835</b>
	Table 300: Output Control Keys for the monitor interface Command . . . . .	2845
	Table 301: Output Control Keys for the monitor interface traffic Command . . . . .	2846
	Table 302: monitor interface Output Fields . . . . .	2847
	Table 303: request diagnostics tdr Output Fields . . . . .	2855
	Table 304: show diagnostics tdr Output Fields . . . . .	2857
	Table 305: show forwarding-options enhanced-hash-key Output Fields . . . . .	2861
	Table 306: show interfaces diagnostics optics Output Fields . . . . .	2864
	Table 307: show interfaces ge- Output Fields . . . . .	2879
	Table 308: show interfaces irb Output Fields . . . . .	2890
	Table 309: show interfaces mc-ae Output Fields . . . . .	2896
	Table 310: show interfaces me0 Output Fields . . . . .	2898
	Table 311: show interfaces queue Output Fields . . . . .	2905
	Table 312: show interfaces xe- Output Fields . . . . .	2912
	Table 313: show lacp interfaces Output Fields . . . . .	2926
<b>Part 15</b>	<b>BGP</b>	
<b>Chapter 48</b>	<b>Overview . . . . .</b>	<b>2939</b>
	Table 314: Supported Junos OS Layer 3 Protocol Statements and Features . . . . .	2939
	Table 315: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported . . . . .	2940
<b>Chapter 49</b>	<b>Configuration . . . . .</b>	<b>2943</b>
	Table 316: BGP Routing Configuration Summary . . . . .	2944
	Table 317: BGP Global Settings . . . . .	2945
<b>Chapter 50</b>	<b>Administration . . . . .</b>	<b>3035</b>
	Table 318: Summary of Key BGP Routing Output Fields . . . . .	3035
	Table 319: show bgp bmp Output Fields . . . . .	3043
	Table 320: show bgp group Output Fields . . . . .	3046
	Table 321: show bgp neighbor Output Fields . . . . .	3053
	Table 322: show bgp summary Output Fields . . . . .	3067
	Table 323: show policy damping Output Fields . . . . .	3073
<b>Part 16</b>	<b>IS-IS</b>	
<b>Chapter 51</b>	<b>Overview . . . . .</b>	<b>3077</b>
	Table 324: Supported Junos OS Layer 3 Protocol Statements and Features . . . . .	3077

	Table 325: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported . . . . .	3078
<b>Chapter 52</b>	<b>Configuration . . . . .</b>	<b>3081</b>
	Table 326: Default Metric Values for Routes Exported into IS-IS . . . . .	3114
<b>Chapter 53</b>	<b>Administration . . . . .</b>	<b>3141</b>
	Table 327: show isis adjacency Output Fields . . . . .	3151
	Table 328: show isis authentication Output Fields . . . . .	3154
	Table 329: show isis backup coverage Output Fields . . . . .	3156
	Table 330: show isis backup label-switched-path Output Fields . . . . .	3158
	Table 331: show isis backup spf results Output Fields . . . . .	3161
	Table 332: show isis database Output Fields . . . . .	3165
	Table 333: show isis hostname Output Fields . . . . .	3176
	Table 334: show isis interface Output Fields . . . . .	3179
	Table 335: show isis overview Output Fields . . . . .	3182
	Table 336: show isis route Output Fields . . . . .	3186
	Table 337: show isis spf Output Fields . . . . .	3189
	Table 338: show isis statistics Output Fields . . . . .	3195
<b>Part 17</b>	<b>OSPF</b>	
<b>Chapter 54</b>	<b>Overview . . . . .</b>	<b>3199</b>
	Table 339: Supported Junos OS Layer 3 Protocol Statements and Features . .	3199
	Table 340: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported . . . . .	3200
<b>Chapter 55</b>	<b>Configuration . . . . .</b>	<b>3207</b>
	Table 341: OSPF Routing Configuration Summary . . . . .	3208
	Table 342: Edit OSPF Global Settings . . . . .	3209
<b>Chapter 56</b>	<b>Administration . . . . .</b>	<b>3277</b>
	Table 343: Summary of Key OSPF Routing Output Fields . . . . .	3277
	Table 344: show (ospf   ospf3) interface Output Fields . . . . .	3291
	Table 345: show (ospf   ospf3) io-statistics Output Fields . . . . .	3296
	Table 346: show (ospf   ospf3) log Output Fields . . . . .	3298
	Table 347: show (ospf   ospf3) neighbor Output Fields . . . . .	3302
	Table 348: show ospf overview Output Fields . . . . .	3308
	Table 349: show (ospf   ospf3) route Output Fields . . . . .	3313
	Table 350: show (ospf   ospf3) statistics Output Fields . . . . .	3318
	Table 351: show ospf database Output Fields . . . . .	3323
	Table 352: show ospf3 database Output Fields . . . . .	3331
<b>Part 18</b>	<b>RIP</b>	
<b>Chapter 57</b>	<b>Overview . . . . .</b>	<b>3343</b>
	Table 353: Supported Junos OS Layer 3 Protocol Statements and Features . .	3343
	Table 354: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported . . . . .	3344
<b>Chapter 58</b>	<b>Configuration . . . . .</b>	<b>3347</b>

	Table 355: RIP Routing Configuration Summary . . . . .	3348
	Table 356: Edit RIP Global Settings . . . . .	3348
<b>Chapter 59</b>	<b>Administration . . . . .</b>	<b>3393</b>
	Table 357: Summary of Key RIP Routing Output Fields . . . . .	3393
	Table 358: show rip general-statistics Output Fields . . . . .	3397
	Table 359: show rip neighbor Output Fields . . . . .	3400
	Table 360: show rip statistics Output Fields . . . . .	3402
	Table 361: show ripng general-statistics Output Fields . . . . .	3406
	Table 362: show ripng neighbor Output Fields . . . . .	3408
	Table 363: show ripng statistics Output Fields . . . . .	3410
<b>Part 19</b>	<b>Routing Options</b>	
<b>Chapter 60</b>	<b>Overview . . . . .</b>	<b>3415</b>
	Table 364: Supported Junos OS Layer 3 Protocol Statements and Features . .	3415
	Table 365: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported . . . . .	3416
<b>Chapter 61</b>	<b>Configuration . . . . .</b>	<b>3421</b>
	Table 366: Static Routing Configuration Summary . . . . .	3423
<b>Chapter 62</b>	<b>Administration . . . . .</b>	<b>3537</b>
	Table 367: Filtering Route Messages . . . . .	3537
	Table 368: Summary of Key Routing Information Output Fields . . . . .	3538
	Table 369: show as-path Output Fields . . . . .	3543
	Table 370: show as-path domain Output Fields . . . . .	3546
	Table 371: show as-path summary Output Fields . . . . .	3549
	Table 372: show ipv6 neighbors Output Fields . . . . .	3551
	Table 373: show route Output Fields . . . . .	3554
	Table 374: show route damping Output Fields . . . . .	3578
	Table 375: show route detail Output Fields . . . . .	3582
	Table 376: Next-hop Types Output Field Values . . . . .	3587
	Table 377: State Output Field Values . . . . .	3588
	Table 378: Communities Output Field Values . . . . .	3590
	Table 379: show route export Output Fields . . . . .	3601
	Table 380: show route extensive Output Fields . . . . .	3604
	Table 381: show route flow validation Output Fields . . . . .	3621
	Table 382: show route instance Output Fields . . . . .	3629
	Table 383: show route martians Output Fields . . . . .	3641
	Table 384: show route receive-protocol Output Fields . . . . .	3668
	Table 385: show route resolution Output Fields . . . . .	3677
	Table 386: show route summary Output Fields . . . . .	3693
	Table 387: show route terse Output Fields . . . . .	3712
<b>Part 20</b>	<b>Multicast</b>	
<b>Chapter 64</b>	<b>Configuration . . . . .</b>	<b>3727</b>
	Table 388: IGMP Snooping Configuration Fields . . . . .	3736
	Table 389: Supported Tracing Operations for IGMP Snooping . . . . .	3737

<b>Chapter 65</b>	<b>Administration</b>	<b>3823</b>
	Table 390: Summary of IGMP Snooping Output Fields	3824
	Table 391: mtrace Output Fields	3849
	Table 392: mtrace from-source Output Fields	3853
	Table 393: mtrace monitor Output Fields	3855
	Table 394: mtrace to-gateway Output Fields	3858
	Table 395: show igmp group Output Fields	3860
	Table 396: show igmp interface Output Fields	3864
	Table 397: show igmp statistics Output Fields	3868
	Table 398: show igmp-snooping membership Output Fields	3871
	Table 399: show igmp-snooping route Output Fields	3874
	Table 400: show igmp-snooping statistics Output Fields	3876
	Table 401: show multicast flow-map Output Fields	3878
	Table 402: show multicast interface Output Fields	3880
	Table 403: show multicast minfo Output Fields	3882
	Table 404: show multicast next-hops Output Fields	3885
	Table 405: show multicast pim-to-igmp-proxy Output Fields	3887
	Table 406: show multicast pim-to-mld-proxy Output Fields	3889
	Table 407: show multicast route Output Fields	3892
	Table 408: show multicast rpf Output Fields	3899
	Table 409: show multicast scope Output Fields	3902
	Table 410: show multicast sessions Output Fields	3904
	Table 411: show multicast usage Output Fields	3908
	Table 412: show pim bootstrap Output Fields	3910
	Table 413: show pim interfaces Output Fields	3912
	Table 414: show pim join Output Fields	3917
	Table 415: show pim neighbors Output Fields	3930
	Table 416: show pim rps Output Fields	3934
	Table 417: show pim source Output Fields	3941
	Table 418: show pim statistics Output Fields	3944
<b>Part 21</b>	<b>Network Managment and Monitoring</b>	
<b>Chapter 66</b>	<b>Overview</b>	<b>3959</b>
	Table 419: Mirroring Terminologies	3965
	Table 420: Configuration Guidelines for Port Mirroring and Analyzers on EX4300 Switches	3967
	Table 421: Network Analytics CLI Changes	3982
	Table 422: GPB Stream Format Message Header Information	3986
	Table 423: Streamed Queue Statistics Data Output Fields	3989
	Table 424: Streamed Traffic Statistics Data Output Fields	3989
	Table 425: Output Fields for Queue Statistics in Local Analytics File	3991
	Table 426: Output Fields for Traffic Statistics in Local Analytics File	3992
<b>Chapter 67</b>	<b>Configuration</b>	<b>3995</b>
	Table 427: SNMP Configuration Page	4041
	Table 428: Port Mirroring Configuration Settings	4048
	Table 429: RPM Probe Owner, Concurrent Probes, and Probe Servers Configuration Fields	4052

**Chapter 68**

Table 430: Performance Probe Tests Configuration Fields . . . . .	4053
Table 431: Unsupported [edit snmp] Configuration Statements on EX Series Switches . . . . .	4083
Table 432: Unsupported [edit forwarding-options] Subhierarchies on EX Series Switches . . . . .	4156
<b>Administration . . . . .</b>	<b>4263</b>
Table 433: J-Web Ping Host Field Summary . . . . .	4264
Table 434: Traceroute field summary . . . . .	4266
Table 435: Match Conditions for the monitor traffic Command . . . . .	4272
Table 436: Logical Operators for the monitor traffic Command . . . . .	4273
Table 437: Arithmetic and Relational Operators for the monitor traffic Command . . . . .	4275
Table 438: show pfe statistics bridge Output Fields . . . . .	4284
Table 439: traceroute Output Fields . . . . .	4291
Table 440: show services rpm active-servers Output Fields . . . . .	4293
Table 441: show services rpm history-results Output Fields . . . . .	4294
Table 442: show services rpm probe-results Output Fields . . . . .	4297
Table 443: show snmp health-monitor Output Fields . . . . .	4312
Table 444: show snmp inform-statistics Output Fields . . . . .	4319
Table 445: show snmp mib Output Fields . . . . .	4322
Table 446: show snmp rmon Output Fields . . . . .	4324
Table 447: show smp rmon history Output Fields . . . . .	4328
Table 448: show snmp statistics Output Fields . . . . .	4332
Table 449: show snmp v3 Output Fields . . . . .	4337
Table 450: show forwarding-options port-mirroring Output Fields . . . . .	4339
Table 451: show forwarding-options analyzer Output Fields . . . . .	4341
Table 452: monitor start Command Output Fields . . . . .	4343
Table 453: show analytics collector Command Output Fields . . . . .	4346
Table 454: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D15 and Later) . . . . .	4348
Table 455: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D10 and earlier) . . . . .	4349
Table 456: show analytics queue-statistics Command Output Fields . . . . .	4352
Table 457: show analytics status Command Output Fields . . . . .	4354
Table 458: show analytics traffic-statistics Command Output Fields . . . . .	4358
Table 459: show sflow Output Fields . . . . .	4361
Table 460: show sflow interface Output Fields . . . . .	4363
Table 461: show sflow collector Output Fields . . . . .	4365
Table 462: monitor ethernet delay-measurement one-way Output Fields . . .	4373
Table 463: monitor ethernet delay-measurement two-way Output Fields . . .	4373
Table 464: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields . . . . .	4377
Table 465: show oam ethernet connectivity-fault-management forwarding-state Output Fields . . . . .	4380
Table 466: show oam ethernet connectivity-fault-management interfaces Output Fields . . . . .	4384
Table 467: show oam ethernet connectivity-fault-management linktrace path-database Output Fields . . . . .	4390

	Table 468: show oam ethernet connectivity-fault-management mep-database Output Fields . . . . .	4393
	Table 469: show oam ethernet connectivity-fault-management mip Output Fields . . . . .	4399
	Table 470: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields . . . . .	4400
	Table 471: show oam ethernet link-fault-management Output Fields . . . . .	4405
	Table 472: show uplink-failure-detection Output Fields . . . . .	4410
<b>Chapter 69</b>	<b>Troubleshooting . . . . .</b>	<b>4413</b>
	Table 473: show pfe statistics bridge Output Fields . . . . .	4415
<b>Part 22</b>	<b>PoE</b>	
<b>Chapter 70</b>	<b>Overview . . . . .</b>	<b>4423</b>
	Table 474: PoE Version Support . . . . .	4424
	Table 475: PoE Power Budget for EX2200 Switches . . . . .	4426
	Table 476: PoE Power Budget for EX3200 Switch Models . . . . .	4426
	Table 477: PoE Power Budget EX3300 Switch Models . . . . .	4427
	Table 478: PoE Power Budget for EX4200 Switch Models . . . . .	4427
	Table 479: PoE Power Budget for EX4300 Switch Models . . . . .	4428
	Table 480: Class of Powered Device and Power Levels . . . . .	4430
	Table 481: Maximum Power per Port in Static Mode . . . . .	4431
	Table 482: LLDP Power Priority Conversion . . . . .	4432
<b>Chapter 71</b>	<b>Configuration . . . . .</b>	<b>4433</b>
	Table 483: Components of the PoE Configuration Topology . . . . .	4434
	Table 484: Components of the PoE Configuration Topology . . . . .	4436
	Table 485: Configurable PoE Options and Default Settings . . . . .	4441
	Table 486: PoE Edit Settings . . . . .	4448
	Table 487: System Settings . . . . .	4448
	Table 488: Edit PoE Settings . . . . .	4449
	Table 489: FPC PoE Settings . . . . .	4449
<b>Chapter 72</b>	<b>Administration . . . . .</b>	<b>4467</b>
	Table 490: show poe controller Output Fields . . . . .	4481
	Table 491: show poe interface Output Fields . . . . .	4484
	Table 492: show poe notification-control Output Fields . . . . .	4488
	Table 493: show poe telemetries interface Output Fields . . . . .	4490
<b>Chapter 73</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>4493</b>
	Table 494: Troubleshooting a PoE Interface . . . . .	4493
<b>Part 23</b>	<b>Port Security</b>	
<b>Chapter 74</b>	<b>Overview . . . . .</b>	<b>4497</b>
	Table 495: DHCPv6 Messages and DHCPv4 Equivalent Messages . . . . .	4509
<b>Chapter 75</b>	<b>Configuration . . . . .</b>	<b>4529</b>
	Table 496: Components of the Port Security Topology . . . . .	4531
	Table 497: Components of the Port Security Topology . . . . .	4536

	Table 498: Port Security Settings on VLANs . . . . .	4542
	Table 499: Port Security on Interfaces . . . . .	4544
	Table 500: Unsupported [edit vlans] Configuration Statements on EX Series Switches . . . . .	4584
<b>Chapter 76</b>	<b>Administration . . . . .</b>	<b>4657</b>
	Table 501: show dhcp-security arp inspection statistics Output Fields . . . . .	4666
	Table 502: show dhcp-security binding Output Fields . . . . .	4668
	Table 503: show dhcp-security binding ip-source-guard Output Fields . . . . .	4671
	Table 504: show dhcp-security ipv6 binding Output Fields . . . . .	4673
	Table 505: show dhcp-security ipv6 statistics Output Fields . . . . .	4676
	Table 506: show dhcp-security nd inspection statistics Output Fields . . . . .	4678
	Table 507: show security macsec connections Output Fields . . . . .	4680
	Table 508: show security macsec statistics Output Fields . . . . .	4682
	Table 509: show security mka sessions Output Fields . . . . .	4686
	Table 510: show security mka statistics Output Fields . . . . .	4688
<b>Part 24</b>	<b>Routing Policy and Packet Filtering</b>	
<b>Chapter 77</b>	<b>Overview . . . . .</b>	<b>4693</b>
	Table 511: Elements of a Firewall Filter Configuration . . . . .	4705
	Table 512: Firewall Filter Match Conditions Supported on EX Series Switches . .	4705
	Table 513: Actions for Firewall Filters . . . . .	4712
	Table 514: Action Modifiers for Firewall Filters . . . . .	4713
	Table 515: Match Conditions for Firewall Filters on Loopback Interfaces for IPv4 and IPv6 Traffic—Support per Switch . . . . .	4716
	Table 516: Actions for Firewall Filters on Loopback Interfaces for IPv4 and IPv6 Traffic—Support per Switch . . . . .	4717
	Table 517: Action Modifiers for Firewall Filters on Loopback Interfaces for IPv4 and IPv6 Traffic—Support per Switch . . . . .	4718
	Table 518: Bind Points Associated with Firewall Filter Types . . . . .	4719
	Table 519: Support for IPv4 and IPv6 Firewall Filters on Switches . . . . .	4720
	Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches . . . . .	4721
	Table 521: Firewall Filter Match Conditions Supported for IPv6 Traffic on Switches . . . . .	4734
	Table 522: Firewall Filter Match Condition Supported for Non-IP Traffic on Switches . . . . .	4742
	Table 523: Firewall Filter Actions Supported for IPv4 Traffic on Switches . . .	4743
	Table 524: Firewall Filter Actions Supported for IPv6 Traffic on Switches . . .	4746
	Table 525: Firewall Filter Action Modifiers Supported for IPv4 Traffic on Switches . . . . .	4749
	Table 526: Firewall Filter Action Modifiers Supported for IPv6 Traffic on Switches . . . . .	4755
	Table 527: Logical Operators for Matching Multiple Bit-Field Operators . . . .	4765
	Table 528: Policer Actions . . . . .	4768
<b>Chapter 78</b>	<b>Configuration . . . . .</b>	<b>4773</b>
	Table 529: Configuration Components: Firewall Filters . . . . .	4774
	Table 530: Configuration Components: VLANs . . . . .	4776



	Table 531: Configuration Components: Switch Ports on a 48-Port All-PoE Switch . . . . .	4777
	Table 532: Create a New Filter . . . . .	4814
	Table 533: Create a New Term . . . . .	4814
	Table 534: Advanced Options for Terms . . . . .	4815
	Table 535: Policies Global Configuration Parameters . . . . .	4823
	Table 536: Terms Configuration Parameters . . . . .	4824
	Table 537: Supported Options for Firewall Filter Statements . . . . .	4832
	Table 538: Firewall Filter Statements That Are Not Supported by Junos OS for EX Series Switches . . . . .	4834
<b>Chapter 79</b>	<b>Administration . . . . .</b>	<b>4873</b>
	Table 539: show firewall Output Fields . . . . .	4880
	Table 540: show firewall log Output Fields . . . . .	4886
	Table 541: show policer Output Fields . . . . .	4889
	Table 542: show policy Output Fields . . . . .	4891
	Table 543: show policy conditions Output Fields . . . . .	4893
<b>Part 25</b>	<b>Spanning-Tree Protocols</b>	
<b>Chapter 81</b>	<b>Overview . . . . .</b>	<b>4903</b>
	Table 544: Selecting a Spanning-Tree Protocol . . . . .	4905
	Table 545: Selecting a Spanning-Tree Protocol . . . . .	4909
	Table 546: Selecting a Spanning-Tree Protocol . . . . .	4911
	Table 547: Selecting a Spanning-Tree Protocol . . . . .	4913
<b>Chapter 82</b>	<b>Configuration . . . . .</b>	<b>4919</b>
	Table 548: Components of the Topology for Configuring RSTP . . . . .	4921
	Table 549: Components of the Topology for Configuring MSTP on EX Series Switches . . . . .	4938
	Table 550: Components of the Topology for Configuring BPDU Protection on EX Series Switches . . . . .	4958
	Table 551: Components of the Topology for Configuring BPDU Protection on EX Series Switches . . . . .	4963
	Table 552: Components of the Topology for Configuring Loop Protection on EX Series Switches . . . . .	4968
	Table 553: Components of the Topology for Configuring Root Protection on EX Series Switches . . . . .	4973
	Table 554: Spanning-Tree Protocol Configuration Parameters . . . . .	4978
<b>Chapter 83</b>	<b>Administration . . . . .</b>	<b>5027</b>
	Table 555: Summary of Spanning-Tree Protocols Output Fields . . . . .	5027
	Table 556: show spanning-tree bridge Output Fields . . . . .	5035
	Table 557: show spanning-tree bridge Output Fields . . . . .	5040
	Table 558: show spanning-tree interface Output Fields . . . . .	5046
	Table 559: show spanning-tree Interface Output Fields . . . . .	5049
	Table 560: show spanning-tree mstp configuration Output Fields . . . . .	5055
	Table 561: show spanning-tree mstp configuration Output Fields . . . . .	5057
	Table 562: show spanning-tree statistics Output Fields . . . . .	5059
	Table 563: show spanning-tree statistics Output Fields . . . . .	5061

<b>Part 26</b>	<b>Virtual Chassis</b>	
<b>Chapter 84</b>	<b>Overview</b>	<b>5065</b>
	Table 564: Maximum Member Switch Support for Virtual Chassis by Junos OS Release	5067
	Table 565: Virtual Chassis Fabric Summary	5074
	Table 566: Virtual Chassis Summary	5076
	Table 567: Maximum Member Links in LAGs Over Optical Interface VCPs	5084
	Table 568: Automatic Software Update Support	5092
	Table 569: Effects of a Routing Engine Switchover	5095
<b>Chapter 86</b>	<b>Administration</b>	<b>5149</b>
	Table 570: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration	5150
	Table 571: show system uptime Output Fields	5171
	Table 572: show virtual-chassis active-topology Output Fields	5174
	Table 573: show virtual-chassis device-topology Output Fields	5179
	Table 574: show virtual-chassis protocol adjacency Output Fields	5187
	Table 575: show virtual-chassis protocol database Output Fields	5190
	Table 576: show virtual-chassis protocol interface Output Fields	5195
	Table 577: show virtual-chassis protocol route Output Fields	5197
	Table 578: show virtual-chassis protocol statistics Output Fields	5200
	Table 579: show virtual-chassis Output Fields	5203
	Table 580: show virtual-chassis vc-path Output Fields	5207
	Table 581: show virtual-chassis vc-port Output Fields	5209
	Table 582: show virtual-chassis vc-port diagnostics optics Output Fields	5214
	Table 583: show virtual-chassis vc-port statistics Output Fields	5228
<b>Part 27</b>	<b>Virtual Chassis Fabric</b>	
<b>Chapter 88</b>	<b>Overview</b>	<b>5241</b>
	Table 584: Virtual Chassis Fabric Summary	5255
	Table 585: Virtual Chassis Summary	5256
<b>Chapter 90</b>	<b>Administration</b>	<b>5313</b>
	Table 586: show forwarding-options enhanced-hash-key Output Fields	5327
	Table 587: show virtual-chassis active-topology Output Fields	5330
	Table 588: show virtual-chassis device-topology Output Fields	5335
	Table 589: show virtual-chassis mode Output Fields	5342
	Table 590: show virtual-chassis protocol adjacency Output Fields	5346
	Table 591: show virtual-chassis protocol database Output Fields	5349
	Table 592: show virtual-chassis protocol interface Output Fields	5354
	Table 593: show virtual-chassis protocol route Output Fields	5356
	Table 594: show virtual-chassis protocol statistics Output Fields	5359
	Table 595: show virtual-chassis Output Fields	5362
	Table 596: show virtual-chassis vc-port Output Fields	5366
	Table 597: show virtual-chassis vc-port diagnostics optics Output Fields	5371
	Table 598: show virtual-chassis vc-port statistics Output Fields	5385

# About the Documentation

- Documentation and Release Notes on page cix
- Supported Platforms on page cix
- Using the Examples in This Manual on page cix
- Documentation Conventions on page cxi
- Documentation Feedback on page cxiii
- Requesting Technical Support on page cxiii

## Documentation and Release Notes

---

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

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

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

## Supported Platforms

---

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

- EX Series

## Using the Examples in This Manual

---

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

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

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

## Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

## Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

## Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

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

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

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

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

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

## Documentation Feedback

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

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

## Requesting Technical Support

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

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

## Self-Help Online Tools and Resources

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

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

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

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

## Opening a Case with JTAC

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

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

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



## PART 1

# Getting Started with ELS

- [Getting Started with ELS on page 3](#)



## CHAPTER 1

# Getting Started with ELS

- [Getting Started on page 3](#)

## Getting Started

---

- [Getting Started with Enhanced Layer 2 Software on page 3](#)
- [Enhanced Layer 2 CLI Configuration Statement and Command Changes on page 17](#)

## Getting Started with Enhanced Layer 2 Software

- [Understanding Enhanced Layer 2 Software Support on page 3](#)
- [Using the ELS Translator Tool on page 4](#)
- [Configuring a VLAN on page 5](#)
- [Configuring the Native VLAN Identifier on page 6](#)
- [Configuring Layer 2 Interfaces on page 6](#)
- [Configuring Layer 3 Interfaces on page 6](#)
- [Configuring an IRB Interface on page 7](#)
- [Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface on page 7](#)
- [Enhanced Layer 2 CLI Configuration Statement and Command Changes on page 8](#)

## Understanding Enhanced Layer 2 Software Support

---

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

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

**Table 3: ELS Support**

Device	Initial ELS Release
EX4300 switches	13.2X50-D10
EX4600 switches	13.2X51-D25

Table 3: ELS Support (*continued*)

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

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

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

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



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

### Using the ELS Translator Tool

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

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

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

1. Log onto your device using the console port.



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

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



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

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



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

## Configuring a VLAN

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

To configure a VLAN:

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

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

2. Assign at least one interface to the VLAN:

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

---

### Configuring the Native VLAN Identifier

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

To configure the native VLAN ID:

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

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

2. Configure the native VLAN ID:

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

3. Assign the interface to the native VLAN ID:

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

---

### Configuring Layer 2 Interfaces

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

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

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

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

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

---

### Configuring Layer 3 Interfaces

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

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

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

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

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

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

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

### Configuring an IRB Interface

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

To configure an IRB interface:

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

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

2. Create an IRB logical interface:

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

3. Associate the IRB interface with the VLAN:

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

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

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

To configure an aggregated Ethernet interface:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

To configure LACP:

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

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

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

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

---

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

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





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

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

**Table 4: Enhanced Layer 2 CLI Changes**

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

Table 4: Enhanced Layer 2 CLI Changes (*continued*)

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

Table 4: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   secure-access-port {     interface (all   <i>interface-name</i>) {       (dhcp-trusted   no-dhcp-trusted );       static-ip <i>ip-address</i> {         mac <i>mac-address</i>;         vlan <i>vlan-name</i>;       }     }   }   vlan (all   <i>vlan-name</i>) {     (arp-inspection   no-arp-inspection );     dhcp-option82 {       disable;       circuit-id {         prefix <i>hostname</i>;         use-interface-description;         use-vlan-id;       }       remote-id {         prefix (<i>hostname</i>   mac   none);         use-interface-description;         use-string <i>string</i>;       }       vendor-id [<i>string</i>];     }     (examine-dhcp   no-examine-dhcp);   }   (ip-source-guard   no-ip-source-guard); } </pre>	<pre> vlans <i>vlan-name</i> forwarding-options{   dhcp-security {     arp-inspection;     group <i>group-name</i> {       interface <i>interface-name</i> {         static-ip <i>ip-address</i> {           mac <i>mac-address</i>;         }       }     }     overrides {       no-option-82;       trusted;     }   }   ip-source-guard;   no-dhcp-snooping;   option-82 {     circuit-id {       prefix {         host-name;         routing-instance-name;       }       use-interface-description (device           logical);       use-vlan-id;     }     remote-id {       host-name;       use-interface-description (device           logical);       use-string <i>string</i>;     }     vendor-id {       use-string <i>string</i>;     }   } } </pre>	<p>Statements moved to different hierarchy.</p> <p><b>NOTE:</b> The statement <b>examine-dhcp</b> does not exist in the changed hierarchy. Instead, DHCP snooping is enabled automatically when other DHCP security features are enabled on a VLAN. See “<a href="#">Configuring Port Security (CLI Procedure)</a>” on page 4539 for additional information.</p>
<pre> ethernet-switching-options {   secure-access-port {     dhcp-snooping-file {       location <i>local_pathname</i>   <i>remote_URL</i>;       timeout <i>seconds</i>;       write-interval <i>seconds</i>;     }   } } </pre>	<pre> system [   processes [     dhcp-service     dhcp-snooping-file <i>local_pathname</i>         <i>remote_URL</i>;     write-interval <i>interval</i>;   ] } </pre>	<p>Statement moved to different hierarchy.</p>
<pre> ethernet-switching-options {   secure-access-port vlan (all   <i>vlan-name</i>) {     mac-move-limit   } } </pre>	<pre> vlans <i>vlan-name</i> switch-options {   mac-move-limit } </pre>	<p>Statement moved to different hierarchy.</p>

Table 4: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   static {     vlan <i>vlan-id</i> {       mac <i>mac-address</i> next-hop         <i>interface-name</i>;       ...     }   } } </pre>	<pre> vlangs {   <i>vlan-name</i> {     switch-options {       interface <i>interface-name</i> {         static-mac <i>mac-address</i>;         ...       }     }   } } </pre>	Statement replaced with new statement and moved to different hierarchy.
<pre> ethernet-switching-options {   storm-control {     (...)   } } </pre>	<pre> forwarding-options {   storm-control-profiles <i>profile-name</i> {     (...)   } }  interfaces <i>interface-name</i> unit <i>number</i> family <i>family</i>   ethernet-switching {     storm-control <i>storm-control-profile</i>;   } </pre>	Storm control configuration is done in two steps. The first step is to create a storm control profile at the [edit forwarding-options] hierarchy, and the second step is to bind the profile to a logical interface at the [edit interfaces] hierarchy. See <a href="#">“Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches”</a> on page 2201 for additional information.
<pre> ethernet-switching-options {   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt;       &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable           no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;;     ...   } } </pre>	—	Statements removed.
<pre> ethernet-switching-options {   unknown-unicast-forwarding {     (...)   } } </pre>	<pre> switch-options on page 2212 {   unknown-unicast-forwarding {     (...)   } } </pre>	Hierarchy renamed.

Table 4: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   voip {     interface (all   [<i>interface-name</i>         access-ports]) {       forwarding-class (assured-forwarding           best-effort   expedited-forwarding           network-control);       vlan <i>vlan-name</i>;       ...     }   } } </pre>	<pre> switch-options {   voip {     interface (all   [<i>interface-name</i>         access-ports]) {       forwarding-class (assured-forwarding           best-effort   expedited-forwarding           network-control);       vlan <i>vlan-name</i>;       ...     }   } } </pre>	Hierarchy renamed.
<pre> interfaces <i>interface-name</i> {   ether-options {     link-mode <i>mode</i>;     speed (auto-negotiation   <i>speed</i>)   } } </pre>	<pre> interfaces <i>interface-name</i> {   link-mode <i>mode</i>;   speed <i>speed</i>) } </pre>	Statements moved to different hierarchy.
<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       native-vlan-id <i>vlan-id</i>     }   } } </pre>	<pre> interfaces <i>interface-name</i> {   native-vlan-id <i>vlan-id</i> } </pre>	Statement moved to different hierarchy.
<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       port-mode <i>mode</i>     }   } } </pre>	<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       interface-mode <i>mode</i>     }   } } </pre>	Statement replaced with a new statement.
<pre> interfaces vlan </pre>	<pre> interfaces irb </pre>	Statement replaced with a new statement.

Table 4: Enhanced Layer 2 CLI Changes (*continued*)

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

Table 4: Enhanced Layer 2 CLI Changes (*continued*)

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

Table 4: Enhanced Layer 2 CLI Changes (*continued*)

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



Table 4: Enhanced Layer 2 CLI Changes (*continued*)

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

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

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

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

Table 5: Enhanced Layer 2 CLI Changes

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   analyzer {     name {       ...     }   } } </pre>	<pre> forwarding-options {   analyzer {     name {       ...     }   } } </pre>	Statements moved to different hierarchy.
<pre> ethernet-switching-options {   authentication-whitelist {     ...   } } </pre>	<pre> switch-options {   ...   authentication-whitelist {     ...   } } </pre>	Hierarchy renamed.

Table 5: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   bpdu-block {     ...   } } </pre>	<pre> protocols {   layer2-control {     bpdu-block {       ...     }   } } </pre>	Statement moved to different hierarchy.
<pre> ethernet-switching-options {   dot1q-tunneling {     ether-type (0x8100   0x88a8   0x9100);     ...   } } </pre>	<pre> interfaces interface-name {   ether-options {     ethernet-switch-profile {       tag-protocol-id [tpids];     }   } }  interfaces interface-name {   aggregated-ether-options {     ethernet-switch-profile {       tag-protocol-id [tpids];     }   } } </pre>	Statement replaced with new statement and moved to different hierarchy.
<pre> ethernet-switching-options {   interfaces interface-name {     no-mac-learning;     ...   } } </pre>	<pre> switch-options {   interfaces interface-name {     no-mac-learning;     ...   } } </pre>	Hierarchy renamed.
<pre> ethernet-switching-options {   mac-notification {     notification-interval seconds;     ...   } } </pre>	—	Statements deleted.
<pre> ethernet-switching-options {   mac-table-aging-time seconds;   ... } </pre>	<pre> protocols {   l2-learning {     global-mac-table-aging-time seconds;     ...   } } </pre>	Statement replaced with new statement and moved to different hierarchy.
<pre> ethernet-switching-options {   nonstop-bridging; } </pre>	<pre> protocols {   layer2-control {     nonstop-bridging {     }   } } </pre>	Statement moved to different hierarchy.

Table 5: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   port-error-disable {     disable-timeout <i>timeout</i>;     ...   } } </pre>	<pre> interfaces <i>interface-name</i> family   ethernet-switching {     <i>recovery-timeout</i> <i>seconds</i>;   } </pre>	Statement replaced with a new statement.
<pre> ethernet-switching-options {   <i>redundant-trunk-group</i> {     group <i>name</i> {       description;       interface <i>interface-name</i> {         primary;       }       preempt-cutover-timer <i>seconds</i>;       ...     }   } } </pre>	<pre> switch-options {   <i>redundant-trunk-group</i> {     group <i>name</i> {       description;       interface <i>interface-name</i> {         primary;       }       preempt-cutover-timer <i>seconds</i>;       ...     }   } } </pre>	Hierarchy renamed.

Table 5: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   secure-access-port {     interface (all   interface-name) {       (dhcp-trusted   no-dhcp-trusted );       static-ip ip-address {         mac mac-address;         vlan vlan-name;       }     }   }   vlan (all   vlan-name) {     (arp-inspection   no-arp-inspection );     dhcp-option82 {       disable;       circuit-id {         prefix hostname;         use-interface-description;         use-vlan-id;       }       remote-id {         prefix (hostname   mac   none);         use-interface-description;         use-string string;       }       vendor-id [string];     }     (examine-dhcp   no-examine-dhcp);   }   (ip-source-guard   no-ip-source-guard); } </pre>	<pre> vlans vlan-name forwarding-options{   dhcp-security {     arp-inspection;     group group-name {       interface interface-name {         static-ip ip-address {           mac mac-address;         }       }     }     overrides {       no-option-82;       trusted;     }   }   ip-source-guard;   no-dhcp-snooping;   option-82 {     circuit-id {       prefix {         host-name;         routing-instance-name;       }       use-interface-description (device           logical);       use-vlan-id;     }     remote-id {       host-name;       use-interface-description (device           logical);       use-string string;     }     vendor-id {       use-string string;     }   } } </pre>	<p>Statements moved to different hierarchy.</p> <p><b>NOTE:</b> The statement <b>examine-dhcp</b> does not exist in the changed hierarchy. Instead, DHCP snooping is enabled automatically when other DHCP security features are enabled on a VLAN. See “<a href="#">Configuring Port Security (CLI Procedure)</a>” on page 4539 for additional information.</p>
<pre> ethernet-switching-options {   secure-access-port {     dhcp-snooping-file {       location local_pathname   remote_URL;       timeout seconds;       write-interval seconds;     }   } } </pre>	<pre> system [   processes [     dhcp-service     dhcp-snooping-file local_pathname         remote_URL;     write-interval interval;   ] ] </pre>	<p>Statement moved to different hierarchy.</p>
<pre> ethernet-switching-options {   secure-access-port vlan (all   vlan-name {     mac-move-limit   } } </pre>	<pre> vlans vlan-name switch-options {   mac-move-limit } </pre>	<p>Statement moved to different hierarchy.</p>

Table 5: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   static {     vlan <i>vlan-id</i> {       mac <i>mac-address</i> next-hop         <i>interface-name</i>;       ...     }   } } </pre>	<pre> vlangs {   <i>vlan-name</i> {     switch-options {       interface <i>interface-name</i> {         static-mac <i>mac-address</i>;         ...       }     }   } } </pre>	Statement replaced with new statement and moved to different hierarchy.
<pre> ethernet-switching-options {   storm-control {     (...)   } } </pre>	<pre> forwarding-options {   storm-control-profiles <i>profile-name</i> {     (...)   } }  interfaces <i>interface-name</i> unit <i>number</i> family   ethernet-switching {     storm-control <i>storm-control-profile</i>;   } </pre>	Storm control configuration is done in two steps. The first step is to create a storm control profile at the [edit forwarding-options] hierarchy, and the second step is to bind the profile to a logical interface at the [edit interfaces] hierarchy. See <a href="#">“Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches”</a> on <a href="#">page 2201</a> for additional information.
<pre> ethernet-switching-options {   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt;       &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable           no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;;     ...   } } </pre>	—	Statements removed.
<pre> ethernet-switching-options {   unknown-unicast-forwarding {     (...)   } } </pre>	<pre> switch-options on page 2212 {   unknown-unicast-forwarding {     (...)   } } </pre>	Hierarchy renamed.

Table 5: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   voip {     interface (all   [<i>interface-name</i>         access-ports]) {       forwarding-class (assured-forwarding           best-effort   expedited-forwarding           network-control);       vlan <i>vlan-name</i>;       ...     }   } } </pre>	<pre> switch-options {   voip {     interface (all   [<i>interface-name</i>         access-ports]) {       forwarding-class (assured-forwarding           best-effort   expedited-forwarding           network-control);       vlan <i>vlan-name</i>;       ...     }   } } </pre>	Hierarchy renamed.
<pre> interfaces <i>interface-name</i> {   ether-options {     link-mode <i>mode</i>;     speed (auto-negotiation   <i>speed</i>)   } } </pre>	<pre> interfaces <i>interface-name</i> {   link-mode <i>mode</i>;   speed <i>speed</i>) } </pre>	Statements moved to different hierarchy.
<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       native-vlan-id <i>vlan-id</i>     }   } } </pre>	<pre> interfaces <i>interface-name</i> {   native-vlan-id <i>vlan-id</i> } </pre>	Statement moved to different hierarchy.
<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       port-mode <i>mode</i>     }   } } </pre>	<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       interface-mode <i>mode</i>     }   } } </pre>	Statement replaced with a new statement.
<pre> interfaces vlan </pre>	<pre> interfaces irb </pre>	Statement replaced with a new statement.

Table 5: Enhanced Layer 2 CLI Changes (*continued*)

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

Table 5: Enhanced Layer 2 CLI Changes (*continued*)

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



Table 5: Enhanced Layer 2 CLI Changes (*continued*)

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

Table 5: Enhanced Layer 2 CLI Changes (*continued*)

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

**Related Documentation**

- [EX4300 Switches Hardware Overview](#)

## PART 2

# System Setup

- [Overview on page 29](#)
- [Configuration on page 33](#)
- [Administration on page 79](#)
- [Troubleshooting Procedures on page 305](#)



## CHAPTER 2

# Overview

- [Software Overview on page 29](#)

## Software Overview

---

- [Understanding Software Infrastructure and Processes on page 29](#)

### Understanding Software Infrastructure and Processes

Each switch runs the Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches on its general-purpose processors. Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the chassis.

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

With the J-Web interface and the command-line interface (CLI) to the Junos OS, you configure switching features and routing protocols and set the properties of network interfaces on your switch. After activating a software configuration, use either the J-Web or CLI user interface to monitor the switch, manage operations, and diagnose protocol and network connectivity problems.

- [Routing Engine and Packet Forwarding Engine on page 29](#)
- [Junos OS Processes on page 30](#)

### Routing Engine and Packet Forwarding Engine

---

A switch has two primary software processing components:

- **Packet Forwarding Engine**—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.
- **Routing Engine**—Provides three main functions:
  - Creates the packet forwarding switch fabric for the switch, providing route lookup, filtering, and switching on incoming data packets, then directing outbound packets to the appropriate interface for transmission to the network

- Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
- Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

### Junos OS Processes

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

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

Table 6 on page 30 describes the primary Junos OS processes.

**Table 6: Junos OS Processes**

Process	Name	Description
Chassis process	chassisd	<p>Detects hardware on the system that is used to configure network interfaces.</p> <p>Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered.</p> <p>Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.</p>
Ethernet switching process	eswd	<p>Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree protocol and access port security. The process is also responsible for managing Ethernet switching interfaces, VLANs, and VLAN interfaces.</p> <p>Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.</p>
Forwarding process	pfem	<p>Defines how routing protocols operate on the switch. The overall performance of the switch is largely determined by the effectiveness of the forwarding process.</p>
Interface process	dcd	<p>Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.</p>
Management process	mgd	<p>Provides communication between the other processes and an interface to the configuration database.</p> <p>Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured.</p> <p>Interacts with the other processes when commands are issued through one of the user interfaces on the switch.</p> <p>If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.</p>
Routing protocol process	rpd	<p>Defines how routing protocols such as RIP, OSPF, and BGP operate on the device, including selecting routes and maintaining forwarding tables.</p>

- Related Documentation**
- For more information about processes, see *Junos OS Network Operations Guide*
  - For more information about basic system parameters, supported protocols, and software processes, see *Junos OS System Basics Configuration Guide*





## CHAPTER 3

# Configuration

- [Initial Configuration on page 33](#)
- [Configuration Statements on page 37](#)

### Initial Configuration

---

- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\) on page 33](#)
- [Configuring Date and Time for the EX Series Switch \(J-Web Procedure\) on page 35](#)
- [Configuring System Identity for an EX Series Switch \(J-Web Procedure\) on page 36](#)

### Configuring the LCD Panel on EX Series Switches (CLI Procedure)

This topic applies to hardware devices in the EX Series product family, which includes switches and the XRE200 External Routing Engine, that support the LCD panel interface.

The LCD panel on the front panel of EX Series switches displays a variety of information about the switch in the Status menu and provides the Maintenance menu to allow you to perform basic operations such as initial setup and reboot. You can disable these menus or individual menu options if you do not want switch users to use them. You can also set a custom message that will be displayed on the panel.

This topic describes:

- [Disabling or Enabling Menus and Menu Options on the LCD Panel on page 33](#)
- [Configuring a Custom Display Message on page 34](#)

#### Disabling or Enabling Menus and Menu Options on the LCD Panel

---

By default, the Maintenance menu, the Status menu, and the options in those menus in the LCD panel are enabled. Users can configure and troubleshoot the switch using the Maintenance menu and view certain details about the switch using the Status menu.

If you do not want users to be able to use those menus or use some of the menu options, you can disable the menus or individual menu options. You can re-enable the menus or menu options.

Issue the `show chassis lcd` operational mode command to see which menus and menu options are currently enabled.



**NOTE:** On some platforms you must specify an FPC slot number in these commands. See the [lcd-menu](#) statement for details.

To disable a menu:

```
[edit]
user@switch# set chassis lcd-menu menu-item menu-name disable
```

To enable a menu:

```
[edit]
user@switch# delete chassis lcd-menu menu-item menu-name disable
```

To disable a menu option:

```
[edit]
user@switch# set chassis lcd-menu menu-item menu-option disable
```

To enable a menu option:

```
[edit]
user@switch# delete chassis lcd-menu menu-item menu-option disable
```

### Configuring a Custom Display Message

---

You can configure the second line of the LCD to display a custom message temporarily for 5 minutes or permanently.

To display a custom message temporarily:

- On an EX3200 switch, a standalone EX3300 switch, a standalone EX4200 switch, a standalone EX4500 switch, an EX8200 switch, or an XRE200 External Routing Engine:

```
user@switch> set chassis display message message
```

- On an EX3300, EX4200, or EX4500 switch in a Virtual Chassis configuration:

```
user@switch> set chassis display message message fpc-slot slot-number
```

To display a custom message permanently:

- On an EX3200 switch, a standalone EX3300 switch, a standalone EX4200 switch, a standalone EX4500 switch, an EX8200 switch, or an XRE200 External Routing Engine:

```
user@switch> set chassis display message message permanent
```

- On an EX3300, EX4200, or EX4500 switch in a Virtual Chassis configuration:

```
user@switch> set chassis display message message fpc-slot slot-number permanent
```



**NOTE:** The buttons on the LCD panel are disabled when the LCD is configured to display a custom message.

To disable the display of the custom message:

```
user@switch> clear chassis display message
```

You can view the custom message by issuing the command [show chassis lcd](#).

- Related Documentation
- *LCD Panel in EX3200 Switches*
  - *LCD Panel in EX3300 Switches*
  - *LCD Panel in EX4200 Switches*
  - *LCD Panel in EX4500 Switches*
  - *LCD Panel in an EX6200 Switch*
  - *LCD Panel in an EX8200 Switch*
  - *LCD Panel in an XRE200 External Routing Engine*

Configuring Date and Time for the EX Series Switch (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

To configure date and time on an EX Series switch:

1. Select **Configure > System Properties > Date & Time**.
2. To modify the information, click **Edit**. Enter information into the Edit Date & Time page as described in [Table 7 on page 35](#).
3. Click one of the following options:
  - To apply the configuration, click **OK**.
  - To cancel your entries and return to the System Properties page, click **Cancel**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

Table 7: Date and Time Settings

Time	Function	Your Action
Time Zone	Identifies the timezone that the switching platform is located in.	Select the appropriate time zone from the list.

Table 7: Date and Time Settings (*continued*)

Time	Function	Your Action
Set Time	Synchronizes the system time with that of the NTP server. You can also manually set the system time and date.	<p>To immediately set the time, Click one of the following options:</p> <ul style="list-style-type: none"> <li>• <b>Synchronize with PC time</b>—The switch synchronizes the time with that of the PC.</li> <li>• <b>NTP Servers</b>—The switch sends a request to the NTP server and synchronizes the system time.</li> <li>• <b>Manual</b>—A pop-up window allows you to select the current date and time from a list.</li> </ul>

**Related Documentation**

- *J-Web User Interface for EX Series Switches Overview*

## Configuring System Identity for an EX Series Switch (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

To configure identification details for an EX Series switch:

1. Select **Configure > System Properties > System Identity**. The System Identity page displays configuration details.
2. To modify the configuration, click **Edit**. Enter information into the System Identity page as described in [Table 8 on page 36](#).



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

Table 8: System Identity Configuration Summary

Field	Function	Your Action
Host Name	Defines the hostname of the switching platform.	Type the hostname.
Domain Name	Defines the network or subnetwork that the machine belongs to.	Type the domain name.

Table 8: System Identity Configuration Summary (*continued*)

Field	Function	Your Action
Root Password	Sets the root password that user <i>root</i> can use to log in to the switching platform.	Type a plain-text password. The system encrypts the password.  <b>NOTE:</b> After a root password has been defined, it is required when you log in to the J-Web user interface or the CLI.
Confirm Root Password	Verifies that the root password has been typed correctly.	Retype the password.
DNS Name Servers	Specifies a DNS server for the switching platform to use to resolve hostnames into addresses.	To add an IP address, click <b>Add</b> .  To edit an IP address, click <b>Edit</b> .  To delete an IP address, click <b>Delete</b> .
Domain Search	Specifies the domains to be searched.	To add a domain, click <b>Add</b> .  To edit a domain click <b>Edit</b> .  To delete a domain, click <b>Delete</b> .

**Related Documentation**

- [Configuring Date and Time for the EX Series Switch \(J-Web Procedure\) on page 35](#)

## Configuration Statements

- [arp \(System\) on page 39](#)
- [authentication-key on page 40](#)
- [auxiliary on page 41](#)
- [boot-server \(NTP\) on page 42](#)
- [broadcast on page 43](#)
- [broadcast-client on page 44](#)
- [console \(System Ports\) on page 45](#)
- [default-address-selection on page 46](#)
- [domain-name on page 47](#)
- [gre-path-mtu-discovery on page 47](#)
- [host-name on page 48](#)
- [icmpv4-rate-limit on page 48](#)
- [icmpv6-rate-limit on page 49](#)
- [inet6-backup-router on page 50](#)
- [internet-options on page 51](#)
- [ipip-path-mtu-discovery on page 52](#)

- [ipv6-duplicate-addr-detection-transmits](#) on page 52
- [ipv6-path-mtu-discovery](#) on page 53
- [ipv6-path-mtu-discovery-timeout](#) on page 53
- [ipv6-reject-zero-hop-limit](#) on page 54
- [lcd-menu](#) on page 55
- [location \(System\)](#) on page 57
- [menu-item](#) on page 59
- [multicast-client](#) on page 62
- [no-multicast-echo](#) on page 63
- [no-ping-record-route](#) on page 64
- [no-ping-time-stamp](#) on page 64
- [no-redirects \(IPv4 Traffic\)](#) on page 65
- [no-tcp-rfc1323-paws](#) on page 66
- [no-tcp-rfc1323](#) on page 66
- [ntp](#) on page 67
- [path-mtu-discovery](#) on page 68
- [peer \(NTP\)](#) on page 69
- [port-type](#) on page 70
- [ports](#) on page 71
- [power](#) on page 72
- [processes](#) on page 73
- [server \(NTP\)](#) on page 74
- [tcp-drop-synfin-set](#) on page 75
- [traceoptions \(SBC Configuration Process\)](#) on page 76
- [trusted-key](#) on page 78

## arp (System)

<b>Syntax</b>	<pre>arp {     aging-timer <i>minutes</i>;     gratuitous-arp-delay <i>seconds</i>;     gratuitous-arp-on-ifup;     interfaces {         <i>interface-name</i> {             aging-timer <i>minutes</i>;         }     }     passive-learning;     purging; }</pre> <p>For EX-Series switches:</p> <pre>arp {     aging-timer <i>minutes</i>; }</pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Specify ARP options. You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.</p> <p>For EX-Series switches, set only the time interval between ARP updates.</p>
<b>Options</b>	<p><b>aging-timer</b>—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, on routers only, metro Ethernet environments), increasing the time between updates can improve system performance.</p> <p><b>passive-learning</b> (QFX-Series only)—Configure switches to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests.</p> <p><b>Default:</b> 20 minutes</p> <p><b>Range:</b> 1 to 240 minutes</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses</i></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>

- For more information about ARP updates, see the [Junos OS System Basics Configuration Guide](#).

## authentication-key

---

<b>Syntax</b>	<code>authentication-key <i>key-number</i> type <i>type</i> value <i>password</i>;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">ntp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	<p>Configure Network Time Protocol (NTP) authentication keys so that the router or switch can send authenticated packets. If you configure the router or switch to operate in authenticated mode, you must configure a key.</p> <p>Both the keys and the authentication scheme (MD5) must be identical between a set of peers sharing the same key number.</p>
<b>Options</b>	<p><b><i>key-number</i></b>—Positive integer that identifies the key.</p> <p><b><i>type type</i></b>—Authentication type. It can only be <b>md5</b>.</p> <p><b><i>value password</i></b>—The key itself, which can be from 1 through 8 ASCII characters. If the key contains spaces, enclose it in quotation marks.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring NTP Authentication Keys</a></li><li>• <a href="#">broadcast on page 43</a></li><li>• <a href="#">peer on page 69</a></li><li>• <a href="#">server on page 74</a></li><li>• <a href="#">trusted-key on page 78</a></li></ul>



## auxiliary

---

<b>Syntax</b>	<pre> auxiliary {   disable;   insecure;   type <i>terminal-type</i>;   port-type (mini-usb   rj45); } </pre>
<b>Hierarchy Level</b>	[edit system <b>ports</b> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Configure the characteristics of the auxiliary port.</p> <p>Remaining statement is explained separately.</p>
<b>Default</b>	<b>disable</b> is the default option.
<b>Options</b>	<p><b>disable</b>—Disable the port.</p> <p><b>insecure</b>—Disable super user access or root logins to establish terminal connection.</p> <p><b>type <i>terminal-type</i></b>—Type of terminal that is connected to the port.</p> <p><b>Range:</b> <b>ansi</b>, <b>vt100</b>, <b>small-xterm</b>, <b>xterm</b></p> <p><b>Default:</b> The terminal type is unknown, and the user is prompted for the terminal type. The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring Junos OS to Set Console and Auxiliary Port Properties</i></li> </ul>

## boot-server (NTP)

---

<b>Syntax</b>	<code>boot-server (address   hostname);</code>
<b>Hierarchy Level</b>	[edit system <a href="#">ntp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	<p>Configure the server that NTP queries when the router or switch boots to determine the local date and time.</p> <p>When you boot the router or switch, it issues an <b>ntpdate</b> request, which polls a network server to determine the local date and time. You need to configure a server that the router or switch uses to determine the time when the router or switch boots. You can either configure an IP address or a hostname for the boot server. If you configure a hostname instead of an IP address, the <b>ntpdate</b> request resolves the hostname to an IP address when the router or switch boots up.</p> <p>If you configure an NTP boot server, then when the router or switch boots, it immediately synchronizes with the boot server even if the NTP process is explicitly disabled or if the time difference between the client and the boot server exceeds the threshold value of 1000 seconds.</p>
<b>Options</b>	<ul style="list-style-type: none"><li>• <b>address</b>—The IP address of an NTP boot server.</li><li>• <b>hostname</b>—The hostname of an NTP boot server.</li></ul>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Synchronizing and Coordinating Time Distribution Using NTP</i></li></ul>

## broadcast


<b>Syntax</b>	<code>broadcast address &lt;key key-number&gt; &lt;routing-instance-name routing-instance-name&gt; &lt;ttl value&gt; &lt;version value&gt;;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">ntp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. <b>routing-instance-name</b> option added in Junos OS Release 14.1
<b>Description</b>	Configure the local router or switch to operate in broadcast mode with the remote system at the specified <b>address</b> . In this mode, the local router or switch sends periodic broadcast messages to a client population at the specified broadcast or multicast <b>address</b> . Normally, you include this statement only when the local router or switch is operating as a transmitter.
<b>Options</b>	<p><b>address</b>—The broadcast address on one of the local networks or a multicast address assigned to NTP. You must specify an address, not a hostname. If the multicast address is used, it must be <b>224.0.1.1</b>.</p> <p><b>key key-number</b>—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number. <b>Range:</b> Any unsigned 32-bit integer</p> <p><b>routing-instance-name routing-instance-name</b>—(Optional) The routing instance name in which the interface has address in the broadcast subnet. <b>Default:</b> The default routing instance is used to broadcast packets.</p> <p><b>ttl value</b>—(Optional) Time-to-live (TTL) value to use. <b>Range:</b> 1 through 255 <b>Default:</b> 1</p> <p><b>version value</b>—(Optional) Specify the version number to be used in outgoing NTP packets. <b>Range:</b> 1 through 4 <b>Default:</b> 4</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring the NTP Time Server and Time Services</i></li> </ul>

## **broadcast-client**

---

<b>Syntax</b>	<code>broadcast-client;</code>
<b>Hierarchy Level</b>	<code>[edit system <a href="#">ntp</a>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the local router or switch to listen for broadcast messages on the local network to discover other servers on the same subnet.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Router or Switch to Listen for Broadcast Messages Using NTP</i></li></ul>

## console (System Ports)

<b>Syntax</b>	<pre>console {   disable;   insecure;   log-out-on-disconnect;   type <i>terminal-type</i>; }</pre>
<b>Hierarchy Level</b>	[edit system <a href="#">ports</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>disable</b> option added in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Configure the characteristics of the console port.
<b>Default</b>	The console port is enabled and its speed is 9600 baud.
<b>Options</b>	<p><b>disable</b>—Disable console login connections.</p> <p><b>insecure</b>—Disable root login connections to the console and auxiliary ports. Configuring the console port as insecure also prevents superusers and anyone with a user identifier (UID) of 0 from establishing terminal connections in multiuser mode. This option can be used to prevent a user from attempting password recovery by booting into single-user mode, if the user does not know the root password.</p> <p><b>log-out-on-disconnect</b>—Log out the session when the data carrier on the console port is lost.</p>
	<p> <b>NOTE:</b> The <b>log-out-on-disconnect</b> option is not operational on MX80 routers. On MX80 routers you must manually log out from the console with the <b>request system logout u0</b> command.</p>
	<p><b>type <i>terminal-type</i></b>—Type of terminal that is connected to the port.</p> <p><b>Range:</b> <b>ansi, vt100, small-xterm, xterm</b></p> <p><b>Default:</b> The terminal type is unknown, and the user is prompted for the terminal type.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring Junos OS to Set Console and Auxiliary Port Properties</i></li> </ul>

## default-address-selection

---

<b>Syntax</b>	default-address-selection;
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Use the loopback interface, <b>lo0</b> , as the source address for all locally generated IP packets when the packet is sent through a routed interface, and also when the packet is sent through a local interface such as <b>fxp0</b> . The <b>lo0</b> interface is the interface to the router's or switch's Routing Engine.
<b>Default</b>	<p>The default address is used as the source address for all locally generated IP packets on outgoing interfaces that are unnumbered. If an outgoing interface is numbered, the default address is chosen using the following sequence:</p> <ul style="list-style-type: none"><li>• The primary address on the loopback interface <b>lo0</b> that is <i>not</i> <b>127.0.0.1</b> is used.</li><li>• The primary address for the primary interface or the preferred address (if configured) for the primary interface is used.</li></ul> <p>By default, the primary address on an interface is selected as the numerically lowest local address configured on the interface.</p> <p>An interface's <i>primary address</i> is used by default as the local address for broadcast and multicast packets sourced locally and sent out through the interface. An interface's <i>preferred address</i> is the default local address used for packets sourced by the local router or switch to destinations on the subnet. By default, the numerically lowest local address configured for the interface is chosen as the preferred address on the subnet.</p> <p>To configure a different primary address or preferred address, include the <b>primary</b> or <b>preferred</b> statement at the <b>[edit interfaces interface-name unit logical-unit-number family family address address]</b> or <b>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family address address]</b> hierarchy levels.</p> <p>For more information about default, primary, and preferred addresses for an interface, see "Configuring Default, Primary, and Preferred Addresses and Interfaces" in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets</i></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>

## domain-name

---

<b>Syntax</b>	<code>domain-name <i>domain-name</i>;</code>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the name of the domain in which the router or switch is located. This is the default domain name that is appended to hostnames that are not fully qualified.
<b>Options</b>	<i>domain-name</i> —Name of the domain.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Reaching a Domain Name System Server</i></li> </ul>

## gre-path-mtu-discovery

---

<b>Syntax</b>	<code>(gre-path-mtu-discovery   no-gre-path-mtu-discovery);</code>
<b>Hierarchy Level</b>	[edit system <a href="#">internet-options</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure path MTU discovery for outgoing GRE tunnel connections: <ul style="list-style-type: none"> <li>• <b>gre-path-mtu-discovery</b>—Path MTU discovery is enabled.</li> <li>• <b>no-gre-path-mtu-discovery</b>—Path MTU discovery is disabled.</li> </ul>
<b>Default</b>	Path MTU discovery is enabled.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Junos OS for Path MTU Discovery on Outgoing GRE Tunnel Connections</i></li> </ul>

## host-name

---

<b>Syntax</b>	<code>host-name <i>hostname</i>;</code>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Set the hostname of the router or switch.
<b>Options</b>	<i>hostname</i> —Name of the router or switch.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Understanding Hostnames</i></li><li>• <i>Configuring the Hostname of the Router or Switch</i></li></ul>

## icmpv4-rate-limit

---

<b>Syntax</b>	<code>icmpv4-rate-limit {     bucket-size <i>seconds</i>;     packet-rate <i>pps</i>; }</code>
<b>Hierarchy Level</b>	[edit system <a href="#">internet-options</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure rate-limiting parameters for ICMPv4 messages sent.
<b>Options</b>	<b>bucket-size <i>seconds</i></b> —Number of seconds in the rate-limiting bucket. <b>Range:</b> 0 through 4294967295 seconds <b>Default:</b> 5  <b>packet-rate <i>pps</i></b> —Rate-limiting packets earned per second. <b>Range:</b> 0 through 4294967295 pps <b>Default:</b> 1000
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages</i></li></ul>



## icmpv6-rate-limit

---

<b>Syntax</b>	icmpv6-rate-limit { bucket-size <i>seconds</i> ; packet-rate <i>packet-rate</i> ; }
<b>Hierarchy Level</b>	[edit system <a href="#">internet-options</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure rate-limiting parameters for ICMPv6 messages sent.
<b>Options</b>	<p><b>bucket-size <i>seconds</i></b>—Number of seconds in the rate-limiting bucket.  <b>Range:</b> 0 through 4294967295 seconds  <b>Default:</b> 5</p> <p><b>packet-rate <i>pps</i></b>—Rate-limiting packets earned per second.  <b>Range:</b> 0 through 4294967295 pps  <b>Default:</b> 1000</p>
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages</i></li> </ul>

## inet6-backup-router

---

<b>Syntax</b>	<code>inet6-backup-router <i>address</i> &lt;destination <i>destination-address</i>&gt;;</code>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Set a default router (running IP version 6 [IPv6]) to use while the local router or switch (running IPv6) is booting and if the routing protocol processes fail to start. The Junos OS removes the route to this router or switch as soon as the software starts.
<b>Options</b>	<p><b><i>address</i></b>—Address of the default router.</p> <p><b><i>destination destination-address</i></b>—(Optional) Destination address that is reachable through the backup router. You can include this option to achieve network reachability while loading, configuring, and recovering the router or switch, but without the risk of installing a default route in the forwarding table.</p> <p><b>Default:</b> All hosts (default route) are reachable through the backup router.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring a Backup Router</i></li></ul>

## internet-options

<b>Syntax</b>	<pre> internet-options {   (gre-path-mtu-discovery   no-gre-path-mtu-discovery);   icmpv4-rate-limit bucket-size <i>bucket-size</i> packet-rate <i>packet-rate</i>;   icmpv6-rate-limit bucket-size <i>bucket-size</i> packet-rate <i>packet-rate</i>;   (ipip-path-mtu-discovery   no-ipip-path-mtu-discovery);   ipv6-duplicate-addr-detection-transmits;   (ipv6-reject-zero-hop-limit   no-ipv6-reject-zero-hop-limit);   (ipv6-path-mtu-discovery   no-ipv6-path-mtu-discovery);   ipv6-path-mtu-discovery-timeout;   no-tcp-rfc1323;   no-tcp-rfc1323-paws;   (path-mtu-discovery   no-path-mtu-discovery);   source-port upper-limit &lt;<i>upper-limit</i>&gt;;   (source-quench   no-source-quench);   tcp-drop-synfin-set;   tcp-mss <i>mss-value</i>; } </pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Configure system IP options to protect against certain types of DoS attacks.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages</i></li> <li>• <i>Configuring the Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages</i></li> <li>• <i>Configuring the Junos OS for IP-IP Path MTU Discovery on IP-IP Tunnel Connections</i></li> <li>• <i>Configuring the Junos OS for Path MTU Discovery on Outgoing GRE Tunnel Connections</i></li> <li>• <i>Configuring the Junos OS for Path MTU Discovery on Outgoing TCP Connections</i></li> <li>• <i>Configuring the Junos OS for IPv6 Duplicate Address Detection Attempts</i></li> <li>• <i>Configuring the Junos OS for Acceptance of IPv6 Packets with a Zero Hop Limit</i></li> <li>• <i>Configuring the Junos OS to Ignore ICMP Source Quench Messages</i></li> <li>• <i>Configuring the Junos OS to Enable the Router or Switch to Drop Packets with the SYN and FIN Bits Set</i></li> <li>• <i>Configuring the Junos OS to Disable TCP RFC 1323 Extensions</i></li> <li>• <i>Configuring the Junos OS to Disable the TCP RFC 1323 PAWS Extension</i></li> <li>• <i>Configuring the Junos OS to Extend the Default Port Address Range</i></li> </ul>

- *Configuring TCP MSS for Session Negotiation*

## ipip-path-mtu-discovery

---

<b>Syntax</b>	(ipip-path-mtu-discovery   no-ipip-path-mtu-discovery);
<b>Hierarchy Level</b>	[edit system internet-options]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure path MTU discovery for outgoing IP-IP tunnel connections: <ul style="list-style-type: none"><li>• <b>ipip-path-mtu-discovery</b>—Path MTU discovery is enabled.</li><li>• <b>no-ipip-path-mtu-discovery</b>—Path MTU discovery is disabled.</li></ul>
<b>Default</b>	Path MTU discovery is enabled.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS for IP-IP Path MTU Discovery on IP-IP Tunnel Connections</i></li><li>• <a href="#">internet-options on page 51</a></li></ul>

## ipv6-duplicate-addr-detection-transmits

---

<b>Syntax</b>	ipv6-duplicate-addr-detection-transmits;
<b>Hierarchy Level</b>	[edit system <a href="#">internet-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches.
<b>Description</b>	Control the number of attempts for IPv6 duplicate address detection.  The range of values supported for ipv6-duplicate-addr-detection-transmits is from 0 to 20.
<b>Default</b>	The default value is 3.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Junos OS for IPv6 Duplicate Address Detection Attempts</i></li></ul>

## ipv6-path-mtu-discovery

<b>Syntax</b>	(ipv6-path-mtu-discovery   no-ipv6-path-mtu-discovery);
<b>Hierarchy Level</b>	[edit system <a href="#">internet-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 9.2 for EX Series switches.
<b>Description</b>	Configure path MTU discovery for IPv6 packets: <ul style="list-style-type: none"> <li>• <b>ipv6-path-mtu-discovery</b>—IPv6 path MTU discovery is enabled.</li> <li>• <b>no-ipv6-path-mtu-discovery</b>—IPv6 path MTU discovery is disabled.</li> </ul>
<b>Default</b>	IPv6 path MTU discovery is enabled.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Junos OS for IPv6 Path MTU Discovery</i></li> </ul>

## ipv6-path-mtu-discovery-timeout

<b>Syntax</b>	ipv6-path-mtu-discovery-timeout <i>minutes</i> ;
<b>Hierarchy Level</b>	[edit system <a href="#">internet-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 9.2 for EX Series switches.
<b>Description</b>	Set the IPv6 path MTU discovery timeout interval.
<b>Options</b>	<i>minutes</i> —IPv6 path MTU discovery timeout. <b>Default:</b> 10 minutes
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Junos OS for IPv6 Path MTU Discovery</i></li> </ul>

## ipv6-reject-zero-hop-limit

---

<b>Syntax</b>	(ipv6-reject-zero-hop-limit   no-ipv6-reject-zero-hop-limit);
<b>Hierarchy Level</b>	[edit system <a href="#">internet-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches.
<b>Description</b>	Enable and disable rejecting incoming IPv6 packets with a zero hop limit value in their header.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Junos OS for Acceptance of IPv6 Packets with a Zero Hop Limit</i></li></ul>

## lcd-menu

**Syntax** EX3200, EX3300, EX4200, or EX4500 switch:

```
lcd-menu fpc slot-number {
  menu-item (menu-name | menu-option) <disable>;
}
```

EX6200 or EX8200 switch or XRE200 External Routing Engine:

```
lcd-menu {
  menu-item (menu-name | menu-option) <disable>;
}
```

**Hierarchy Level** [edit chassis]

**Release Information** Statement introduced in Junos OS Release 10.2 for EX Series switches.

**Description** Disable or enable the Maintenance menu or the Status menu in the LCD panel.

**Options** none—(EX6200 and EX8200 switches and XRE200 External Routing Engines only) Disable or enable the specified menu or menu options.

**fpc slot-number**—(EX3200, EX3300, EX4200, and EX4500 switches only) Disable or enable the specified menu or menu options, where **slot-number** is:

- 0—On standalone switches.
- 0–9—On a device in a Virtual Chassis. The value is the member ID of the device.



**NOTE:** This option is not available on an EX8200 Virtual Chassis. The LCD panel on an XRE200 External Routing Engine provides information for the XRE200 External Routing Engine only.

**disable**—(Optional) Disable the specified menu.

The remaining statement is explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-level—To add this statement to the configuration.

**Related Documentation**

- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\) on page 33](#)
- [LCD Panel in EX3200 Switches](#)
- [LCD Panel in EX3300 Switches](#)
- [LCD Panel in EX4200 Switches](#)
- [LCD Panel in EX4500 Switches](#)
- [LCD Panel in an EX6200 Switch](#)

- *LCD Panel in an EX8200 Switch*
- *LCD Panel in an XRE200 External Routing Engine*



## location (System)

<b>Syntax</b>	<pre>location {   altitude <i>feet</i>;   building <i>name</i>;   country-code <i>code</i>;   floor <i>number</i>;   hcoord <i>horizontal-coordinate</i>;   lata <i>transport-area</i>;   latitude <i>degrees</i>;   longitude <i>degrees</i>;   npa-nxx <i>number</i>;   postal-code <i>postal-code</i>;   rack <i>number</i>;   vcoord <i>vertical-coordinate</i>; }</pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Configure the system location in various formats.
<b>Options</b>	<p><b>altitude <i>feet</i></b>—Number of feet above sea level.</p> <p><b>building <i>name</i></b>—Name of building. The name of the building can be 1 to 28 characters in length. If the string contains spaces, enclose it in quotation marks (" ").</p> <p><b>country-code <i>code</i></b>—Two-letter country code.</p> <p><b>floor <i>number</i></b>—Floor in the building.</p> <p><b>hcoord <i>horizontal-coordinate</i></b>—Bellcore Horizontal Coordinate.</p> <p><b>lata <i>transport-area</i></b>—Local Access Transport Area.</p> <p><b>latitude <i>degrees</i></b>—Latitude in degree format.</p> <p><b>longitude <i>degrees</i></b>—Longitude in degree format.</p> <p><b>npa-nxx <i>number</i></b>—First six digits of the phone number (area code and exchange).</p> <p><b>postal-code <i>postal-code</i></b>—Postal code.</p> <p><b>rack <i>number</i></b>—Rack number.</p> <p><b>vcoord <i>vertical-coordinate</i></b>—Bellcore Vertical Coordinate.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

**Related Documentation** • *Specifying the Physical Location of the Router or Switch*

## menu-item

---

<b>Syntax</b>	<code>menu-item (<i>menu-name</i>   <i>menu-option</i>) &lt;disable&gt;;</code>
<b>Hierarchy Level</b>	<code>[edit chassis <b>lcd-menu</b>],</code> <code>[edit chassis <b>lcd-menu</b> <i>fpc slot-number</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.2 for EX Series switches.
<b>Description</b>	<p>Disable or enable the Maintenance menu, the Status menu, or an individual option in one of those menus in the LCD panel.</p> <p>On EX3200, EX3300, EX4200, and EX4500 switches, you use <b>menu-item</b> at the <code>[edit chassis <b>lcd-menu</b> <i>fpc slot-number</i>]</code> hierarchy level.</p> <p>On EX6200 and EX8200 switches, and on XRE200 External Routing Engines, you use <b>menu-item</b> at the <code>[edit chassis <b>lcd-menu</b>]</code> hierarchy level.</p>
<b>Options</b>	<p><b><i>menu-name</i></b>—Name of the LCD menu:</p> <ul style="list-style-type: none"> <li>• <b>maintenance-menu</b></li> <li>• <b>status-menu</b></li> </ul> <p><b><i>menu-option</i></b>—Specific option on one of the LCD menus. You must include the quotation marks when you type the option. <a href="#">Table 9 on page 60</a> describes the different menu options of the LCD menus supported on the switches.</p>

Table 9: Menu Options of the LCD Menus Supported on the Switches

Menu	Menu Options	Option Descriptions	Platforms Supported
maintenance-menu	"maintenance-menu halt-menu"	System halt option	All switches except EX2200
	"maintenance-menu system-reboot"	System reboot option	All switches except EX2200
	"maintenance-menu rescue-config"	Load rescue option	All switches except EX2200
	"maintenance-menu vc-uplink-config"	Request VC port option for a device in a Virtual Chassis configuration	EX3300, EX4200, and EX4500 switches and XRE200 External Routing Engines only
	"maintenance-menu factory-default"	Factory default option	All switches except EX2200
status-menu	"status-menu vcp-status"	Virtual Chassis port (VCP) status for a device in a Virtual Chassis configuration	EX3300, EX4200, and EX4500 switches and XRE200 External Routing Engines only
	"status-menu sf-status1-menu"	Status of the switch fabric on the Switch Fabric and Routing Engine (SRE) module in slot SRE0 on EX8208 switches	EX8208 and EX8216 switches only
		Status of the switch fabric on the Switch Fabric (SF) modules in slots SF0 and SF1 on EX8216 switches	
	"status-menu sf-status2-menu"	Status of the switch fabric on the SRE module in slot SRE1 on EX8208 switches	EX8208 and EX8216 switches only
		Status of the switch fabric on the SF modules in slots SF2–SF5 on EX8216 switches	
	"status-menu sf-status3-menu"	Status of the switch fabric on the SF modules in slots SF6 and SF7 on EX8216 switches	EX8216 switches only
	"status-menu power-status"	Status of the power supply or power supplies	EX3200, EX3300, EX4200, and EX4500 switches

Table 9: Menu Options of the LCD Menus Supported on the Switches (*continued*)

Menu	Menu Options	Option Descriptions	Platforms Supported
			and XRE200 External Routing Engines only
	"status-menu psu-status1-menu"	Status of the power supplies in slots P0 and P1	EX8208 and EX8216 switches only
	"status-menu psu-status2-menu"	Status of the power supplies in slots P2–P5	EX8208 and EX8216 switches only
	"status-menu environ-menu"	Status of the fan; current chassis temperature	All switches (except EX2200) and XRE200 External Routing Engine
	"status-menu show-version"	The version of Junos OS loaded on the switch	All switches except EX2200

**disable**—(Optional) Disable the specified menu.

**Required Privilege Level** view-level—To view this statement in the configuration.  
control-level—To add this statement to the configuration.

**Related Documentation**

- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\) on page 33](#)
- *LCD Panel in EX3200 Switches*
- *LCD Panel in EX3300 Switches*
- *LCD Panel in EX4200 Switches*
- *LCD Panel in EX4500 Switches*
- *LCD Panel in EX4550 Switches*
- *LCD Panel in an EX6200 Switch*
- *LCD Panel in an EX8200 Switch*
- *LCD Panel in an XRE200 External Routing Engine*

## multicast-client

---

<b>Syntax</b>	<code>multicast-client &lt;address&gt;;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">ntp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	For NTP, configure the local router or switch to listen for multicast messages on the local network to discover other servers on the same subnet.
<b>Options</b>	<b>address</b> —(Optional) One or more IP addresses. If you specify addresses, the router or switch joins those multicast groups. <b>Default:</b> 224.0.1.1.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Router or Switch to Listen for Multicast Messages Using NTP</i></li></ul>

## no-multicast-echo

```
Syntax  no-multicast-echo {
        arp {
            aging-timer minutes;
            gratuitous-arp-delay seconds;
            gratuitous-arp-on-ifup;
            interfaces {
                interface-name {
                    aging-timer minutes;
                }
            }
            passive-learning;
            purging;
        }
        host-name hostname;
        location {
            altitude feet;
            building name;
            country-code code;
            floor number;
            hcoord horizontal-coordinate;
            lata service-area;
            latitude degrees;
            longitude degrees;
            npa-nxx number;
            postal-code postal-code;
            rack number;
            vcoord vertical-coordinate;
        }
    }
    license {
        autoupdate URL;
    }
    renew before-expiration (number | interval number)
}
}
```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced in Junos OS Release 8.1.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Disable the Routing Engine from responding to ICMP echo requests sent to multicast group addresses.

**Default** The Routing Engine responds to ICMP echo requests sent to multicast group addresses.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping Packets*

---

## no-ping-record-route

---

<b>Syntax</b>	no-ping-record-route;
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the Junos OS to disable the reporting of the IP address in ping responses.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <i>Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses</i>

---

## no-ping-time-stamp

---

<b>Syntax</b>	no-ping-time-stamp;
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the Junos OS to disable the recording of timestamps in ping responses.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <i>Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses</i>



## no-redirects (IPv4 Traffic)

<b>Syntax</b>	no-redirects;
<b>Hierarchy Level</b>	[edit system], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for EX Series switches.
<b>Description</b>	<p>Stop protocol redirect messages for IPv4 traffic from being sent on the entire switch or on an interface on the router or switch.</p> <p>To disable the sending of protocol redirect messages for the entire router or switch, include the <b>no-redirects</b> statement at the [edit system] hierarchy level.</p> <p>To disable the sending of protocol redirect messages on a specific interface, include the <b>no-redirects</b> statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>] hierarchy level.</p>
<b>Default</b>	The router or switch sends redirect messages.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Junos OS to Disable Protocol Redirect Messages on the Router or Switch</i></li> <li>• <i>Understanding the Protocol Redirect Mechanism on EX Series Switches</i></li> <li>• <i>Configuring Junos OS to Disable Sending Protocol Redirect Messages on EX Series Switches (CLI Procedure)</i></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>

## no-tcp-rfc1323-paws

---

<b>Syntax</b>	no-tcp-rfc1323-paws;
<b>Hierarchy Level</b>	[edit system internet-options]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the Junos OS to disable the RFC 1323 Protection Against Wrapped Sequence (PAWS) number extension.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS to Disable the TCP RFC 1323 PAWS Extension</i></li></ul>

## no-tcp-rfc1323

---

<b>Syntax</b>	no-tcp-rfc1323;
<b>Hierarchy Level</b>	[edit system internet-options]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the Junos OS to disable RFC 1323 TCP extensions.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS to Disable TCP RFC 1323 Extensions</i></li></ul>

## ntp

<b>Syntax</b>	<pre> ntp {   authentication-key number type type value password;   boot-server address;   broadcast &lt;address&gt; &lt;key key-number&gt; &lt;routing-instance-name routing-instance-name&gt;     &lt;version value&gt; &lt;ttl value&gt;;   broadcast-client;   multicast-client &lt;address&gt;;   peer address &lt;key key-number&gt; &lt;version value&gt; &lt;prefer&gt;;   server address &lt;key key-number&gt; &lt;version value&gt; &lt;prefer&gt;;   source-address source-address &lt;routing-instance routing-instance-name&gt;;   trusted-key [ key-numbers ]; } </pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Configure NTP on the router or switch.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Synchronizing and Coordinating Time Distribution Using NTP</i></li> </ul>

## path-mtu-discovery

---

<b>Syntax</b>	(path-mtu-discovery   no-path-mtu-discovery);
<b>Hierarchy Level</b>	[edit system <a href="#">internet-options</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure path MTU discovery for outgoing Transmission Control Protocol (TCP) connections: <ul style="list-style-type: none"><li>• <b>path-mtu-discovery</b>—Path MTU discovery is enabled.</li><li>• <b>no-path-mtu-discovery</b>—Path MTU discovery is disabled.</li></ul>
<b>Default</b>	Path MTU discovery is enabled.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS for Path MTU Discovery on Outgoing TCP Connections</i></li></ul>

## peer (NTP)

<b>Syntax</b>	<code>peer address &lt;key key-number&gt; &lt;version value&gt; &lt;prefer&gt;;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">ntp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	For NTP, configure the local router or switch to operate in symmetric active mode with the remote system at the specified address. In this mode, the local router or switch and the remote system can synchronize with each other. This configuration is useful in a network in which either the local router or switch or the remote system might be a better source of time.
<b>Options</b>	<p><b>address</b>—Address of the remote system. You must specify an address, not a hostname.</p> <p><b>key key-number</b>—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number.</p> <p><b>Range:</b> Any unsigned 32-bit integer</p> <p><b>prefer</b>—(Optional) Mark the remote system as the preferred host, which means that if all other factors are equal, this remote system is chosen for synchronization among a set of correctly operating systems.</p> <p><b>version value</b>—(Optional) Specify the NTP version number to be used in outgoing NTP packets.</p> <p><b>Range:</b> 1 through 4</p> <p><b>Default:</b> 4</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the NTP Time Server and Time Services</i></li> </ul>

## port-type

---


<b>Syntax</b>	port-type (mini-usb   rj45);
<b>Hierarchy Level</b>	[edit system ports <a href="#">auxiliary</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for EX Series switches.
<b>Description</b>	(EX2200-C and EX4550 switch only) Set the RJ-45 console port or the Mini-USB console port as the active console port.
<b>Default</b>	The RJ-45 console port is the active port.
<b>Options</b>	<b>mini-usb</b> —Set the Mini USB type-B console port as the active console port. <b>rj45</b> —Set the RJ-45 console port as the active console port.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Console Port Type (CLI Procedure)</i></li></ul>

## ports

<b>Syntax</b>	<pre> ports {   auxiliary {     disable;     insecure;     type <i>terminal-type</i>;     port-type (mini-usb   rj45);   }   console {     disable;     insecure;     log-out-on-disconnect;     type <i>terminal-type</i>;   } } </pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Configure the properties of the console and auxiliary ports. The ports are located on the router's craft interface.</p> <p>See the switch's hardware documentation for port locations.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring Junos OS to Set Console and Auxiliary Port Properties</i></li> </ul>

## power

---

<b>Syntax</b>	power (off   on);
<b>Hierarchy Level</b>	[edit <a href="#">chassis fpc slot</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4 for EX Series switches.
<b>Description</b>	On an EX6200 or EX8200 switch, turn a specified Flexible PIC Concentrator (FPC) on or off. See <a href="#">“EX Series Switches Hardware and CLI Terminology Mapping” on page 315</a> .
<div> <b>NOTE:</b> On an EX6200 switch, the power statement has no effect when you configure it for an uplink port FPC on the Switch Fabric and Routing Engine (SRE) module. If you configure the statement for those FPCs, the configuration will be committed, but a message that informs you that the configuration has no effect is logged in the system log. You cannot turn the power on and off for these FPCs.</div>	
<b>Options</b>	<ul style="list-style-type: none"><li>• on—Turn on the specified FPC.</li><li>• off—Turn off the specified FPC.</li></ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Removing a Line Card from an EX6200 Switch</i></li></ul>



## processes

<b>Syntax</b>	<pre>processes {   process-name (enable   disable) failover (alternate-media   other-routing-engine);   timeout seconds; }</pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure which Junos OS processes are running on the router or switch.



**CAUTION:** Never disable any of the software processes unless instructed to do so by a customer support engineer.

<b>Default</b>	All processes are enabled by default.
<b>Options</b>	<p><b>(enable   disable)</b>—(Optional) Enable or disable a specified process.</p> <p><b>failover (alternate-media   other-routing-engine)</b>—(Optional) For routers or switches with redundant Routing Engines only, switch to backup media if a process fails repeatedly. If a process fails four times within 30 seconds, the router or switch reboots from the alternate media or the other Routing Engine.</p> <p><b>process-name</b>—One of the valid process names. You can obtain a complete list of process names by using the CLI command completion feature. After specifying a process name, command completion also indicates any additional options for that process.</p> <p><b>timeout seconds</b>—(Optional) How often the system checks the watchdog timer, in seconds. If the watchdog timer has not been checked in the specified number of seconds, the system reloads. If you set the time value too low, it is possible for the system to reboot immediately after it loads.</p> <p><b>Values:</b> 15, 60, or 180</p> <p><b>Default:</b> 180 seconds (rounded up to 291 seconds by the Junos kernel)</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Disabling Junos OS Processes</li> </ul>

## server (NTP)

---

<b>Syntax</b>	<code>server address &lt;key key-number&gt; &lt;version value&gt; &lt;prefer&gt;;</code>
<b>Hierarchy Level</b>	[edit system ntp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	<p>For NTP, configure the local router or switch to operate in client mode with the remote system at the specified <b>address</b>. In this mode, the local router or switch can be synchronized with the remote system, but the remote system can never be synchronized with the local router or switch.</p> <p>If the NTP client time drifts so that the difference in time from the NTP server exceeds 128 milliseconds, the client is automatically stepped back into synchronization. If the offset between the NTP client and server exceeds the 1000-second threshold, the client still synchronizes with the server, but it also generates a system log message noting that the threshold was exceeded.</p>
<b>Options</b>	<p><b>address</b>—Address of the remote system. You must specify an address, not a hostname.</p> <p><b>key key-number</b>—(Optional) Use the specified key number to encrypt authentication fields in all packets sent to the specified address.</p> <p><b>Range:</b> Any unsigned 32-bit integer</p> <p><b>prefer</b>—(Optional) Mark the remote system as preferred host, which means that if all other things are equal, this remote system is chosen for synchronization among a set of correctly operating systems.</p> <p><b>version value</b>—(Optional) Specify the version number to be used in outgoing NTP packets.</p> <p><b>Range:</b> 1 through 4</p> <p><b>Default:</b> 4</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the NTP Time Server and Time Services</i></li></ul>

---

## tcp-drop-synfin-set

---

<b>Syntax</b>	tcp-drop-synfin-set;
<b>Hierarchy Level</b>	[edit system internet-options]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the router or switch to drop packets that have both the SYN and FIN bits set.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS to Enable the Router or Switch to Drop Packets with the SYN and FIN Bits Set</i></li><li>• <i>TCP Headers with SYN and FIN Flags Set</i></li></ul>

## traceoptions (SBC Configuration Process)

---

<b>Syntax</b>	<pre>traceoptions {   file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regex</i>&gt; &lt;size <i>size</i>&gt;     &lt;world-readable   no-world-readable&gt;;   flag <i>flag</i>; }</pre>
<b>Hierarchy Level</b>	[edit system processes sbc-configuration-process]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Configure trace options for the session border controller (SBC) process of the border signaling gateway (BSG).
<b>Options</b>	<p><b>file <i>filename</i></b>—Name of the file that receives the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory <b>/var/log</b>. You can include the following file options:</p> <ul style="list-style-type: none"><li>• <b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</li></ul> <p>If you specify a maximum number of files, you must also specify a maximum file size with the <b>size</b> option and a filename.</p> <p><b>Range:</b> 2 through 1000 <b>Default:</b> 3 files</p> <ul style="list-style-type: none"><li>• <b>match <i>regex</i></b>—(Optional) Refine the output to include lines that contain the regular expression.</li><li>• <b>no-world-readable</b>—(Optional) Disable unrestricted file access.</li><li>• <b>size <i>size</i></b>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <b>trace-file</b> reaches this size, it is renamed <b>trace-file.0</b>. When the trace-file again reaches its maximum size, <b>trace-file.0</b> is renamed <b>trace-file.1</b> and <b>trace-file</b> is renamed <b>trace-file.0</b>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the files option and filename.</li></ul> <p><b>Syntax:</b> <b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify GB. <b>Range:</b> 10 KB through 1 GB <b>Default:</b> 128 KB</p> <ul style="list-style-type: none"><li>• <b>world-readable</b>—(Optional) Enable unrestricted file access.</li></ul>

**flag flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all trace-level**—Trace all SBC process operations.
- **common trace-level**—Trace common events.
- **configuration trace-level**—Trace configuration events.
- **device-monitor trace-level**—Trace device monitor events.
- **ipc trace-level**—Trace IPC events.
- **memory—pool trace-level**—Trace memory pool events.
- **trace-level**—Trace level options are related to the severity of the event being traced. When you choose a trace level, messages at that level and higher levels are captured. Enter one of the following trace levels as the **trace-level**:
  - **debug**—Log all code flow of control.
  - **error**—Log failures with a short-term effect.
  - **info**—Log summary for normal operations, such as the policy decisions made for a call.
  - **trace**—Log program trace START and EXIT macros.
  - **warning**—Log failure recovery events or failure of an external entity.
- **ui trace-level**—Trace user interface operations.

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
---------------------------------	---

<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• See “Troubleshooting the IMMSG” in the <i>Junos Multiplay Solutions Guide</i></li> <li>• <i>System Management Configuration Statements</i></li> </ul>
------------------------------	--

## trusted-key

---

<b>Syntax</b>	<code>trusted-key [ <i>key-numbers</i> ];</code>
<b>Hierarchy Level</b>	[edit system <a href="#">ntp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	For NTP, configure the keys you are allowed to use when you configure the local router or switch to synchronize its time with other systems on the network.
<b>Options</b>	<i>key-numbers</i> —One or more key numbers. Each key can be any 32-bit unsigned integer except 0.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring NTP Authentication Keys</i></li><li>• <a href="#">authentication-key on page 40</a></li><li>• <a href="#">broadcast on page 43</a></li><li>• <a href="#">peer on page 69</a></li><li>• <a href="#">server on page 74</a></li></ul>

## CHAPTER 4

# Administration

- [Operational Commands on page 79](#)

### Operational Commands

---

- [clear chassis display message](#)
- [clear system reboot](#)
- [configure](#)
- [op](#)
- [request chassis pic](#)
- [request chassis routing-engine master](#)
- [request system halt](#)
- [request system logout](#)
- [request system power-off](#)
- [request system reboot](#)
- [request system reboot](#)
- [request system scripts convert](#)
- [request system scripts refresh-from commit](#)
- [request system scripts refresh-from event](#)
- [request system scripts refresh-from op](#)
- [request system storage cleanup](#)
- [restart](#)
- [set chassis display message](#)
- [set date](#)
- [show chassis fan](#)
- [show chassis firmware](#)
- [show chassis lcd](#)
- [show configuration](#)
- [show host](#)
- [show ntp associations](#)

- `show ntp status`
- `show system firmware`
- `show system reboot`
- `show system software`
- `show system storage`
- `show system switchover`
- `show system uptime`
- `show system users`
- `show system virtual-memory`
- `show task replication`
- `show version`
- `show version fpc`



## clear chassis display message

<b>List of Syntax</b>	<a href="#">Syntax on page 81</a> <a href="#">Syntax (TX Matrix Router) on page 81</a> <a href="#">Syntax (TX Matrix Plus Router) on page 81</a> <a href="#">Syntax (QFabric Systems) on page 81</a>
<b>Syntax</b>	clear chassis display message
<b>Syntax (TX Matrix Router)</b>	clear chassis display message <lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	clear chassis display message <lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (QFabric Systems)</b>	clear chassis display message <node-device <i>name</i>   interconnect-device <i>name</i> >
<b>Release Information</b>	<p>Command introduced in Junos OS Release 7.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option for the TX Matrix Plus routers introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>(M40e, M160, M320, T Series routers, EX Series, and QFabric systems only) Clear or stop a text message on the craft interface display, which is on the front of the router or switch or on the LCD panel display on the router or switch. The craft interface alternates the display of text messages with standard craft interface messages, switching between messages every 2 seconds. By default, on both the router and the switch, the text message is displayed for 5 minutes. The craft interface display has four 20-character lines. The LCD panel display has two 16-character lines, and text messages appear only on the second line.</p>
<b>Options</b>	<p><b>none</b>—Clear or stop a text message on the craft interface display.</p> <p><b>interconnect-device <i>name</i></b>—(QFabric systems only) (Optional) On a QFabric system, clear or stop a text message on the LCD panel display on the specified Interconnect device.</p> <p><b>lcc <i>number</i></b>—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> <li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li> <li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li> </ul>

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**node-device *name***—(QFabric systems only) (Optional) On a QFabric system, clear or stop a text message on the LCD panel display on the specified Node device in a Node group.

**scc**—(TX Matrix routers only) (Optional) Clear or stop a text message on the craft interface on the TX Matrix router (switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Clear or stop a text message on the craft interface on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

**Required Privilege Level** clear

**Related Documentation**

- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\) on page 33](#)
- [set chassis display message on page 146](#)
- [show chassis craft-interface](#)

**List of Sample Output** [clear chassis display message on page 82](#)

**Output Fields** See [show chassis craft-interface](#) for an explanation of output fields.

## Sample Output

### clear chassis display message

The following example displays and then clears the text message on the craft interface display:

```
user@host> show chassis craft-interface
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
Host OK LED:    On
Host fail LED:  Off
FPCs           0  1  2  3  4  5  6  7
-----
Green  ..  *..  *  *.
Red    .....
LCD screen:
      +-----+
      |NOC contact Dusty|
      |(888) 526-1234   |
      +-----+

user@host> clear chassis display message

user@host> show chassis craft-interface
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
```

```
Host OK LED:  On
Host fail LED: Off
FPCs      0  1  2  3  4  5  6  7
-----
Green  ..  *..  *  *.
Red    .....
LCD screen:
+-----+
|host    |
|Up: 0+17:05:47|
|        |
|Temperature OK|
+-----+
```

## clear system reboot

---

<b>List of Syntax</b>	<a href="#">Syntax on page 84</a> <a href="#">Syntax (EX Series Switches) on page 84</a> <a href="#">Syntax (TX Matrix Router) on page 84</a> <a href="#">Syntax (TX Matrix Plus Router) on page 84</a> <a href="#">Syntax (QFX Series) on page 84</a>
<b>Syntax</b>	clear system reboot <both-routing-engines>
<b>Syntax (EX Series Switches)</b>	clear system reboot <all-members> <both-routing-engines> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	clear system reboot <both-routing-engines> <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	clear system reboot <both-routing-engines> <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (QFX Series)</b>	clear system reboot <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Clear any pending system software reboots or halts. When issued on a TX Matrix router without any options, the default behavior clears all pending system software reboots or halts on all T640 routers connected to the TX Matrix router. When issued on a TX Matrix Plus router without any options, the default behavior clears all pending system software reboots or halts on all T1600 or T4000 routers connected to the TX Matrix Plus router.
<b>Options</b>	<b>none</b> —Clear all pending system software reboots or halts.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Clear all halt or reboot requests for all the Routing Engines in the chassis.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, clear all halt or reboot requests for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, clear all halt or reboot requests on the l connected T1600 or T4000 LCCs.

**all-members**—(EX4200 switches only) (Optional) Clear all halt or reboot requests on all members of the Virtual Chassis configuration.

**both-routing-engines**—(Systems with multiple Routing Engines) (Optional) Clear all halt or reboot requests on both Routing Engines. On a TX Matrix router, clear both Routing Engines on all chassis connected to the TX Matrix router. Likewise, on a TX Matrix Plus router, clear both Routing Engines on all chassis connected to the TX Matrix Plus router.

**infrastructure *name***—(QFabric systems) (Optional) Clear all halt or reboot requests on the fabric control Routing Engines or fabric manager Routing Engines.

**interconnect-device *name***—(QFabric systems) (Optional) Clear all halt or reboot requests on the Interconnect device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, clear all halt or reboot requests for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, clear all halt or reboot requests for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches only) (Optional) Clear all halt or reboot requests on the local Virtual Chassis member.

**member *member-id***—(EX4200 switches only) (Optional) Clear all halt or reboot requests on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

**node-group *name***—(QFabric systems) (Optional) Clear all halt or reboot requests on the Node group.

**scc**—(TX Matrix routers only) (Optional) Clear all halt or reboot requests for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Clear all halt or reboot requests for the TX Matrix Plus router. Replace *number* with 0.

**Required Privilege Level**      maintenance

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">request system reboot on page 113</a></li><li>• <i>request system reboot</i></li><li>• <i>Rebooting and Halting a Device</i></li><li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear system reboot on page 87</a> <a href="#">clear system reboot (TX Matrix Router) on page 87</a> <a href="#">clear system reboot (QFX Series) on page 87</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear system reboot

```
user@host> clear system reboot
reboot requested by root at Sat Dec 12 19:37:34 1998
[process id 17855]
Terminating...
```

### clear system reboot (TX Matrix Router)

```
user@host> clear system reboot
scc-re0:
-----
No shutdown/reboot scheduled.
lcc0-re0:
-----
No shutdown/reboot scheduled.
lcc2-re0:
-----
No shutdown/reboot scheduled.
```

### clear system reboot (QFX Series)

```
user@switch> clear system reboot node-group node1
No shutdown/reboot scheduled.
```

## configure

---

<b>Syntax</b>	<code>configure</code> <code>&lt;batch&gt;</code> <code>&lt;dynamic&gt;</code> <code>&lt;exclusive&gt;</code> <code>&lt;private&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Enter configuration mode. When this command is entered without any optional keywords, everyone can make configuration changes and commit all changes made to the configuration.
<b>Options</b>	<p><b>none</b>—Enter configuration mode.</p> <p><b>batch</b>—(Optional) Work in the batch commit mode where commit operations are executed in batches.</p> <p><b>dynamic</b>—(Optional) Configure routing policies and certain routing policy objects in a dynamic database that is not subject to the same verification required in the standard configuration database. As a result, the time it takes to commit changes to the dynamic database is much shorter than for the standard configuration database. You can then reference these policies and policy objects in routing policies you configure in the standard database.</p> <p><b>exclusive</b>—(Optional) Lock the candidate configuration for as long as you remain in configuration mode, allowing you to make changes without interference from other users. Other users can enter and exit configuration mode, but they cannot change the configuration.</p> <p><b>private</b>—(Optional) Allow multiple users to edit different parts of the configuration at the same time and to commit only their own changes, or to roll back without interfering with one another's changes. You cannot commit changes in configure private mode when another user is in configure exclusive mode.</p>
<b>Additional Information</b>	For more information about the different methods of entering configuration mode and the restrictions that apply, see the <i>Junos OS Administration Library for Routing Devices</i> .
<b>Required Privilege Level</b>	configure
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show configuration on page 187</a></li></ul>
<b>List of Sample Output</b>	<a href="#">configure on page 89</a>
<b>Output Fields</b>	When you enter this command, you are placed in configuration mode and the system prompt changes from <i>hostname&gt;</i> to <i>hostname#</i> .



## Sample Output

configure

```
user@host> configure
Entering configuration mode
[edit]
user@host#
```

## op

---

<b>Syntax</b>	<code>op filename</code> <code>&lt;detail&gt;</code> <code>&lt;argument-name argument-value&gt;</code> <code>&lt;key (md5   sha-256   sha1) key-value</code> <code>&lt;url url&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>key</b> option introduced in Junos OS Release 10.0. <b>url</b> option introduced in Junos OS Release 10.0.
<b>Description</b>	Execute an op script stored in one of the following locations: <ul style="list-style-type: none"><li>• On the router or switch in the <code>/var/db/scripts/op</code> directory</li><li>• At a remote URL</li></ul>
<b>Options</b>	<b>detail</b> —(Optional) Display detailed output.  <b>argument-name argument-value</b> —(Optional) Specify one or more arguments to the script. For each argument you include on the command line, you must specify a corresponding value for the argument.  <b>key (md5   sha-256   sha1) key-value</b> —(Optional) With the <code>&lt;url&gt;</code> option, specify a checksum hash to verify the integrity of the script. You can include the <code>&lt;key&gt;</code> option if the <b>checksum</b> statement is included at the <code>[edit system scripts op file filename]</code> hierarchy level.  <b>url url</b> —(Optional) Specify a URL where the script is located.
<b>Additional Information</b>	For more information about Junos op scripts, see the <i>Junos OS Automation Library</i> .
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Executing an Op Script</i> in the <i>Junos OS Automation Library</i></li><li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Automation Library</i></li><li>• <i>checksum</i></li><li>• <a href="#">file checksum md5 on page 650</a></li><li>• <a href="#">file checksum sha-256 on page 652</a></li><li>• <a href="#">file checksum sha1 on page 651</a></li></ul>
<b>List of Sample Output</b>	<a href="#">op on page 91</a> <a href="#">op url on page 91</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

op

```
user@host> op script1 interface ge-0/2/0.0 protocol inet
```

op url

```
user@host> op url https://www.juniper.net/fa/2009-04-01.01.slax key md5  
8de24d09e1d90b2581bb937d2a5ad590 interface ge-0/2/0.0 protocol inet
```

## request chassis pic

---

<b>List of Syntax</b>	<a href="#">Syntax on page 92</a> <a href="#">Syntax (ACX4000 Series Routers) on page 92</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 92</a>
<b>Syntax</b>	<code>request chassis pic (offline   online) fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
<b>Syntax (ACX4000 Series Routers)</b>	<code>request chassis pic (offline   online) fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	<code>request chassis pic (offline   online) fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> &lt;lcc <i>number</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.3 for ACX4000 Routers. Command introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Control the operation of the PIC.



**NOTE:** The `request chassis pic (offline | online) fpc-slot slot number pic-slot slot-number` command is not supported for built-in PICs on MX Series routers.

To view a list of built-in PICs on the router or switch chassis, use the `show chassis hardware` command.



**NOTE:** This command is not supported on MX960 and MX2020 routers with MPC5EQ.



**NOTE:** T1600 routers and TX Matrix Plus routers with 100-Gigabit Ethernet PICs require two adjacent PIC slots, 0 and 1, for each PIC. Therefore, only online and offline command options to PIC slot 0 are allowed. Use of the online and offline command options for PIC slot 1 with the described router and PIC combination is not allowed.



**NOTE:** In T Series routers, when the PIC state is set from offline to online or vice-versa before the processing is complete for the previous command, you are provided feedback on the status of your request. The following sample messages are displayed if you try to set a PIC offline or online:

```
user@switch> request chassis pic fpc-slot 1 pic-slot 0 online
fpc 1 pic 0 online initiated, use "show chassis fpc pic-status" to verify
```

```
user@switch> request chassis pic fpc-slot 1 pic-slot 0 online
FPC 1 PIC 0 already transitioning to online
```

When the same PIC is set to a different state while the transition is in progress, you are provided feedback on the status of your request.

```
user@switch> request chassis pic fpc-slot 1 pic-slot 0 offline
FPC 1, PIC 0 already transitioning to online. Please retry later.
```

**Options**    **offline**—Take the PIC offline.

**online**—Bring the PIC online.

**fpc-slot *slot-number***—Flexible PIC Concentrator (FPC) slot number. Replace *slot-number* with a value appropriate for your router or switch:

- ACX4000 routers—1 or 2.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—0.
  - EX4200 switches in a Virtual Chassis configuration—0 through 9 (switch's member ID).
  - EX8208 switches—0 through 7 (line card).
  - EX8216 switches—0 through 15 (line card).
- M5, M7i, M10, and M10i routers—0 or 1.
- M20 routers—0 through 3.
- M40 and M40e routers—0 through 7.
- M120 routers—0 through 5.
- M160 routers—0 through 7.
- M320 routers—0 through 7.
- MX 5, MX10, and MX40 routers—0 or 1.
- MX80 routers—0 or 1.
- MX240 routers—0 through 2
- MX480 routers—0 through 5
- MX2020 routers—0 through 19.

- MX2010 routers—0 through 9.
- MX960 routers—0 through 11.
- PTX5000 routers—0 or 1.
- T Series routers—0 through 7.
- TX Matrix and TX Matrix Plus routers only—On a TX Matrix router, if you specify the number of the T640 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, replace **slot-number** with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the **number** of the T1600 or T4000 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, for the FPC slot number, replace **slot-number** with a value from 0 through 31. On a TX Matrix Plus router with 3D SIBs to assign the FPC slot number, replace **slot-number** with a value from 0 through 63. For example, the following commands have the same result:

```
user@host> request chassis pic fpc-slot 1 lcc 1 pic-slot 0 offline
user@host> request chassis pic fpc-slot 9 pic-slot 0 offline
```

- QFX5100 standalone switches—0.

**pic-slot slot-number**—PIC slot number.

- EX3200 and EX4200 switches—0 for built-in network interfaces and 1 for interfaces on uplink modules.
- EX8208 and EX8216 switches—0.
- M Series routers—0, 1, 2, or 3
- MX960 router—**slot-number** corresponds to the slot number of the Packet Forwarding Engine.
- PTX5000 routers—0 or 1.
- T320 router—0 or 1.
- T640 router—0, 1, 2, or 3.
- T1600 router —0, 1, 2, or 3.
- T4000 router—0, 1, 2, or 3.
- QFX5100 standalone switches—0, 1, or 2. PIC 0 is used for all interfaces that are not configured on expansion modules, and PIC 1 and PIC 2 are used for interfaces configured on expansion modules.

**lcc number**—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**Required Privilege Level** maintenance

**Related Documentation**

- *show chassis hardware*
- [show chassis pic on page 1211](#)
- *Configuring the PIC Type*

**List of Sample Output** [request chassis pic on page 95](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

`request chassis pic`

```
user@host> request chassis pic pic-slot 0 online fpc-slot 0
FPC 0, PIC 0 is already online
```

## request chassis routing-engine master

---

<b>List of Syntax</b>	<a href="#">Syntax on page 96</a> <a href="#">Syntax (M Series, MX Series, T Series Routers) on page 96</a> <a href="#">Syntax (TX Matrix Routers) on page 96</a> <a href="#">Syntax (TX Matrix Plus Routers) on page 96</a> <a href="#">Syntax (MX Series Virtual Chassis) on page 96</a> <a href="#">Syntax (QFX Series) on page 96</a>
<b>Syntax</b>	request chassis routing-engine master (acquire   release   switch) <force> <no-confirm>
<b>Syntax (M Series, MX Series, T Series Routers)</b>	request chassis routing-engine master (acquire   release   switch <check>) <no-confirm>
<b>Syntax (TX Matrix Routers)</b>	request chassis routing-engine master (acquire   release   switch) (lcc <i>number</i>   scc   all-chassis) <force> <no-confirm>
<b>Syntax (TX Matrix Plus Routers)</b>	request chassis routing-engine master (acquire   release   switch) (lcc <i>number</i>   sfc   all-chassis   all-lcc) <force> <no-confirm>
<b>Syntax (MX Series Virtual Chassis)</b>	request chassis routing-engine master (acquire   release   switch <check>) <all-members> <local> <member <i>member-id</i> > <no-confirm>
<b>Syntax (QFX Series)</b>	request chassis routing-engine master (release   switch) <check> <interconnect-device <i>name</i> > <node-group <i>name</i> > <no-confirm>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <b>all-chassis</b> option added in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.3 for QFX Series. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
<b>Description</b>	For routers or switches with multiple Routing Engines, control which Routing Engine is the master.





**CAUTION:** (Routing matrix based on the TX Matrix or TX Matrix Plus routers only) Within the routing matrix, we recommend that all Routing Engines run the same Junos OS Release. If you run different releases on the Routing Engines and a change in mastership occurs on any backup Routing Engine in the routing matrix, one or all routers (in a routing matrix based on the TX Matrix router or in a routing matrix based on a TX Matrix Plus router) might become logically disconnected from the TX Matrix router and cause data loss. For more information, see the [TX Matrix Router Hardware Guide](#) or the *Junos OS High Availability Library for Routing Devices*.



**NOTE:** Successive graceful Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

If the router or switch displays a warning message similar to “Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset,” do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines that were not ready for graceful switchover are reset. None of the Flexible PIC concentrators (FPCs) should spontaneously restart. We recommend that you wait until the warning no longer appears and then proceed with the switchover.

You will receive an error message stating “Command aborted. Not ready for mastership switch, try after n seconds” when this command is re-entered before 240 seconds have elapsed on EX Series switches.



**NOTE:** On a QFabric system, to avoid traffic loss on the network Node group, switch mastership of the routing engine to the backup routing engine, and then reboot.

**Options** **acquire**—Attempt to become the master Routing Engine.

**release**—Request that the other Routing Engine become the master.

**switch**—Toggle mastership between Routing Engines.



**NOTE:** The **acquire** option should be used with caution because acquiring a Routing Engine may result in a corrupted database. If possible, use the **switch** option instead.

The **acquire**, **release**, and **switch** options have the following suboptions:

**all-chassis**—(TX Matrix and TX Matrix Plus routers only) On a routing matrix composed of a TX Matrix router and the attached T640 routers, switch mastership on all the Routing Engines in the routing matrix. Likewise, on a routing matrix composed of a TX Matrix Plus router and the attached T1600 or T4000 routers, switch mastership on all the Routing Engines in the routing matrix.

**all-lcc**—(TX Matrix Plus routers only) Request to acquire mastership for all line-card chassis (LCC).

**all-members**—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines in all member routers of the Virtual Chassis configuration.

**check**—(QFabric systems, MX104, MX480, MX960, MX2010, and MX2020 routers, and PTX5000 routers only) (Optional) Available only with the **switch** option. Check graceful switchover status of the standby Routing Engine before toggling mastership between Routing Engines.

**interconnect-device *name***—(QFabric systems only) (Optional) Control Routing Engine mastership on the Routing Engines on an Interconnect device.

**lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines in the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines of the specified member in the Virtual Chassis Configuration. Replace *member-id* with a value of 0 or 1.

**no-confirm**—(Optional) Do not request confirmation for the switch.

**node-group *name***—(QFabric systems only) (Optional) Control Routing Engine mastership on the Routing Engines on a Node group.

**scc**—(TX Matrix routers only) TX Matrix (switch-card chassis).

**sfc**—(TX Matrix Plus routers only) TX Matrix Plus router (or switch-fabric chassis).

**force**—(Optional) Available only with the **acquire** option. Force the change to a new master Routing Engine.



**NOTE:** The **force** option is not supported on the M Series, MX Series, or T Series routers.

**Additional Information** Because both Routing Engines are always running, the transition from one to the other as the master Routing Engine is immediate. However, the changeover interrupts communication to the System and Switch Board (SSB). The SSB takes several seconds to reinitialize the Flexible PIC Concentrators (FPCs) and restart the PICs. Interior gateway protocol (IGP) and BGP convergence times depend on the specific network environment.

By default, the Routing Engine in slot 0 (**RE0**) is the master and the Routing Engine in slot 1 (**RE1**) is the backup. To change the default master Routing Engine, include the **routing-engine** statement at the **[edit chassis redundancy]** hierarchy level in the configuration. For more information, see the *Junos OS Administration Library for Routing Devices*

To have the backup Routing Engine become the master Routing Engine, use the **request chassis routing-engine master switch** command. If you use this command to change the master and then restart the chassis software for any reason, the master reverts to the default setting.



**NOTE:** Although the configurations on the two Routing Engines do not have to be the same and are not automatically synchronized, we recommend making both configurations the same.

**Required Privilege Level** maintenance

**Related Documentation**

- [show chassis routing-engine on page 1227](#)
- *Configuring Routing Engine Redundancy*
- *Switching the Global Master and Backup Roles in a Virtual Chassis Configuration*

**List of Sample Output**

- [request chassis routing-engine master acquire on page 100](#)
- [request chassis routing-engine master switch on page 100](#)
- [request chassis routing-engine master switch check on page 100](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request chassis routing-engine master acquire

```
user@host> request chassis routing-engine master acquire

warning: Traffic will be interrupted while the PFE is re-initialized

warning: The other routing engine's file system could be corrupted

Reset other routing engine and become master ? [yes,no] (no)
```

### request chassis routing-engine master switch

```
user@host> request chassis routing-engine master switch

warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between Routing Engines ? [yes,no] (no) yes

Resolving mastership...
Complete. The other Routing Engine becomes the master.
```

Switch mastership back to the local Routing Engine:

```
user@host> request chassis routing-engine master switch

warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between routing engines ? [yes,no] (no) yes

Resolving mastership...
Complete. The local routing engine becomes the master.
```

### request chassis routing-engine master switch check

Usage shown for M Series, MX Series, and T Series routers.

```
{master}[edit]
```

```
user@host> request chassis routing-engine master switch check
```

```
warning: Standby Routing Engine is not ready for graceful switchover.
```

```
{master}[edit]
```

```
user@host> request chassis routing-engine master switch check
Switchover Ready
```

You can similarly check the backup Routing Engine.

## request system halt

<b>List of Syntax</b>	<a href="#">Syntax on page 101</a> <a href="#">Syntax (EX Series Switches) on page 101</a> <a href="#">Syntax (PTX Series) on page 101</a> <a href="#">Syntax (TX Matrix Router) on page 101</a> <a href="#">Syntax (TX Matrix Plus Router) on page 101</a> <a href="#">Syntax (MX Series Router) on page 102</a> <a href="#">Syntax (QFX Series) on page 102</a>
<b>Syntax</b>	<pre>request system halt &lt;at <i>time</i>&gt; &lt;backup-routing-engine&gt; &lt;both-routing-engines&gt; &lt;other-routing-engine&gt; &lt;in <i>minutes</i>&gt; &lt;media (compact-flash   disk   removable-compact-flash   usb)&gt; &lt;message "<i>text</i>"&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>request system halt &lt;all-members&gt; &lt;at <i>time</i>&gt; &lt;backup-routing-engine&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;local&gt; &lt;media (external   internal)&gt; &lt;member <i>member-id</i>&gt; &lt;message "<i>text</i>"&gt; &lt;other-routing-engine&gt; &lt;slice <i>slice</i>&gt;</pre>
<b>Syntax (PTX Series)</b>	<pre>request system halt &lt;at <i>time</i>&gt; &lt;backup-routing-engine&gt; &lt;both-routing-engines&gt; &lt;other-routing-engine&gt; &lt;in <i>minutes</i>&gt; &lt;media (compact-flash   disk)&gt; &lt;message "<i>text</i>"&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>request system halt &lt;all-lcc   lcc <i>number</i>   scc&gt; &lt;at <i>time</i>&gt; &lt;backup-routing-engine&gt; &lt;both-routing-engines&gt; &lt;other-routing-engine&gt; &lt;in <i>minutes</i>&gt; &lt;media (compact-flash   disk)&gt; &lt;message "<i>text</i>"&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>request system halt &lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt;</pre>

	<div>&lt;at <i>time</i>&gt; &lt;backup-routing-engine&gt; &lt;both-routing-engines&gt; &lt;other-routing-engine&gt; &lt;in <i>minutes</i>&gt; &lt;media (compact-flash   disk)&gt; &lt;message "text"&gt;</div>
Syntax (MX Series Router)	<div>request system halt &lt;all-members&gt; &lt;at <i>time</i>&gt; &lt;backup-routing-engine&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;local&gt; &lt;media (external   internal)&gt; &lt;member <i>member-id</i>&gt; &lt;message "text"&gt; &lt;other-routing-engine&gt;</div>
Syntax (QFX Series)	<div>request system halt &lt;all-members&gt; &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;director-device <i>director-device-id</i>&gt; &lt;in <i>minutes</i>&gt; &lt;local&gt; &lt;media &gt; &lt;member <i>member-id</i>&gt; &lt;message "text"&gt; &lt;other-routing-engine&gt; &lt;slice <i>slice</i>&gt;</div>
Release Information	<div>Command introduced before Junos OS Release 7.4. <b>other-routing-engine</b> option introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. <b>director-device</b> option introduced for QFabric systems in Junos OS Release 12.2. <b>backup-routing-engine</b> option introduced in Junos OS Release 13.1.</div>
Description	<div>Stop the router or switch software.</div>



**NOTE:** When you issue this command on an individual component—for example, a Node device—in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member of a Node group, issue this command with the **member** option on the Node device CLI, because you cannot issue this command from the QFabric CLI. Also, issuing this command might cause traffic loss on an individual component.

When you issue this command on a QFX5100 switch, you are not prompted to reboot. You must power cycle the switch to reboot.

**Options** **none**—Stop the router or switch software immediately.

**all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Halt all chassis.

**all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, halt all T1600 or T4000 routers connected to the TX Matrix Plus router.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Halt all members of the Virtual Chassis configuration.

**at time** —(Optional) Time at which to stop the software, specified in one of the following ways:

- **now**—Stop the software immediately. This is the default.
- **+minutes**—Number of minutes from now to stop the software.
- **yymmddhhmm**—Absolute time at which to stop the software, specified as year, month, day, hour, and minute.
- **hh:mm**—Absolute time on the current day at which to stop the software.

**backup-routing-engine**—(Optional) Halt the backup Routing Engine. This command halts the backup Routing Engine, regardless from which Routing Engine the command is executed. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. If you issue the command from the backup Routing Engine, the backup Routing Engine is halted.

**both-routing-engines**—(Optional) Halt both Routing Engines at the same time.

**director-device** *director-device-id*—(QFabric systems only) Halt a specific Director device.

**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, halt a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Halt the local Virtual Chassis member.

**in *minutes***—(Optional) Number of minutes from now to stop the software. This option is an alias for the at *+minutes* option.

**media (compact-flash | disk | removable-compact-flash | usb)**—(Optional) Boot medium for the next boot. (The options **removable-compact-flash** and **usb** pertain to J Series routers only.)

**media (external | internal)**—(EX Series and QFX Series switches and MX Series routers only) (Optional) Halt the boot media:

- **external**—Halt the external mass storage device.
- **internal**—Halt the internal flash device.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Halt the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**message "*text*"**—(Optional) Message to display to all system users before stopping the software.

**other-routing-engine**—(Optional) Halt the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

**scc**—(TX Matrix routers only) (Optional) Halt the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Halt the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

**slice *slice***—(EX Series and QFX Series switches only) (Optional) Halt a partition on the boot media. This option has the following suboptions:

- 1—Halt partition 1.
- 2—Halt partition 2.



- **alternate**—Reboot from the alternate partition.

**Additional Information** On the M7i router, the **request system halt** command does not immediately power down the Packet Forwarding Engine. The power-down process can take as long as 5 minutes.

On a TX Matrix router and TX Matrix Plus router if you issue the **request system halt** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are halted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are halted.



**NOTE:** If you have a router or switch with two Routing Engines and you want to shut the power off to the router or switch or remove a Routing Engine, you must first halt the backup Routing Engine (if it has been upgraded), and then halt the master Routing Engine. To halt a Routing Engine, issue the **request system halt** command. You can also halt both Routing Engines at the same time by issuing the **request system halt both-routing-engines** command.

**Required Privilege Level** maintenance

**Related Documentation**

- [clear system reboot on page 84](#)
- [request system power-off on page 108](#)
- *Rebooting and Halting a Device*
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

- [request system halt on page 106](#)
- [request system halt \(In 2 Hours\) on page 106](#)
- [request system halt \(Immediately\) on page 106](#)
- [request system halt \(At 1:20 AM\) on page 106](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system halt

```
user@host> request system halt
Halt the system ? [yes,no] (no) yes

*** FINAL System shutdown message from root@section2 ***
System going down IMMEDIATELY
Terminated
...
syncing disks... 11 8 done
The operating system has halted.
Please press any key to reboot.
```

### request system halt (In 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request that the system stop 2 hours from now:

```
user@host> request system halt at +120
user@host> request system halt in 120
user@host> request system halt at 19:00
```

### request system halt (Immediately)

```
user@host> request system halt at now
```

### request system halt (At 1:20 AM)

To stop the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system halt at yymdd120
request system halt at 120
Halt the system at 120? [yes,no] (no) yes
```

## request system logout

<b>Syntax</b>	<code>request system logout (pid <i>pid</i>   terminal <i>terminal</i>   user <i>username</i>) &lt;all&gt;</code>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Log out users from the router or switch and the configuration database. If a user held the <b>configure exclusive</b> lock, this command clears the exclusive lock.
<b>Options</b>	<p><b>all</b>—(Optional) Log out all sessions owned by a particular PID, terminal session, or user. (On a TX Matrix or TX Matrix Plus router, this command is broadcast to all chassis.)</p> <p><b>pid <i>pid</i></b>—Log out the user session using the specified management process identifier (PID). The PID type must be management process.</p> <p><b>terminal <i>terminal</i></b>—Log out the user for the specified terminal session.</p> <p><b>user <i>username</i></b>—Log out the specified user.</p>
<b>Required Privilege Level</b>	configure
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Junos OS Administration Library for Routing Devices</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system logout on page 107</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output



### request system logout

```
user@host> request system logout user tammy all
Connection closed by foreign host.
```

## request system power-off

---

<b>List of Syntax</b>	<a href="#">Syntax on page 108</a> <a href="#">Syntax (EX Series Switches) on page 108</a> <a href="#">Syntax (TX Matrix Router) on page 108</a> <a href="#">Syntax (TX Matrix Plus Router) on page 108</a> <a href="#">Syntax (MX Series Router) on page 108</a> <a href="#">Syntax (QFX Series) on page 109</a>
<b>Syntax</b>	<code>request system power-off</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;other-routing-engine&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;media (compact-flash   disk   removable-compact-flash   usb)&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code>
<b>Syntax (EX Series Switches)</b>	<code>request system power-off</code> <code>&lt;all-members&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;local&gt;</code> <code>&lt;media (external   internal)&gt;</code> <code>&lt;member <i>member-id</i>&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code> <code>&lt;other-routing-engine&gt;</code> <code>&lt;slice <i>slice</i>&gt;</code>
<b>Syntax (TX Matrix Router)</b>	<code>request system power-off</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;other-routing-engine&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;media (compact-flash   disk)&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code>
<b>Syntax (TX Matrix Plus Router)</b>	<code>request system power-off</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;other-routing-engine&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;media (compact-flash   disk)&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code>
<b>Syntax (MX Series Router)</b>	<code>request system power-off</code> <code>&lt;all-members&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;local&gt;</code>

	<pre> &lt;media (external   internal)&gt; &lt;member <i>member-id</i>&gt; &lt;message "<i>text</i>"&gt; &lt;other-routing-engine&gt; </pre>
<b>Syntax (QFX Series)</b>	<pre> request system power-off &lt;at <i>time</i>&gt; &lt;in <i>minutes</i>&gt; &lt;media (external   internal)&gt; &lt;message "<i>text</i>"&gt; &lt;slice <i>slice</i>&gt; </pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Power off the software.
	<p> <b>NOTE:</b> When you issue this command on an individual component in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member, use the <code>member</code> option. You cannot issue this command from the QFabric CLI.</p>
	<p> <b>NOTE:</b> For a standalone chassis (such as MX Series, PTX Series, and T Series routers), the request to power off the system is applicable only to the Routing Engines. When you request to power off both Routing Engines, all the FPCs in the chassis shut down after approximately 10 minutes and the chassis fans run at full speed. The FPCs shut down because they no longer have communication with the Routing Engines and an Inter-Integrated Circuit (I2C) timeout occurred.</p>
<b>Options</b>	<p><b>none</b>—Power off the router or switch software immediately.</p> <p><b>all-chassis</b>—(Optional) (TX Matrix and TX Matrix Plus router only) Power off all Routing Engines in the chassis.</p> <p><b>all-icc</b>—(Optional) (TX Matrix and TX Matrix Plus router only) On a TX Matrix router, power off all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, power off all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router.</p> <p><b>all-members</b>—(EX4200 switches and MX Series routers only) (Optional) Power off all members of the Virtual Chassis configuration.</p> <p><b>at <i>time</i></b>—(Optional) Time at which to power off the software, specified in one of the following ways:</p>

- **now**—Power off the software immediately. This is the default.
- **+minutes**—Number of minutes from now to power off the software.
- **yymmddhhmm**—Absolute time at which to power off the software, specified as year, month, day, hour, and minute.
- **hh:mm**—Absolute time on the current day at which to power off the software.

**both-routing-engines**—(Optional) Power off both Routing Engines at the same time.

**in minutes**—(Optional) Number of minutes from now to power off the software. This option is an alias for the **at +minutes** option.

**lcc number**—(Optional) (TX Matrix and TX Matrix Plus router only) On a TX Matrix router, power off a T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, power off a specific router that is connected to the TX Matrix Plus router. Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Power off the local Virtual Chassis member.

**media (compact-flash | disk | removable-compact-flash | usb)**—(Optional) Boot medium for the next boot. (The options **removable-compact-flash** and **usb** pertain to the J Series routers only.)

**media (external | internal)**—(EX Series and QFX Series switches and MX Series routers only) (Optional) Power off the boot media:

- **external**—Power off the external mass storage device.
- **internal**—Power off the internal flash device.

**member member-id**—(EX4200 switches and MX Series routers only) (Optional) Power off the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**message "text"**—(Optional) Message to display to all system users before powering off the software.

**other-routing-engine**—(Optional) Power off the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

**scc**—(Optional) (TX Matrix router only) Power off only the master Routing Engine or the backup Routing Engine on the TX Matrix router (or switch-card chassis). If you issue the command from the master Routing Engine, the master SCC is powered off. If you issue the command from the backup Routing Engine, the backup SCC is powered off.

**sfc number**—(Optional) (TX Matrix Plus router only) Power off only the master Routing Engine or the backup Routing Engine on the TX Matrix Plus router (or switch-fabric chassis). If you issue the command from the master Routing Engine, the master SFC is powered off. If you issue the command from the backup Routing Engine, the backup SFC is powered off. Replace *number* with zero.

**slice slice**—(EX Series and QFX Series switches only) (Optional) Power off a partition on the boot media. This option has the following suboptions:

- **1**—Power off partition 1.
- **2**—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

**Additional Information** On a routing matrix composed of a TX Matrix router and T640 routers, if you issue the **request system power-off** command on the TX Matrix master Routing Engine, all the master Routing Engines connected to the routing matrix are powered off. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are powered off.

Likewise, on a routing matrix composed of a TX Matrix Plus router and T1600 routers, if you issue the **request system power-off** command on the TX Matrix Plus master Routing Engine, all the master Routing Engines connected to the routing matrix are powered off. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are powered off.

If you issue the **request system power-off both-routing-engines** command on the TX Matrix or TX Matrix Plus router, all the Routing Engines on the routing matrix are powered off.

**Required Privilege Level** maintenance

**List of Sample Output** [request system power-off on page 112](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system power-off

```
user@host> request system power-off message "This router will be powered off in 30 minutes.  
Please save your data and log out immediately."  
warning: This command will not halt the other routing-engine.  
If planning to switch off power, use the both-routing-engines option.  
Power Off the system ? [yes,no] (no) yes  
  
*** FINAL System shutdown message from remote@nutmeg ***  
System going down IMMEDIATELY  
  
This router will be powered off in 30 minutes. Please save your data and log out  
immediately.  
  
Shutdown NOW!  
[pid 5177]
```



## request system reboot

<b>List of Syntax</b>	<a href="#">Syntax on page 113</a> <a href="#">Syntax (EX Series Switches) on page 113</a> <a href="#">Syntax (TX Matrix Router) on page 113</a> <a href="#">Syntax (TX Matrix Plus Router) on page 113</a> <a href="#">Syntax (MX Series Router) on page 113</a>
<b>Syntax</b>	<pre>request system reboot &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;media (compact-flash   disk   removable-compact-flash   usb)&gt; &lt;message "<i>text</i>"&gt; &lt;other-routing-engine&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>request system reboot &lt;all-members&gt; &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;local&gt; &lt;media (external   internal)&gt; &lt;member <i>member-id</i>&gt; &lt;message "<i>text</i>"&gt; &lt;other-routing-engine&gt; &lt;slice <i>slice</i>&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>request system reboot &lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt; &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;media (compact-flash   disk)&gt; &lt;message "<i>text</i>"&gt; &lt;other-routing-engine&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>request system reboot &lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt; &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;media (compact-flash   disk)&gt; &lt;message "<i>text</i>"&gt; &lt;other-routing-engine&gt; &lt;partition (1   2   alternate)&gt;</pre>
<b>Syntax (MX Series Router)</b>	<pre>request system reboot &lt;all-members&gt; &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;local&gt;</pre>

<media (external | internal)>  
<member *member-id*>  
<message "*text*">  
<other-routing-engine>

<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Option <b>other-routing-engine</b> introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Option <b>sfc</b> introduced for the TX Matrix Plus router in Junos OS Release 9.6. Option <b>both-routing-engines</b> introduced in Junos OS Release 12.1.
<b>Description</b>	Reboot the software.
<b>Options</b>	<p><b>none</b>—Reboot the software immediately.</p> <p><b>all-chassis</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all routers connected to the TX Matrix or TX Matrix Plus router, respectively.</p> <p><b>all-lcc</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively.</p> <p><b>all-members</b>—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on all members of the Virtual Chassis configuration.</p> <p><b>at time</b>—(Optional) Time at which to reboot the software, specified in one of the following ways:</p> <ul style="list-style-type: none"><li>• <b>now</b>—Stop or reboot the software immediately. This is the default.</li><li>• <b>+minutes</b>—Number of minutes from now to reboot the software.</li><li>• <b>yymmddhhmm</b>—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.</li><li>• <b>hh:mm</b>—Absolute time on the current day at which to stop the software, specified in 24-hour time.</li></ul> <p><b>both-routing-engines</b>—(Optional) Reboot both Routing Engines at the same time.</p> <p><b>in minutes</b>—(Optional) Number of minutes from now to reboot the software. This option is an alias for the <b>at +minutes</b> option.</p> <p><b>lcc number</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis number. Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"><li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li><li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li></ul>

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the local Virtual Chassis member.

**media (compact-flash | disk | removable-compact-flash | usb)**—(Optional) Boot medium for next boot. (The options **removable-compact-flash** and **usb** pertain to the J Series routers only.)

**media (external | internal)**—(EX Series switches and MX Series routers only) (Optional) Reboot the boot media:

- **external**—Reboot the external mass storage device.
- **internal**—Reboot the internal flash device.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

**message "*text*"**—(Optional) Message to display to all system users before stopping or rebooting the software.

**other-routing-engine**—(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is rebooted.

**partition**—(TX Matrix Plus routers only) (Optional) Reboot using the specified partition on the boot media. This option has the following suboptions:

- 1—Reboot from partition 1.
- 2—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition.

**scc**—(TX Matrix routers only) (Optional) Reboot the Routing Engine on the TX Matrix switch-card chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Reboot the Routing Engine on the TX Matrix Plus switch-fabric chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted. Replace ***number*** with 0.

**slice *slice***—(EX Series switches only) (Optional) Reboot a partition on the boot media. This option has the following suboptions:

- 1—Power off partition 1.
- 2—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

**Additional Information** Reboot requests are recorded in the system log files, which you can view with the **show log** command (see [show log](#)). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the **show system processes** command (see [show system processes](#)).

On a TX Matrix or TX Matrix Plus router, if you issue the **request system reboot** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are rebooted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are rebooted.



**NOTE:** Before issuing the **request system reboot** command on a TX Matrix Plus router with no options or the **all-chassis**, **all-lcc**, **lcc number**, or **sfc** options, verify that master Routing Engine for all routers in the routing matrix are in the same slot number. If the master Routing Engine for a line-card chassis is in a different slot number than the master Routing Engine for a TX Matrix Plus router, the line-card chassis might become logically disconnected from the routing matrix after the **request system reboot** command.



**NOTE:** To reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the master Routing Engine.

**Required Privilege Level** maintenance

**Related Documentation**

- [clear system reboot on page 84](#)
- [request system halt on page 101](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

- [request system reboot on page 117](#)
- [request system reboot \(at 2300\) on page 117](#)
- [request system reboot \(in 2 Hours\) on page 117](#)
- [request system reboot \(Immediately\) on page 117](#)
- [request system reboot \(at 1:20 AM\) on page 117](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system reboot

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

### request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes
```

```
shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

### request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
user@host> request system reboot in 120
user@host> request system reboot at 19:00
```

### request system reboot (Immediately)

```
user@host> request system reboot at now
```

### request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

## request system reboot

---

<b>Syntax</b>	<pre>request system reboot &lt;all-members   local   member member-id&gt; &lt;at time&gt; &lt;in minutes&gt; &lt;media (external   internal)&gt; &lt;message "text"&gt; &lt;slice (1   2   alternate)&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Option <b>partition</b> changed to <b>slice</b> in Junos OS Release 10.0 for EX Series switches.
<b>Description</b>	<p>Reboot the Junos OS.</p> <p>Reboot requests are recorded in the system log files, which you can view with the <b>show log</b> command. You can view the process names with the <b>show system processes</b> command.</p>
<b>Options</b>	<p><b>none</b>—Reboots the software immediately.</p> <p><b>all-members   local   member member-id</b>—(EX4200 switch only) (Optional) Specify which member of the Virtual Chassis to reboot:</p> <ul style="list-style-type: none"><li>• <b>all-members</b>—Reboots each switch that is a member of the Virtual Chassis.</li><li>• <b>local</b>—Reboots the local switch, meaning the switch you are logged into, only.</li><li>• <b>member member-id</b>—Reboots the specified member switch of the Virtual Chassis.</li></ul> <p><b>at time</b>—(Optional) Time at which to reboot the software, specified in one of the following ways:</p> <ul style="list-style-type: none"><li>• <b>+minutes</b>—Number of minutes from now to reboot the software.</li><li>• <b>hh:mm</b>—Absolute time on the current day at which to reboot the software, specified in 24-hour time.</li><li>• <b>now</b>—Stop or reboot the software immediately. This is the default.</li><li>• <b>yymmddhhmm</b>—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.</li></ul> <p><b>in minutes</b>—(Optional) Number of minutes from now to reboot the software. This option is an alias for the <b>at +minutes</b> option.</p> <p><b>media (external   internal)</b>—(Optional) Boot medium for the next boot. The external option reboots the switch using a software package stored on an external boot source, such as a USB flash drive. The internal option reboots the switch using a software package stored in an internal memory source.</p> <p><b>message "text"</b>—(Optional) Message to display to all system users before rebooting the software.</p>

**slice (1 | 2 | alternate)**—(Optional) Reboot using the specified partition on the boot media.

This option has the following suboptions:

- **1**—Reboot from partition 1.
- **2**—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition, which is the partition that did not boot the switch at the last bootup.

**Required Privilege Level** maintenance

**Related Documentation** [• clear system reboot on page 84](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system reboot

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

### request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes

shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

### request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
user@host> request system reboot in 120
user@host> request system reboot at 19:00
```

### request system reboot (Immediately)

```
user@host> request system reboot at now
```

### request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

## request system scripts convert

---

<b>Syntax</b>	<b>request system scripts convert</b> (slax-to-xslt   xslt-to-slax) <b>source</b> <i>source/filename</i> <b>destination</b> <i>destination/&lt;filename&gt;</i> <b>&lt;partial&gt;</b> <b>&lt;version (1.0   1.1)&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 8.2. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>partial</b> and <b>version</b> options added in Junos OS Release 12.2.
<b>Description</b>	Convert an Extensible Stylesheet Language Transformations (XSLT) script to Stylesheet Language, Alternative syntaX (SLAX), or convert a SLAX script to XSLT.
<b>Options</b>	<b>destination</b> <i>destination/&lt;filename&gt;</i> —Specify a destination for the converted file. Optionally, you can specify a filename for the converted file. If you do not specify a filename, the software assigns one automatically. The default destination filename is <b>SLAX-Conversion-Temp</b> or <b>slax-temp</b> depending on the Junos OS release, with a randomly generated, five-character, alpha-numeric extension. For example, the software converts a source file called <b>test.xml</b> to <b>slax-temp.kWwQk</b> . The software converts a source file called <b>test1.slax</b> to <b>slax-temp.zN61h</b> .  <b>partial</b> —(Optional) Convert partial script input.  <b>slax-to-xslt</b> —Convert a SLAX script to XSLT.  <b>source</b> <i>source/filename</i> —Specify a source file that you want to convert.  <b>version</b> —(Optional) Specify the SLAX version listed in the version statement of the generated script for XSLT-to-SLAX conversions. Acceptable values are 1.0 and 1.1. The default is 1.1.  <b>xslt-to-slax</b> —Convert an XSLT script to SLAX.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Converting Scripts Between SLAX and XSLT</i></li></ul>
<b>List of Sample Output</b>	<a href="#">request system scripts convert slax-to-xslt on page 120</a> <a href="#">request system scripts convert xslt-to-slax on page 121</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system scripts convert slax-to-xslt

```
user@host> request system scripts convert slax-to-xslt source /var/db/scripts/op/script1.slax
destination /var/db/scripts/op
conversion complete
```



**request system scripts convert xslt-to-slax**

```
user@host> request system scripts convert xslt-to-slax source /var/db/scripts/commit/script1.xml  
destination /var/db/scripts/commit partial version 1.0  
conversion complete
```

## request system scripts refresh-from commit

---

<b>Syntax</b>	<code>request system scripts refresh-from commit file <i>file-name</i> url <i>url-path</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.1 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Automatically download the initial Junos OS configuration and a set of standard commit scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.</p> <p>The Junos XML management protocol equivalent for this operational mode command is:</p> <pre>&lt;request-script-refresh-from&gt;   &lt;type&gt;commit&lt;/type&gt;   &lt;file&gt;file-name&lt;/file&gt;   &lt;URL&gt;URL&lt;/URL&gt; &lt;/request-script-refresh-from&gt;</pre>
<b>Options</b>	<p><b>file <i>file-name</i></b>—Name of the file to be downloaded.</p> <p><b>url <i>url-path</i></b>—URL of the file to be downloaded.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Automatic Refreshing of Scripts on EX Series Switches on page 613</a></li><li>• <a href="#">Junos OS Junos XML Management Protocol Guide</a></li><li>• <a href="#">Junos OS NETCONF XML Management Protocol Guide</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request system scripts refresh-from commit file config.txt url http://host1.juniper.net on page 122</a>

### Sample Output

`request system scripts refresh-from commit file config.txt url http://host1.juniper.net`

```
user@switch> request system scripts refresh-from commit file config.txt url
http://host1.juniper.net
user@switch>
```

## request system scripts refresh-from event

<b>Syntax</b>	<code>request system scripts refresh-from event file <i>file-name</i> url <i>url-path</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.1 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Automatically download the initial Junos OS configuration and a set of standard event scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.</p> <p>The Junos XML management protocol equivalent for this operational mode command is:</p> <pre>&lt;request-script-refresh-from&gt;   &lt;type&gt;event&lt;/type&gt;   &lt;file&gt;file-name&lt;/file&gt;   &lt;URL&gt;URL&lt;/URL&gt; &lt;/request-script-refresh-from&gt;</pre>
<b>Options</b>	<p><b>file <i>file-name</i></b>—Name of the file to be downloaded.</p> <p><b>url <i>url-path</i></b>—URL of the file to be downloaded.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Automatic Refreshing of Scripts on EX Series Switches on page 613</a></li> <li>• <a href="#">Junos OS Junos XML Management Protocol Guide</a></li> <li>• <a href="#">Junos OS NETCONF XML Management Protocol Guide</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system scripts refresh-from event file config.txt url http://host1.juniper.net on page 123</a>

### Sample Output

`request system scripts refresh-from event file config.txt url http://host1.juniper.net`

```
user@switch> request system scripts refresh-from event file config.txt url http://host1.juniper.net
user@switch>
```

## request system scripts refresh-from op

---

<b>Syntax</b>	<code>request system scripts refresh-from op file <i>file-name</i> url <i>url-path</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.1 for EX Series switches.
<b>Description</b>	<p>Automatically download the initial Junos OS configuration and a set of standard op scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.</p> <p>The Junos XML management protocol equivalent for this operational mode command is:</p> <pre>&lt;request-script-refresh-from&gt;   &lt;type&gt;op&lt;/type&gt;   &lt;file&gt;file-name&lt;/file&gt;   &lt;URL&gt;URL&lt;/URL&gt; &lt;/request-script-refresh-from&gt;</pre>
<b>Options</b>	<p><code>file <i>file-name</i></code>—Name of the file to be downloaded.</p> <p><code>url <i>url-path</i></code>—URL of the file to be downloaded.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Automatic Refreshing of Scripts on EX Series Switches on page 613</a></li><li>• <a href="#">Junos OS Junos XML Management Protocol Guide</a></li><li>• <a href="#">Junos OS NETCONF XML Management Protocol Guide</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request system scripts refresh-from op file config.txt url http://host1.juniper.net on page 124</a>

### Sample Output

`request system scripts refresh-from op file config.txt url http://host1.juniper.net`

```
user@switch> request system scripts refresh-from op file config.txt url http://host1.juniper.net
user@switch>
```

## request system storage cleanup

<b>List of Syntax</b>	<a href="#">Syntax on page 125</a> <a href="#">Syntax (EX Series Switches) on page 125</a> <a href="#">Syntax (MX Series Router) on page 125</a> <a href="#">Syntax (QFX Series) on page 125</a>
<b>Syntax</b>	request system storage cleanup <dry-run>
<b>Syntax (EX Series Switches)</b>	request system storage cleanup <all-members> <dry-run> <local> <member <i>member-id</i> >
<b>Syntax (MX Series Router)</b>	request system storage cleanup <all-members> <dry-run> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	request system storage cleanup <component ( <i>serial number</i>   <i>UUID</i>   all)> <director-group <i>name</i> > <dry-run> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <name-tag <i>name-tag</i> > <node-group <i>name</i> > <prune> <qfabric (component <i>name</i> )   dry-run   name-tag   repository> <repository (core   log)>
<b>Release Information</b>	Command introduced in Junos OS Release 7.4. <b>dry-run</b> option introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Free storage space on the router or switch by rotating log files and proposing a list of files for deletion. User input is required for file deletion. On a QFabric system, you can delete debug files located on individual devices or on the entire QFabric system.
<b>Options</b>	<b>all-members</b> —(EX4200 switches and MX Series routers only) (Optional) Delete files on the Virtual Chassis master Routing Engine only.



**NOTE:** To delete files on the other members of the Virtual Chassis configuration, log in to each backup Routing Engine and delete the files using the `request system storage cleanup local` command.

**component** (*UUID | serial number | all*)—(QFabric systems only) (Optional) Delete files located on individual QFabric system devices or on the entire QFabric system.

**director-group** *name*—(QFabric systems only) (Optional) Delete files on the Director group.

**dry-run**—(Optional) List files proposed for deletion (without deleting them).

**infrastructure** *name*—(QFabric systems only) (Optional) Delete files on the fabric control Routing Engine and fabric manager Routing Engine.

**interconnect-device** *name*—(QFabric systems only) (Optional) Delete files on the Interconnect device.

**local**—(EX4200 switches and MX Series routers only) (Optional) Delete files on the local Virtual Chassis member.

**member** *member-id*—(EX4200 switches and MX Series routers only) (Optional) Delete files on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**name-tag** *name-tag*—(QFabric systems only) (Optional) Delete debug files that match a specific regular expression.

**node-group** *name*—(QFabric systems only) (Optional) Delete files on the Node group.

**prune**—(QFabric systems only) (Optional) Delete debug files located in either the core or log debug repositories of a QFabric system device.

**qfabric component** *name*—(QFabric systems only) (Optional) Delete debug files located in the debug repositories of a QFabric system device.

**repository** (*core | log*)—(QFabric systems only) (Optional) Specify the repository on the QFabric system device for which you want to delete debug files.

**Additional Information** If logging is configured and being used, the **dry-run** option rotates the log files. In that case, the output displays the message “Currently rotating log files, please wait.” If no logging is currently under way, the output displays only a list of files to delete.

**Required Privilege Level** maintenance

**List of Sample Output** [request system storage cleanup dry-run on page 127](#)  
[request system storage cleanup on page 128](#)  
[request system storage cleanup director-group \(QFabric Systems\) on page 128](#)  
[request system storage cleanup infrastructure device-name \(QFabric Systems\) on page 130](#)  
[request system storage cleanup interconnect-device device-name \(QFabric Systems\) on page 131](#)  
[request system storage cleanup node-group group-name \(QFabric Systems\) on page 132](#)

[request system storage cleanup qfabric component device-name \(QFabric Systems\) on page 133](#)

[request system storage cleanup qfabric component device-name repository core \(QFabric Systems\) on page 133](#)

[request system storage cleanup qfabric component all \(QFabric Systems\) on page 133](#)

**Output Fields** [Table 10 on page 127](#) describes the output fields for the **request system storage cleanup** command. Output fields are listed in the approximate order in which they appear.

**Table 10: request system storage cleanup Output Fields**

Field Name	Field Description
List of files to delete:	Shows list of files available for deletion.
Size	Size of the core-dump file.
Date	Last core-dump file modification date and time.
Name	Name of the core-dump file.
Directory to delete:	Shows list of directories available for deletion.
Repository scope:	Repository where core-dump files and log files are stored. The core-dump files are located in the <b>core</b> repository, and the log files are located in the <b>log</b> repository. The default <b>Repository scope</b> is shared since both the <b>core</b> and <b>log</b> repositories are shared by all of the QFabric system devices.
Repository head:	Name of the top-level repository location.
Repository name:	Name of the repository: <b>core</b> or <b>log</b> .
Creating list of debug artifacts to be removed under:	Shows location of files available for deletion.
List of debug artifacts to be removed under:	Shows list of files available for deletion.

## Sample Output

### `request system storage cleanup dry-run`

```
user@host> request system storage cleanup dry-run
Currently rotating log files, please wait.
This operation can take up to a minute.
```

List of files to delete:

	Size	Date	Name
	11.4K	Mar 8 15:00	/var/log/messages.1.gz
	7245B	Feb 5 15:00	/var/log/messages.3.gz
	11.8K	Feb 22 13:00	/var/log/messages.2.gz

```

3926B Mar 16 13:57 /var/log/messages.0.gz
3962B Feb 22 12:47 /var/log/sampled.1.gz
4146B Mar 8 12:20 /var/log/sampled.0.gz
4708B Dec 21 11:39 /var/log/sampled.2.gz
7068B Jan 16 18:00 /var/log/messages.4.gz
13.7K Dec 27 22:00 /var/log/messages.5.gz
890B Feb 22 17:22 /var/tmp/sampled.pkts
65.8M Oct 26 09:10 /var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz
63.1M Oct 26 09:13 /var/sw/pkg/jbundle-7.4R1.7.tgz

```

### request system storage cleanup

```

user@host> request system storage cleanup
Currently rotating log files, please wait.
This operation can take up to a minute.

```

List of files to delete:

	Size	Date	Name
	11.4K	Mar 8 15:00	/var/log/messages.1.gz
	7245B	Feb 5 15:00	/var/log/messages.3.gz
	11.8K	Feb 22 13:00	/var/log/messages.2.gz
	3926B	Mar 16 13:57	/var/log/messages.0.gz
	11.6K	Mar 8 15:00	/var/log/messages.5.gz
	7254B	Feb 5 15:00	/var/log/messages.6.gz
	12.9K	Feb 22 13:00	/var/log/messages.8.gz
	3726B	Mar 16 13:57	/var/log/messages.7.gz
	3962B	Feb 22 12:47	/var/log/sampled.1.gz
	4146B	Mar 8 12:20	/var/log/sampled.0.gz
	4708B	Dec 21 11:39	/var/log/sampled.2.gz
	7068B	Jan 16 18:00	/var/log/messages.4.gz
	13.7K	Dec 27 22:00	/var/log/messages.5.gz
	890B	Feb 22 17:22	/var/tmp/sampled.pkts
	65.8M	Oct 26 09:10	/var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz
	63.1M	Oct 26 09:13	/var/sw/pkg/jbundle-7.4R1.7.tgz

Delete these files ? [yes,no] (yes)

### request system storage cleanup director-group (QFabric Systems)

```

user@switch> request system storage cleanup director-group
List of files to delete:

```

	Size	Date	Name
	4.0K	2011-11-07 05:16:29	/tmp/2064.sfcauth
	4.0K	2011-11-07 05:07:34	/tmp/30804.sfcauth
	4.0K	2011-11-07 04:13:41	/tmp/26792.sfcauth
	4.0K	2011-11-07 04:13:39	/tmp/26432.sfcauth
	0	2011-11-07 07:45:40	/tmp/cluster_cleanup.log
	1.3M	2011-11-07 07:39:11	/tmp/cn_monitor.20111107-052401.log
	4.0K	2011-11-07 07:36:29	/tmp/clustat.28019.log
	4.0K	2011-11-07 07:36:29	/tmp/clustat_x.28019.log
	9.6M	2011-11-07 05:30:24	/tmp/sfc.2.log
	4.0K	2011-11-07 05:28:11	/tmp/mgd-init.1320672491.log
	248K	2011-11-07 05:19:24	/tmp/cn_monitor.20111107-045111.log
	4.0K	2011-11-07 05:17:18	/tmp/clustat.3401.log
	4.0K	2011-11-07 05:17:18	/tmp/clustat_x.3401.log
	8.0K	2011-11-07 04:58:25	/tmp/mgd-init.1320670633.log
	0	2011-11-07 04:54:01	/tmp/mysql_db_install_5.1.37.log
	4.0K	2011-11-07 04:52:08	/tmp/cn_send.log
	0	2011-11-07 04:52:00	/tmp/init_eth0.log



```

4.0K 2011-11-07 04:49:35 /tmp/install_interfaces.sh.log
4.0K 2011-11-07 04:48:15 /tmp/bootstrap.sh.log
160K 2011-11-07 04:47:43 /tmp/bootstrap_cleanup.log
38M 2011-11-07 04:42:42 /tmp/cn_monitor.20111104-110308.log
4.0K 2011-11-07 04:38:47 /tmp/clustat.30913.log
4.0K 2011-11-07 04:38:47 /tmp/clustat_x.30913.log
4.0K 2011-11-07 04:38:03 /tmp/dcf_upgrade.sh.remove.log
4.0K 2011-11-07 04:38:03 /tmp/peer_update.log
4.0K 2011-11-07 04:38:02 /tmp/dcf_upgrade.log
4.0K 2011-11-07 04:38:02 /tmp/perl_mark_upgrade.log
8.0K 2011-11-07 04:13:42 /tmp/install_dcf_rpm.log
4.0K 2011-11-07 04:13:06 /tmp/00_cleanup.sh.1320667986.log
0 2011-11-07 04:13:06 /tmp/ccif_patch_4410_4450.sh.1320667986.log
4.0K 2011-11-07 04:13:06 /tmp/dcf-tools.sh.1320667986.log
0 2011-11-07 04:13:06 /tmp/initial.sh.1320667986.log
0 2011-11-07 04:13:06 /tmp/inventory.sh.1320667986.log
4.0K 2011-11-07 04:13:06 /tmp/qf-db.sh.1320667986.log
4.0K 2011-11-07 04:13:06 /tmp/sfc.sh.1320667986.log
8.0K 2011-11-07 04:13:05 /tmp/jinstall-qfabric.log
8.0K 2011-11-04 11:10:24 /tmp/mgd-init.1320430192.log
4.0K 2011-11-04 11:07:03 /tmp/mysql_dcf_db_install.log
8.0K 2011-11-04 10:55:07 /tmp/ccif_patch_4410_4450.sh.1320429307.log
8.0K 2011-11-04 10:55:07 /tmp/initial.sh.1320429307.log
4.0K 2011-11-04 10:55:07 /tmp/inventory.sh.1320429307.log
8.0K 2011-11-04 10:55:07 /tmp/sfc.sh.1320429307.log
4.0K 2011-11-04 10:54:09 /tmp/ks-script-Ax0tz5.log
4.0K 2011-11-07 04:13:06 /tmp//sfc.sh.1320667986.log
8.0K 2011-11-04 10:55:07 /tmp//sfc.sh.1320429307.log

```

Directory to delete:

```

45M 2011-11-08 10:57:43 /tmp/sfc-captures

```

List of files to delete:

	Size	Date	Name
4.0K	2011-11-08	05:47:47	/tmp/5713.sfcauth
4.0K	2011-11-08	05:14:32	/tmp/14494.sfcauth
4.0K	2011-11-08	05:11:47	/tmp/9978.sfcauth
4.0K	2011-11-08	05:09:37	/tmp/6128.sfcauth
4.0K	2011-11-08	05:04:28	/tmp/29703.sfcauth
4.0K	2011-11-07	11:59:10	/tmp/7811.sfcauth
4.0K	2011-11-07	11:36:08	/tmp/32415.sfcauth
4.0K	2011-11-07	11:30:30	/tmp/22406.sfcauth
4.0K	2011-11-07	11:24:37	/tmp/12131.sfcauth
4.0K	2011-11-07	10:48:42	/tmp/12687.sfcauth
4.0K	2011-11-07	09:27:20	/tmp/31082.sfcauth
4.0K	2011-11-07	07:33:58	/tmp/14633.sfcauth
4.0K	2011-11-07	05:08:25	/tmp/15447.sfcauth
4.0K	2011-11-07	04:12:29	/tmp/26874.sfcauth
4.0K	2011-11-07	04:12:27	/tmp/26713.sfcauth
4.0K	2011-11-07	03:49:17	/tmp/17691.sfcauth
4.0K	2011-11-05	01:32:23	/tmp/5716.sfcauth
4.0K	2011-11-07	08:00:17	/tmp/sfcsnmpd.log
4.0K	2011-11-07	07:57:50	/tmp/cluster_cleanup.log
824K	2011-11-07	07:38:37	/tmp/cn_monitor.20111107-053643.log
4.0K	2011-11-07	07:36:30	/tmp/clustat.18399.log
4.0K	2011-11-07	07:36:30	/tmp/clustat_x.18399.log
4.0K	2011-11-07	07:35:47	/tmp/command_lock.log
4.0K	2011-11-07	05:39:54	/tmp/mgd-init.1320673194.log
92K	2011-11-07	05:19:25	/tmp/cn_monitor.20111107-050412.log
4.0K	2011-11-07	05:17:20	/tmp/clustat.30115.log

```

4.0K  2011-11-07 05:17:20 /tmp/clustat_x.30115.log
8.0K  2011-11-07 05:08:07 /tmp/mgd-init.1320671241.log
4.0K  2011-11-07 05:04:57 /tmp/cn_send.log
0     2011-11-07 05:04:52 /tmp/init_eth0.log
4.0K  2011-11-07 05:02:38 /tmp/install_interfaces.sh.log
4.0K  2011-11-07 05:01:19 /tmp/bootstrap.sh.log
160K  2011-11-07 05:00:47 /tmp/bootstrap_cleanup.log
28M   2011-11-07 04:42:27 /tmp/cn_monitor.20111104-112954.log
4.0K  2011-11-07 04:38:49 /tmp/clustat.6780.log
4.0K  2011-11-07 04:38:49 /tmp/clustat_x.6780.log
4.0K  2011-11-07 04:38:05 /tmp/issue_event.log
4.0K  2011-11-07 04:38:05 /tmp/peer_upgrade_reboot.log
12K   2011-11-07 04:38:05 /tmp/primary_update.log
4.0K  2011-11-07 04:38:04 /tmp/dcf_upgrade.sh.remove.log
4.0K  2011-11-07 04:38:04 /tmp/peer_rexec_upgrade.log
4.0K  2011-11-07 04:13:42 /tmp/peer_install_dcf_rpm.log
4.0K  2011-11-07 04:11:57 /tmp/dcf-tools.sh.1320667917.log
0     2011-11-07 04:11:57 /tmp/initial.sh.1320667917.log
0     2011-11-07 04:11:57 /tmp/inventory.sh.1320667917.log
4.0K  2011-11-07 04:11:57 /tmp/qf-db.sh.1320667917.log
4.0K  2011-11-07 04:11:57 /tmp/sfc.sh.1320667917.log
4.0K  2011-11-07 04:11:56 /tmp/00_cleanup.sh.1320667916.log
0     2011-11-07 04:11:56 /tmp/ccif_patch_4410_4450.sh.1320667916.log
8.0K  2011-11-07 04:11:56 /tmp/jinstall-qfabric.log
4.0K  2011-11-07 04:11:33 /tmp/dcf_upgrade.log
8.0K  2011-11-04 11:53:12 /tmp/mgd-init.1320432782.log
8.0K  2011-11-04 11:06:17 /tmp/ccif_patch_4410_4450.sh.1320429977.log
8.0K  2011-11-04 11:06:17 /tmp/initial.sh.1320429977.log
4.0K  2011-11-04 11:06:17 /tmp/inventory.sh.1320429977.log
8.0K  2011-11-04 11:06:17 /tmp/sfc.sh.1320429977.log
4.0K  2011-11-04 11:05:19 /tmp/ks-script-tnWeb.log
4.0K  2011-11-07 04:11:57 /tmp//sfc.sh.1320667917.log
8.0K  2011-11-04 11:06:17 /tmp//sfc.sh.1320429977.log

```

Directory to delete:

```

49M   2011-11-08 10:45:20 /tmp/sfc-captures

```

### request system storage cleanup infrastructure device-name (QFabric Systems)

```

user@switch> request system storage cleanup infrastructure FC-0
re0:

```

-----

List of files to delete:

	Size	Date	Name
	139B	Nov 8 19:03	/var/log/default-log-messages.0.gz
	5602B	Nov 8 19:03	/var/log/messages.0.gz
	28.4K	Nov 8 10:15	/var/log/messages.1.gz
	35.2K	Nov 7 13:45	/var/log/messages.2.gz
	207B	Nov 7 16:02	/var/log/wtmp.0.gz
	27B	Nov 7 12:14	/var/log/wtmp.1.gz
	184.4M	Nov 7 12:16	/var/sw/pkg/jinstall-dc-re-11.3I20111104_1216_dc-builder-domestic-signed.tgz
	124.0K	Nov 7 15:59	/var/tmp/gres-tp/env.dat
	0B	Nov 7 12:57	/var/tmp/gres-tp/lock
	155B	Nov 7 16:02	/var/tmp/krt_gencfg_filter.txt
	0B	Nov 7 12:35	/var/tmp/last_ccif_update
	1217B	Nov 7 12:15	/var/tmp/loader.conf.preinstall
	184.4M	Nov 6 07:11	/var/tmp/mchassis-install.tgz
	10.8M	Nov 7 12:16	

```

/var/tmp/preinstall/bootstrap-install-11.3I20111104_1216_dc-builder.tar
57.4K Nov 7 12:16 /var/tmp/preinstall/configs-11.3I20111104_1216_dc-builder.tgz

259B Nov 7 12:16 /var/tmp/preinstall/install.conf
734.3K Nov 4 13:46
/var/tmp/preinstall/jboot-dc-re-11.3I20111104_1216_dc-builder.tgz
177.8M Nov 7 12:16
/var/tmp/preinstall/jbundle-dc-re-11.3I20111104_1216_dc-builder-domestic.tgz
124B Nov 7 12:15 /var/tmp/preinstall/metatags
1217B Nov 7 12:16 /var/tmp/preinstall_boot_loader.conf
0B Nov 7 16:02 /var/tmp/rtssdb/if-rtssdb

```

### request system storage cleanup interconnect-device device-name (QFabric Systems)

```

user@switch> request system storage cleanup interconnect IC-WS001
re1:
-----

```

List of files to delete:

	Size	Date	Name
	11B	Nov 7 15:55	/var/jail/tmp/alarmd.ts
	128B	Nov 8 19:06	/var/log/default-log-messages.0.gz
	9965B	Nov 8 19:06	/var/log/messages.0.gz
	15.8K	Nov 8 12:30	/var/log/messages.1.gz
	15.8K	Nov 8 11:00	/var/log/messages.2.gz
	15.7K	Nov 8 07:30	/var/log/messages.3.gz
	15.8K	Nov 8 04:00	/var/log/messages.4.gz
	15.7K	Nov 8 00:30	/var/log/messages.5.gz
	18.7K	Nov 7 21:00	/var/log/messages.6.gz
	17.6K	Nov 7 19:00	/var/log/messages.7.gz
	58.3K	Nov 7 16:00	/var/log/messages.8.gz
	20.3K	Nov 7 15:15	/var/log/messages.9.gz
	90B	Nov 7 15:41	/var/log/wtmp.0.gz
	57B	Nov 7 12:41	/var/log/wtmp.1.gz
	124.0K	Nov 7 15:42	/var/tmp/gres-tp/env.dat
	0B	Nov 7 12:40	/var/tmp/gres-tp/lock
	0B	Nov 7 12:41	/var/tmp/if-rtssdb/env.lock
	12.0K	Nov 7 15:41	/var/tmp/if-rtssdb/env.mem
	132.0K	Nov 7 15:55	/var/tmp/if-rtssdb/shm_usr1.mem
	2688.0K	Nov 7 15:41	/var/tmp/if-rtssdb/shm_usr2.mem
	2048.0K	Nov 7 15:41	/var/tmp/if-rtssdb/trace.mem
	730B	Nov 7 19:57	/var/tmp/juniper.conf+.gz
	155B	Nov 7 15:53	/var/tmp/krt_gencfg_filter.txt
	0B	Nov 7 15:41	/var/tmp/rtssdb/if-rtssdb

```

re0:
-----

```

List of files to delete:

	Size	Date	Name
	11B	Nov 7 15:55	/var/jail/tmp/alarmd.ts
	121B	Nov 8 19:06	/var/log/default-log-messages.0.gz
	16.7K	Nov 8 19:06	/var/log/messages.0.gz
	22.2K	Nov 8 17:45	/var/log/messages.1.gz
	18.4K	Nov 8 17:00	/var/log/messages.2.gz
	21.6K	Nov 8 16:00	/var/log/messages.3.gz
	17.9K	Nov 8 14:30	/var/log/messages.4.gz
	19.4K	Nov 8 13:30	/var/log/messages.5.gz
	18.2K	Nov 8 12:30	/var/log/messages.6.gz

```

20.4K Nov  8 11:30 /var/log/messages.7.gz
21.4K Nov  8 10:15 /var/log/messages.8.gz
21.0K Nov  8 09:00 /var/log/messages.9.gz
19.9K Nov  8 08:13 /var/log/snmp-traps.0.gz
203B Nov  8 15:36 /var/log/wtmp.0.gz
57B Nov  7 12:41 /var/log/wtmp.1.gz
124.0K Nov  7 15:42 /var/tmp/gres-tp/env.dat
0B Nov  7 12:40 /var/tmp/gres-tp/lock
0B Nov  7 12:41 /var/tmp/if-rtssdb/env.lck
12.0K Nov  7 15:41 /var/tmp/if-rtssdb/env.mem
132.0K Nov  7 15:55 /var/tmp/if-rtssdb/shm_usr1.mem
2688.0K Nov  7 15:41 /var/tmp/if-rtssdb/shm_usr2.mem
2048.0K Nov  7 15:41 /var/tmp/if-rtssdb/trace.mem
727B Nov  7 15:54 /var/tmp/juniper.conf+.gz
155B Nov  7 15:55 /var/tmp/krt_gencfg_filter.txt
0B Nov  7 15:41 /var/tmp/rtssdb/if-rtssdb

```

### request system storage cleanup node-group group-name (QFabric Systems)

```

user@switch> request system storage cleanup node-group NW-NG-0
BBAK0372:

```

-----

List of files to delete:

	Size	Date	Name
	126B	Nov  8 19:07	/var/log/default-log-messages.0.gz
	179B	Nov  7 13:32	/var/log/install.0.gz
	22.9K	Nov  8 19:07	/var/log/messages.0.gz
	26.5K	Nov  8 17:30	/var/log/messages.1.gz
	20.5K	Nov  8 13:15	/var/log/messages.2.gz
	33.2K	Nov  7 17:45	/var/log/messages.3.gz
	35.5K	Nov  7 15:45	/var/log/messages.4.gz
	339B	Nov  8 17:10	/var/log/wtmp.0.gz
	58B	Nov  7 12:40	/var/log/wtmp.1.gz
	124.0K	Nov  8 17:08	/var/tmp/gres-tp/env.dat
	0B	Nov  7 12:39	/var/tmp/gres-tp/lock
	0B	Nov  7 12:59	/var/tmp/if-rtssdb/env.lck
	12.0K	Nov  8 17:09	/var/tmp/if-rtssdb/env.mem
	2688.0K	Nov  8 17:09	/var/tmp/if-rtssdb/shm_usr1.mem
	132.0K	Nov  8 17:09	/var/tmp/if-rtssdb/shm_usr2.mem
	2048.0K	Nov  8 17:09	/var/tmp/if-rtssdb/trace.mem
	1082B	Nov  8 17:09	/var/tmp/juniper.conf+.gz
	155B	Nov  7 17:39	/var/tmp/krt_gencfg_filter.txt
	0B	Nov  8 17:09	/var/tmp/rtssdb/if-rtssdb

EE3093:

-----

List of files to delete:

	Size	Date	Name
	11B	Nov  8 17:33	/var/jail/tmp/alarmd.ts
	119B	Nov  8 19:08	/var/log/default-log-messages.0.gz
	180B	Nov  7 17:41	/var/log/install.0.gz
	178B	Nov  7 13:32	/var/log/install.1.gz
	2739B	Nov  8 19:08	/var/log/messages.0.gz
	29.8K	Nov  8 18:45	/var/log/messages.1.gz
	31.8K	Nov  8 17:15	/var/log/messages.2.gz
	20.6K	Nov  8 16:00	/var/log/messages.3.gz
	15.4K	Nov  8 10:15	/var/log/messages.4.gz

```

15.4K Nov  8 02:15 /var/log/messages.5.gz
25.5K Nov  7 20:45 /var/log/messages.6.gz
48.0K Nov  7 17:45 /var/log/messages.7.gz
32.8K Nov  7 13:45 /var/log/messages.8.gz
684B Nov  8 17:02 /var/log/wtmp.0.gz
58B Nov  7 12:40 /var/log/wtmp.1.gz
124.0K Nov  7 17:34 /var/tmp/gres-tp/env.dat
  0B Nov  7 12:40 /var/tmp/gres-tp/lock
  0B Nov  7 12:59 /var/tmp/if-rtbdb/env.lock
12.0K Nov  7 17:39 /var/tmp/if-rtbdb/env.mem
2688.0K Nov  7 17:39 /var/tmp/if-rtbdb/shm_usr1.mem
132.0K Nov  7 17:40 /var/tmp/if-rtbdb/shm_usr2.mem
2048.0K Nov  7 17:39 /var/tmp/if-rtbdb/trace.mem
155B Nov  7 17:40 /var/tmp/krt_gencfg_filter.txt
  0B Nov  7 17:39 /var/tmp/rtbdb/if-rtbdb

```

#### request system storage cleanup qfabric component device-name (QFabric Systems)

```

user@switch> request system storage cleanup qfabric component A0001/YA0197
Repository type: regular
Repository head: /pbstorage
Creating list of debug artifacts to be removed under:
/pbstorage/rdumps/A0001/YA0197
Removing debug artifacts ... (press control C to abort)
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.0.0.05162011123308.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.1.0.05162011123614.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.2.0.05162011123920.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/livekcore.05132011163930.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/tnetd.core.0.1057.05162011124500.gz ...
done
Removing /pbstorage/rdumps/A0001/YA0197/vmcore.05132011120528.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/vmcore.kz ... done
Creating list of debug artifacts to be removed under: /pbstorage/rlogs/A0001/YA0197
Removing debug artifacts ... (press control C to abort)
Removing /pbstorage/rlogs/A0001/YA0197/kdumpinfo.05132011120528 ... done
Removing /pbstorage/rlogs/A0001/YA0197/kernel.tarball.0.1039.051220111234415.tgz
... done
Removing /pbstorage/rlogs/A0001/YA0197/kernel.tarball.1.1039.05132011175544.tgz
... done
Removing /pbstorage/rlogs/A0001/YA0197/tnetd.tarball.0.1057.05162011175453.tgz
... done

```

#### request system storage cleanup qfabric component device-name repository core (QFabric Systems)

```

user@switch> request system storage cleanup qfabric component EE3093 repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
Creating list of debug artifacts to be removed under: /pbdata/export/rdumps/EE3093
NOTE: core repository under /pbdata/export/rdumps/EE3093 empty

```

#### request system storage cleanup qfabric component all (QFabric Systems)

```

user@switch> request system storage cleanup qfabric component all
Repository scope: shared
Repository head: /pbdata/export
Creating list of debug artifacts to be removed under: /pbdata/export/rdumps
NOTE: core repository under /pbdata/export/rdumps/all empty
Creating list of debug artifacts to be removed under: /pbdata/export/rlogs
List of debug artifacts to clean up ... (press control C to abort)
/pbdata/export/rlogs/73747cd8-0710-11e1-b6a4-00e081c5297e/install-11072011125819.log
/pbdata/export/rlogs/77116f18-0710-11e1-a2a0-00e081c5297e/install-11072011125819.log

```

```
/pbdata/export/rlogs/BBAK0372/install-11072011121538.log  
/pbdata/export/rlogs/BBAK0394/install-11072011121532.log  
/pbdata/export/rlogs/EE3093/install-11072011121536.log  
/pbdata/export/rlogs/WS001/YN5999/install-11072011121644.log  
/pbdata/export/rlogs/WS001/YW3803/install-11072011122429.log  
/pbdata/export/rlogs/cd78871a-0710-11e1-878e-00e081c5297e/install-11072011125932.log  
/pbdata/export/rlogs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011125930.log  
/pbdata/export/rlogs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011133211.log  
/pbdata/export/rlogs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011155302.log  
/pbdata/export/rlogs/d31ab7a6-0710-11e1-ad1b-00e081c5297e/install-11072011125931.log  
/pbdata/export/rlogs/d4d0f254-0710-11e1-90c3-00e081c5297e/install-11072011125932.log
```

## restart

### List of Syntax [Syntax on page 135](#)

[Syntax \(ACX Series Routers\) on page 135](#)  
[Syntax \(EX Series Switches\) on page 135](#)  
[Syntax \(Routing Matrix\) on page 136](#)  
[Syntax \(J Series Routing Platform\) on page 136](#)  
[Syntax \(TX Matrix Routers\) on page 136](#)  
[Syntax \(TX Matrix Plus Routers\) on page 136](#)  
[Syntax \(MX Series Routers\) on page 136](#)  
[Syntax \(J Series Routers\) on page 137](#)  
[Syntax \(QFX Series\) on page 137](#)

### Syntax restart

```

<adaptive-services | ancpd-service | application-identification | audit-process |
  auto-configuration | captive-portal-content-delivery | ce-l2tp-service | chassis-control |
  class-of-service | clksyncd-service | database-replication | datapath-trace-service
  | dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
  ecc-error-logging | ethernet-connectivity-fault-management
  | ethernet-link-fault-management | event-processing | firewall
  | general-authentication-service | gracefully | iccp-service | idp-policy | immediately
  | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2cpd-service
  | l2tp-service | l2tp-universal-edge | lacp | license-service | link-management
  | local-policy-decision-function | mac-validation | mib-process | mobile-ip | mountd-service
  | mpls-traceroute | mspd | multicast-snooping | named-service | nfsd-service |
  packet-triggered-subscribers | peer-selection-service | pgcp-service | pgm |
  pic-services-logging | pki-service | ppp | ppp-service | pppoe |
  protected-system-domain-service | redundancy-interface-process | remote-operations |
  root-system-domain-service | routing <logical-system logical-system-name> | sampling
  | sbc-configuration-process | sdk-service | service-deployment | services | services pgcp
  gateway gateway-name | snmp | soft | static-subscribers | statistics-service |
  subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control |
  vrrp | web-management>
<gracefully | immediately | soft>

```

### Syntax (ACX Series Routers)

```

restart
<adaptive-services | audit-process | auto-configuration | autoinstallation | chassis-control |
  class-of-service | clksyncd-service | database-replication | dhcp-service | diameter-service
  | disk-monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management
  | ethernet-link-fault-management | event-processing | firewall
  | general-authentication-service | gracefully | immediately | interface-control |
  ipsec-key-management | l2-learning | lacp | link-management | mib-process | mobile-ip |
  mountd-service | mpls-traceroute | mspd | named-service | nfsd-service | pgm | pki-service
  | ppp | pppoe | redundancy-interface-process | remote-operations | routing | sampling |
  sdk-service | secure-neighbor-discovery | service-deployment | services | snmp | soft
  | statistics-service | subscriber-management | subscriber-management-helper | tunnel-oamd
  | vrrp>

```

### Syntax (EX Series Switches)

```

restart
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp |
  dhcp-service | diameter-service | dot1x-protocol | ethernet-link-fault-management |
  ethernet-switching | event-processing | firewall | general-authentication-service |
  interface-control | kernel-replication | l2-learning | lacp | license-service | link-management

```

	lldpd-service   mib-process   mounstd-service   multicast-snooping   pgm   redundancy-interface-process   remote-operations   routing   secure-neighbor-discovery   service-deployment   sflow-service   snmp   vrrp   web-management>
<b>Syntax (Routing Matrix)</b>	restart <adaptive-services   audit-process   chassis-control   class-of-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   event-processing   firewall   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2tp-service   lacp   link-management   mib-process   pgm   pic-services-logging   ppp   pppoe   redundancy-interface-process   remote-operations   routing <logical-system <i>logical-system-name</i> >   sampling   service-deployment   snmp> <all   all-lcc   lcc <i>number</i> > <gracefully   immediately   soft>
<b>Syntax (J Series Routing Platform)</b>	restart <adaptive-services   audit-process   chassis-control   class-of-service   dhcp   dialer-services   dls   event-processing   firewall   interface-control   ipsec-key-management   isdn-signaling   l2-learning   l2tp-service   mib-process   network-access-service   pgm   ppp   pppoe   remote-operations   routing <logical-system <i>logical-system-name</i> >   sampling   service-deployment   snmp   usb-control   web-management> <gracefully   immediately   soft>
<b>Syntax (TX Matrix Routers)</b>	restart <adaptive-services   audit-process   chassis-control   class-of-service   dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   event-processing   firewall   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2tp-service   lacp   link-management   mib-process   pgm   pic-services-logging   ppp   pppoe   redundancy-interface-process   remote-operations   routing <logical-system <i>logical-system-name</i> >   sampling   service-deployment   snmp   statistics-service> <all-chassis   all-lcc   lcc <i>number</i>   scc> <gracefully   immediately   soft>
<b>Syntax (TX Matrix Plus Routers)</b>	restart <adaptive-services   audit-process   chassis-control   class-of-service   dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   event-processing   firewall   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2tp-service   lacp   link-management   mib-process   pgm   pic-services-logging   ppp   pppoe   redundancy-interface-process   remote-operations   routing <logical-system <i>logical-system-name</i> >   sampling   service-deployment   snmp   statistics-service> <all-chassis   all-lcc   all-sfc   lcc <i>number</i>   sfc <i>number</i> > <gracefully   immediately   soft>
<b>Syntax (MX Series Routers)</b>	restart <adaptive-services   ancpd-service   application-identification   audit-process   auto-configuration   captive-portal-content-delivery   ce-l2tp-service   chassis-control   class-of-service   clksyncd-service   database-replication   datapath-trace-service   dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   ethernet-connectivity-fault-management   ethernet-link-fault-management   event-processing   firewall   general-authentication-service   gracefully   iccp-service   idp-policy   immediately   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2cpd-service   l2tp-service   l2tp-universal-edge   lacp   license-service   link-management   local-policy-decision-function   mac-validation   mib-process   mobile-ip   mounstd-service   mpls-traceroute   msp   multicast-snooping   named-service   nfsd-service



```

packet-triggered-subscribers |peer-selection-service | pgcp-service | pgm |
pic-services-logging | pki-service | ppp | ppp-service | pppoe |
protected-system-domain-service | redundancy-interface-process | remote-operations
|root-system-domain-service | routing |routing <logical-system logical-system-name> |
sampling | sbc-configuration-process | sdk-service |service-deployment |services | services
pgcp gateway gateway-name |snmp |soft |static-subscribers |statistics-service|
subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control|
vrrp |web-management>
<all-members>
<gracefully | immediately | soft>
<local>
<member member-id>

```

**Syntax (J Series  
Routers)**

```

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp | dhcp-service
| dialer-services | diameter-service | dlsr | event-processing | firewall | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
network-access-service | pgm | ppp | pppoe | remote-operations | routing <logical-system
logical-system-name> | sampling | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>

```

**Syntax (QFX Series)**

```

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
diameter-service | dlsr | ethernet-connectivity | event-processing | fibre-channel | firewall
| general-authentication-service | igmp-host-services | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
named-service | network-access-service | nstrace-process | pgm | ppp | pppoe |
redundancy-interface-process | remote-operations |logical-system-name> | routing |
sampling |secure-neighbor-discovery | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>

```

**Release Information**

Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.  
 Command introduced in Junos OS Release 12.2 for ACX Series routers.  
 Options added:

- **dynamic-flow-capture** in Junos OS Release 7.4.
- **dlsr** in Junos OS Release 7.5.
- **event-processing** in Junos OS Release 7.5.
- **ppp** in Junos OS Release 7.5.
- **l2ald** in Junos OS Release 8.0.
- **link-management** in Release 8.0.
- **pgcp-service** in Junos OS Release 8.4.
- **sbc-configuration-process** in Junos OS Release 9.5.
- **services pgcp gateway** in Junos OS Release 9.6.
- **sfc** and **all-sfc** for the TX Matrix Router in Junos OS Release 9.6.

**Description** Restart a Junos OS process.



**CAUTION:** Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

**Options** **none**—Same as **gracefully**.

**adaptive-services**—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

**all-chassis**—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.

**all-lcc**—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.

**all-members**—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.

**all-sfc**—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).

**ancpd-service**—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

**application-identification**—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.

**audit-process**—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.

**auto-configuration**—(Optional) Restart the Interface Auto-Configuration process.

**autoinstallation**—(EX Series switches only) (Optional) Restart the autoinstallation process.

**captive-portal-content-delivery**—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

**ce-l2tp-service**—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

**chassis-control**—(Optional) Restart the chassis management process.

**class-of-service**—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

**clksyncd-service**—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

**database-replication**—(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.

**datapath-trace-service**—(Optional) Restart the packet path tracing process.

**dhcp**—(J Series routers and EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

**dhcp-service**—(Optional) Restart the Dynamic Host Configuration Protocol process.

**dialer-services**—(J Series routers and EX Series switches only) (Optional) Restart the ISDN dial-out process.

**diameter-service**—(Optional) Restart the diameter process.

**disk-monitoring**—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

**dls**—(J Series routers and QFX Series only) (Optional) Restart the data link switching (DLSw) service.

**dot1x-protocol**—(EX Series switches only) (Optional) Restart the port-based network access control process.

**dynamic-flow-capture**—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

**ecc-error-logging**—(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

**ethernet-connectivity-fault-management**—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

**ethernet-link-fault-management**—(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

**ethernet-switching**—(EX Series switches only) (Optional) Restart the Ethernet switching process.

**event-processing**—(Optional) Restart the event process (eventd).

**fibre-channel**—(QFX Series only) (Optional) Restart the Fibre Channel process.

**firewall**—(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

**general-authentication-service**—(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.

**gracefully**—(Optional) Restart the software process.

**iccp-service**—(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.

**idp-policy**—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.

**immediately**—(Optional) Immediately restart the software process.

**interface-control**—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

**ipsec-key-management**—(Optional) Restart the IPsec key management process.

**isdn-signaling**—(J Series routers and QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.

**kernel-replication**—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

**l2-learning**—(Optional) Restart the Layer 2 address flooding and learning process.

**l2cpd-service**—(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

**l2tp-service**—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

**l2tp-universal-edge**—(MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

**lACP**—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG,

and to enable the transmission and reception processes for the link to function in an orderly manner.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**license-service**—(EX Series switches only) (Optional) Restart the feature license management process.

**link-management**—(TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

**lldpd-service**—(EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

**local**—(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

**local-policy-decision-function**—(Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

**mac-validation**—(Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

**member *member-id***—(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**mib-process**—(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.

**mobile-ip**—(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.

**mountd-service**—(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.

**mpls-traceroute**—(Optional) Restart the MPLS Periodic Traceroute process.

**mspd**—(Optional) Restart the Multiservice process.

**multicast-snooping**—(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

**named-service**—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

**network-access-service**—(J Series routers and QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.

**nfsd-service**—(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

**packet-triggered-subscribers**—(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

**peer-selection-service**—(Optional) Restart the Peer Selection Service process.

**pgcp-service**—(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

**pgm**—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

**pic-services-logging**—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

**pki-service**—(Optional) Restart the PKI Service process.

**ppp**—(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

**ppp-service**—(Optional) Restart the Universal Edge PPP process, which is the encapsulation protocol process for transporting IP traffic across Universal Edge routers.

**pppoe**—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

**protected-system-domain-service**—(Optional) Restart the Protected System Domain (PSD) process.

**redundancy-interface-process**—(Optional) Restart the ASP redundancy process.

**remote-operations**—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

**root-system-domain-service**—(Optional) Restart the Root System Domain (RSD) service.

**routing**—(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.

**routing <logical-system *logical-system-name*>**—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

**sampling**—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

**sbc-configuration-process**—(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).

**scc**—(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).

**sdk-service**—(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

**secure-neighbor-discovery**—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

**service-deployment**—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

**services**—(Optional) Restart a service.

**services pgcp gateway *gateway-name***—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

**sflow-service**—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

**snmp**—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

**soft**—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

**static-subscribers**—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

**statistics-service**—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

**subscriber-management**—(Optional) Restart the Subscriber Management process.

**subscriber-management-helper**—(Optional) Restart the Subscriber Management Helper process.

**tunnel-oamd**—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 PDUs across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

**usb-control**—(J Series routers and MX Series routers only) (Optional) Restart the USB control process.

**vrp**—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

**web-management**—(J Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

**Required Privilege Level**

reset

**Related Documentation**

- *Overview of Junos OS CLI Operational Mode Commands*

**List of Sample Output** [restart interfaces on page 144](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```





## set chassis display message

---

<b>List of Syntax</b>	<a href="#">Syntax on page 146</a> <a href="#">Syntax (TX Matrix Router) on page 146</a> <a href="#">Syntax (TX Matrix Plus Router) on page 146</a>
<b>Syntax</b>	set chassis display message " <i>message</i> " <permanent>
<b>Syntax (TX Matrix Router)</b>	set chassis display message " <i>message</i> " ( <i>lcc number</i>   <i>scc</i> ) <permanent>
<b>Syntax (TX Matrix Plus Router)</b>	set chassis display message " <i>message</i> " ( <i>fpc-slot slot-number</i>   <i>lcc number</i>   <i>sfc number</i> ) <permanent>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option for TX Matrix Plus router introduced in Junos OS Release 9.6.
<b>Description</b>	Display or stop a text message on the craft interface display, which is on the front of the router, or on the LCD panel display on the switch. The craft interface alternates the display of text messages with standard craft interface messages three times, switching between messages every 60 seconds.



**NOTE:** On T Series routers, when this command is executed with the **permanent** option, the display of the text message alternates with that of the standard craft interface message continuously every 60 seconds.

By default, on both the router and the switch, the text message is displayed for 5 minutes. The craft interface display has four 20-character lines. The LCD panel display has two 16-character lines, and text messages appear only on the second line.

**Options**    **"message"**—Message to display. On the craft interface display, if the message is longer than 20 characters, it wraps onto the next line. If a word does not fit on one line, the entire word moves down to the next line. Any portion of the message that does not fit on the display is truncated. An empty pair of quotation marks (" ") deletes the text message from the craft interface display. On the LCD panel display, the message is limited to 16 characters.

**fpc-slot slot-number**—(TX Matrix Plus routers and EX4200 and QFX Series only) On the router or switch, display the text message on the craft interface for a specific Flexible PIC Concentrator (FPC). Replace **slot-number** with a value from 0 through 31. On the switch, display the text message for a specific member of a Virtual Chassis, where **fpc-slot slot-number** corresponds to the member ID. Replace **slot-number** with a value from 0 through 9. On the QFX Series, the **slot-number** is always 0. On a TX Matrix Plus router with 3D SIBs replace **slot-number** with a value from 0 through 63.

**lcc number**—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**permanent**—(Optional) Display a text message on the craft interface display or LCD panel display permanently.

**scc**—(TX Matrix routers only) Display the text message on the craft interface display of the TX Matrix router (switch-card chassis).

**sfc number**—(TX Matrix Plus routers only) Display the text message on the craft interface display of the TX Matrix Plus router (or switch-fabric chassis).

**Required Privilege Level** clear

**Related Documentation**

- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\) on page 33](#)
- [clear chassis display message on page 81](#)
- *show chassis craft-interface*
- *Understanding the Implementation of System Log Messages on the QFabric System*

**List of Sample Output** [set chassis display message \(Creating\) on page 147](#)  
[set chassis display message \(Deleting\) on page 148](#)

**Output Fields** See *show chassis craft-interface* for an explanation of output fields.

## Sample Output

### set chassis display message (Creating)

The following example shows how to set the display message and verify the result:

```
user@host> set chassis display message "NOC contact Dusty (888) 555-1234"
message sent
```

```
user@host> show chassis craft-interface
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
Host OK LED:    On
Host fail LED:  Off
FPCs           0 1 2 3 4 5 6 7
```

```
-----  
Green  ..  *..  *  *.  
Red    .....  
LCD screen:  
+-----+  
|NOC contact Dusty|  
|(888) 555-1234  |  
+-----+
```

### set chassis display message (Deleting)

The following example shows how to delete the display message and verify that the message is removed:

```
user@host> set chassis display message ""  
message sent
```

```
user@host> show chassis craft-interface  
Red alarm:      LED off, relay off  
Yellow alarm:   LED off, relay off  
Host OK LED:    On  
Host fail LED:  Off  
FPCs           0  1  2  3  4  5  6  7  
-----  
Green  ..  *..  *  *.  
Red    .....  
LCD screen:  
+-----+  
|host  
|Up: 0+17:05:47  
|  
|Temperature OK  
+-----+
```

## set date

<b>Syntax</b>	<code>set date (date-time ntp &lt;key authentication-key number&gt; &lt;servers&gt; &lt;source-address source-address&gt;)</code>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p><b>key</b> option introduced in Junos OS Release 12.1R2</p>
<b>Description</b>	Set the date and time.
<b>Options</b>	<p><b>date-time</b>—Date and time. Enter this string inside quotation marks.</p> <p><b>ntp</b>—Use a Network Time Protocol (NTP) server to synchronize the current date and time setting on the router or switch.</p> <p><b>key authentication-key number</b>—(Optional) Specify a key number to authenticate the NTP server used to synchronize the date and time. You must specify the same key number used to authenticate the server configured at the <code>[edit system ntp authentication-key number]</code> hierarchy level.</p> <p><b>servers</b>—(Optional) Specify the IP address of one or more NTP servers.</p> <p><b>source-address source-address</b>—Specify the source address that the router or switch uses to contact the remote NTP server.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Setting the Date and Time</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">set date on page 149</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### set date

```
user@host> set date ntp
21 Apr 17:22:02 ntpdate[3867]: step time server 172.17.27.46 offset 8.759252 sec
```

## show chassis fan

---

<b>List of Syntax</b>	<a href="#">Syntax on page 150</a> <a href="#">Syntax (ACX4000 Series Router) on page 150</a> <a href="#">Syntax (MX Series Router) on page 150</a> <a href="#">Syntax (T Series Routers) on page 150</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Router) on page 150</a> <a href="#">Syntax (QFabric Systems) on page 150</a> <a href="#">Syntax (TX Matrix Router) on page 150</a> <a href="#">Syntax (TX Matrix Plus Router) on page 150</a>
<b>Syntax</b>	show chassis fan
<b>Syntax (ACX4000 Series Router)</b>	show chassis fan
<b>Syntax (MX Series Router)</b>	show chassis fan <all-members> <local> <member <i>member-id</i> >
<b>Syntax (T Series Routers)</b>	show chassis fan
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Router)</b>	show chassis fan
<b>Syntax (QFabric Systems)</b>	show chassis fan <interconnect-device <i>name</i> >
<b>Syntax (TX Matrix Router)</b>	show chassis fan <lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show chassis fan <lcc <i>number</i>   sfc <i>number</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.0 on MX Series 3D Universal Edge Routers, M120 routers, and M320 routers, T320 routers, T640 routers, T1600 routers, TX Matrix Routers, and TX Matrix Plus routers. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 11.4 for EX Series switches. Command introduced in Junos OS Release 12.3 for PTX5000 Packet Transport Routers. Command introduced in Junos OS Release 12.1 for T4000 routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for ACX Series Routers. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
<b>Description</b>	(T Series routers, TX Matrix routers, TX Matrix Plus routers, M120 routers, M320 routers, MX104 routers, MX2010 routers, MX2020 routers, MX Series 3D Universal Edge Routers,

QFX3008-I Interconnect devices, EX Series switches, and PTX Series Packet Transport Routers only) Show information about the fan tray and fans.

**Options** **all-members**—(MX Series routers only) (Optional) Display information about the fan tray and fans for all members of the Virtual Chassis configuration.

**local**—(MX Series routers only) (Optional) Display information about the fan tray and fans for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display information about the fan tray and fans for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* variable with a value 0 or 1.

**interconnect-device *name***—(QFX3000-G QFabric systems only) (Optional) Display information about the fan tray and fans for the specified QFX3008-I Interconnect device.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display information about the fan tray and fans for the specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display information about the fan tray and fans for the specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**scc**—(TX Matrix routers only) (Optional) Display information about the fan tray and fans for the TX Matrix router (switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display information about the fan tray and fans for the TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with 0.

**Required Privilege Level** view

**List of Sample Output** [show chassis fan on page 152](#)  
[show chassis fan \(QFabric Systems\) on page 153](#)  
[show chassis fan \(EX Series Switches\) on page 154](#)  
[show chassis fan \(T320 Router\) on page 154](#)  
[show chassis fan \(T640 Router\) on page 155](#)  
[show chassis fan \(T1600 Router\) on page 155](#)

[show chassis fan \(T4000 Core Router\) on page 156](#)  
[show chassis fan \(TX Matrix Router\) on page 156](#)  
[show chassis fan \(TX Matrix Plus Router\) on page 157](#)  
[show chassis fan \(TX Matrix Plus Router with 3D SIBs\) on page 158](#)  
[show chassis fan \(PTX5000 Packet Transport Router\) on page 160](#)  
[show chassis fan \(MX104 Router\) on page 161](#)  
[show chassis fan \(MX2010 Router\) on page 161](#)  
[show chassis fan \(MX2020 Router\) on page 161](#)  
[show chassis fan \(ACX4000 Router\) on page 162](#)  
[show chassis fan \(QFX5100 Switch\) on page 162](#)

**Output Fields** Table 11 on page 152 lists the output fields for the **show chassis fan** command. Output fields are listed in the approximate order in which they appear.

**Table 11: show chassis fan Output Fields**

Field Name	Field Description
<b>Item</b>	Fan item identifier.
<b>Status</b>	Status of the fan: <ul style="list-style-type: none"> <li>• <b>OK</b>—Fan is running properly and within the normal range.</li> <li>• <b>Check</b>—Fan is in <b>Check</b> state because of some fault or alarm condition.</li> </ul>
<b>RPM</b>	(T Series routers, TX Matrix routers, TX Matrix Plus routers, MX Series 3D Universal Edge Routers, QFX3108 Interconnect devices, and EX Series switches only) Fan speed in revolutions per minute (RPM).
<b>% RPM</b>	(MX2010 routers, MX2020 routers, and PTX Series Packet Transport Routers only) Percentage of the fan speed being used.
<b>Measurement</b>	(T Series routers, TX Matrix routers, TX Matrix Plus routers, MX Series 3D Universal Edge Routers, QFX3108 Interconnect devices, and EX Series switches only) Fan speed status based on different chassis cooling requirements: <ul style="list-style-type: none"> <li>• Spinning at high speed</li> <li>• Spinning at intermediate speed</li> <li>• Spinning at normal speed</li> <li>• Spinning at low speed (except EX Series switches)</li> </ul> (MX2010 routers, MX2020 routers, and PTX Series Packet Transport Routers only) Fan speed in revolutions per minute (RPM) for each fan in the fan tray.

## Sample Output

show chassis fan

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
------	--------	-----	-------------



Top Tray Fan 1	OK	3790	Spinning at normal speed
Top Tray Fan 2	OK	3769	Spinning at normal speed
Top Tray Fan 3	OK	3769	Spinning at normal speed
Top Tray Fan 4	OK	3790	Spinning at normal speed
Top Tray Fan 5	OK	3790	Spinning at normal speed
Top Tray Fan 6	OK	3769	Spinning at normal speed
Top Tray Fan 7	OK	3790	Spinning at normal speed
Top Tray Fan 8	OK	3769	Spinning at normal speed
Top Tray Fan 9	OK	3769	Spinning at normal speed
Top Tray Fan 10	OK	3790	Spinning at normal speed
Top Tray Fan 11	OK	3790	Spinning at normal speed
Top Tray Fan 12	OK	3769	Spinning at normal speed
Bottom Tray Fan 1	OK	2880	Spinning at normal speed
Bottom Tray Fan 2	OK	2912	Spinning at normal speed
Bottom Tray Fan 3	OK	2928	Spinning at normal speed
Bottom Tray Fan 4	OK	2896	Spinning at normal speed
Bottom Tray Fan 5	OK	2896	Spinning at normal speed
Bottom Tray Fan 6	OK	2928	Spinning at normal speed

### show chassis fan (QFabric Systems)

```
user@host> show chassis fan interconnect-device interconnect1
```

Item	Status	RPM	Measurement
TFT 0 Fan 0	OK	2849	Spinning at normal speed
TFT 0 Fan 1	OK	2821	Spinning at normal speed
TFT 0 Fan 2	OK	2735	Spinning at normal speed
TFT 0 Fan 3	OK	2815	Spinning at normal speed
TFT 0 Fan 4	OK	2828	Spinning at normal speed
TFT 0 Fan 5	OK	2863	Spinning at normal speed
BFT 1 Fan 0	OK	2941	Spinning at normal speed
BFT 1 Fan 1	OK	3008	Spinning at normal speed
BFT 1 Fan 2	OK	3073	Spinning at normal speed
BFT 1 Fan 3	OK	2925	Spinning at normal speed
BFT 1 Fan 4	OK	2863	Spinning at normal speed
BFT 1 Fan 5	OK	2933	Spinning at normal speed
SFT 0 Fan 0 Rotor 0	OK	15472	Spinning at normal speed
SFT 0 Fan 0 Rotor 1	OK	14477	Spinning at normal speed
SFT 0 Fan 1 Rotor 0	OK	15561	Spinning at normal speed
SFT 0 Fan 1 Rotor 1	OK	14210	Spinning at normal speed
SFT 0 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 0 Fan 2 Rotor 1	OK	14248	Spinning at normal speed
SFT 0 Fan 3 Rotor 0	OK	16463	Spinning at normal speed
SFT 0 Fan 3 Rotor 1	OK	14099	Spinning at normal speed
SFT 1 Fan 0 Rotor 0	OK	15083	Spinning at normal speed
SFT 1 Fan 0 Rotor 1	OK	13533	Spinning at normal speed
SFT 1 Fan 1 Rotor 0	OK	16071	Spinning at normal speed
SFT 1 Fan 1 Rotor 1	OK	14400	Spinning at normal speed
SFT 1 Fan 2 Rotor 0	OK	15517	Spinning at normal speed
SFT 1 Fan 2 Rotor 1	OK	14210	Spinning at normal speed
SFT 1 Fan 3 Rotor 0	OK	16413	Spinning at normal speed
SFT 1 Fan 3 Rotor 1	OK	14400	Spinning at normal speed
SFT 2 Fan 0 Rotor 0	OK	15297	Spinning at normal speed
SFT 2 Fan 0 Rotor 1	OK	14634	Spinning at normal speed
SFT 2 Fan 1 Rotor 0	OK	15561	Spinning at normal speed
SFT 2 Fan 1 Rotor 1	OK	14285	Spinning at normal speed
SFT 2 Fan 2 Rotor 0	OK	15835	Spinning at normal speed
SFT 2 Fan 2 Rotor 1	OK	14400	Spinning at normal speed
SFT 2 Fan 3 Rotor 0	OK	15789	Spinning at normal speed
SFT 2 Fan 3 Rotor 1	OK	14323	Spinning at normal speed
SFT 3 Fan 0 Rotor 0	OK	16314	Spinning at normal speed

SFT 3 Fan 0 Rotor 1	OK	14876	Spinning at normal speed
SFT 3 Fan 1 Rotor 0	OK	15835	Spinning at normal speed
SFT 3 Fan 1 Rotor 1	OK	14323	Spinning at normal speed
SFT 3 Fan 2 Rotor 0	OK	16265	Spinning at normal speed
SFT 3 Fan 2 Rotor 1	OK	14594	Spinning at normal speed
SFT 3 Fan 3 Rotor 0	OK	16071	Spinning at normal speed
SFT 3 Fan 3 Rotor 1	OK	14323	Spinning at normal speed
SFT 4 Fan 0 Rotor 0	OK	15652	Spinning at normal speed
SFT 4 Fan 0 Rotor 1	OK	14438	Spinning at normal speed
SFT 4 Fan 1 Rotor 0	OK	16167	Spinning at normal speed
SFT 4 Fan 1 Rotor 1	OK	14555	Spinning at normal speed
SFT 4 Fan 2 Rotor 0	OK	16023	Spinning at normal speed
SFT 4 Fan 2 Rotor 1	OK	14361	Spinning at normal speed
SFT 4 Fan 3 Rotor 0	OK	16216	Spinning at normal speed
SFT 4 Fan 3 Rotor 1	OK	14438	Spinning at normal speed
SFT 5 Fan 0 Rotor 0	OK	15297	Spinning at normal speed
SFT 5 Fan 0 Rotor 1	OK	14173	Spinning at normal speed
SFT 5 Fan 1 Rotor 0	OK	15472	Spinning at normal speed
SFT 5 Fan 1 Rotor 1	OK	13846	Spinning at normal speed
SFT 5 Fan 2 Rotor 0	OK	15340	Spinning at normal speed
SFT 5 Fan 2 Rotor 1	OK	13917	Spinning at normal speed
SFT 5 Fan 3 Rotor 0	OK	15835	Spinning at normal speed
SFT 5 Fan 3 Rotor 1	OK	13917	Spinning at normal speed
SFT 6 Fan 0 Rotor 0	OK	15743	Spinning at normal speed
SFT 6 Fan 0 Rotor 1	OK	14594	Spinning at normal speed
SFT 6 Fan 1 Rotor 0	OK	16167	Spinning at normal speed
SFT 6 Fan 1 Rotor 1	OK	14634	Spinning at normal speed
SFT 6 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 6 Fan 2 Rotor 1	OK	14516	Spinning at normal speed
SFT 6 Fan 3 Rotor 0	OK	16666	Spinning at normal speed
SFT 6 Fan 3 Rotor 1	OK	14438	Spinning at normal speed
SFT 7 Fan 0 Rotor 0	OK	15517	Spinning at normal speed
SFT 7 Fan 0 Rotor 1	OK	14438	Spinning at normal speed
SFT 7 Fan 1 Rotor 0	OK	15517	Spinning at normal speed
SFT 7 Fan 1 Rotor 1	OK	14361	Spinning at normal speed
SFT 7 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 7 Fan 2 Rotor 1	OK	14555	Spinning at normal speed
SFT 7 Fan 3 Rotor 0	OK	15697	Spinning at normal speed
SFT 7 Fan 3 Rotor 1	OK	14361	Spinning at normal speed

### show chassis fan (EX Series Switches)

user@host> show chassis fan

Item	Status	RPM	Measurement
Fan 1	OK	3477	Spinning at normal speed
Fan 2	OK	3477	Spinning at normal speed
Fan 3	OK	3479	Spinning at normal speed
Fan 4	OK	3508	Spinning at normal speed
Fan 5	OK	3517	Spinning at normal speed
Fan 6	OK	3531	Spinning at normal speed
Fan 7	OK	3439	Spinning at normal speed
Fan 8	OK	3424	Spinning at normal speed
Fan 9	OK	3413	Spinning at normal speed
Fan 10	OK	3439	Spinning at normal speed
Fan 11	OK	3446	Spinning at normal speed
Fan 12	OK	3432	Spinning at normal speed

### show chassis fan (T320 Router)

user@host> show chassis fan

Item	Status	RPM	Measurement
Top Left Front fan	OK	2850	Spinning at normal speed
Top Left Middle fan	OK	2820	Spinning at normal speed
Top Left Rear fan	OK	2970	Spinning at normal speed
Top Right Front fan	OK	2790	Spinning at normal speed
Top Right Middle fan	OK	2640	Spinning at normal speed
Top Right Rear fan	OK	2790	Spinning at normal speed
Bottom Left Front fan	OK	2520	Spinning at normal speed
Bottom Left Middle fan	OK	2610	Spinning at normal speed
Bottom Left Rear fan	OK	2550	Spinning at normal speed
Bottom Right Front fan	OK	2610	Spinning at normal speed
Bottom Right Middle fan	OK	2880	Spinning at normal speed
Bottom Right Rear fan	OK	2790	Spinning at normal speed
Rear Tray Top fan	OK	2130	Spinning at normal speed
Rear Tray Second fan	OK	2190	Spinning at normal speed
Rear Tray Middle fan	OK	2250	Spinning at normal speed
Rear Tray Fourth fan	OK	2220	Spinning at normal speed
Rear Tray Bottom fan	OK	2280	Spinning at normal speed

#### show chassis fan (T640 Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3390	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3390	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	5220	Spinning at normal speed
Rear Tray Second fan	OK	5220	Spinning at normal speed
Rear Tray Third fan	OK	5220	Spinning at normal speed
Rear Tray Fourth fan	OK	5220	Spinning at normal speed
Rear Tray Fifth fan	OK	5220	Spinning at normal speed
Rear Tray Sixth fan	OK	5220	Spinning at normal speed
Rear Tray Seventh fan	OK	5220	Spinning at normal speed
Rear Tray Bottom fan	OK	5220	Spinning at normal speed

#### show chassis fan (T1600 Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3450	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3390	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3390	Spinning at normal speed

Bottom Right Middle fan	OK	3420	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	5190	Spinning at normal speed
Rear Tray Second fan	OK	5190	Spinning at normal speed
Rear Tray Third fan	OK	5190	Spinning at normal speed
Rear Tray Fourth fan	OK	5190	Spinning at normal speed
Rear Tray Fifth fan	OK	5190	Spinning at normal speed
Rear Tray Sixth fan	OK	5190	Spinning at normal speed
Rear Tray Seventh fan	OK	5190	Spinning at normal speed
Rear Tray Bottom fan	OK	5190	Spinning at normal speed

**show chassis fan (T4000 Core Router)**

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	5190	Spinning at high speed
Top Left Middle fan	OK	5220	Spinning at high speed
Top Left Rear fan	OK	5190	Spinning at high speed
Top Right Front fan	OK	5160	Spinning at high speed
Top Right Middle fan	OK	5190	Spinning at high speed
Top Right Rear fan	OK	5160	Spinning at high speed
Bottom Left Front fan	OK	6030	Spinning at high speed
Bottom Left Middle fan	OK	6090	Spinning at high speed
Bottom Left Rear fan	OK	6090	Spinning at high speed
Bottom Right Front fan	OK	6030	Spinning at high speed
Bottom Right Middle fan	OK	6060	Spinning at high speed
Bottom Right Rear fan	OK	6060	Spinning at high speed
Rear Tray Top fan	OK	10000	Spinning at high speed
Rear Tray Second fan	OK	10000	Spinning at high speed
Rear Tray Third fan	OK	10000	Spinning at high speed
Rear Tray Fourth fan	OK	10000	Spinning at high speed
Rear Tray Fifth fan	OK	10000	Spinning at high speed
Rear Tray Sixth fan	OK	10000	Spinning at high speed
Rear Tray Seventh fan	OK	10000	Spinning at high speed
Rear Tray Bottom fan	OK	10000	Spinning at high speed

**show chassis fan (TX Matrix Router)**

```
user@host> show chassis fan
scc-re0:
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3390	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3390	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3450	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3420	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray Top fan	OK	3420	Spinning at normal speed
Rear Tray Second fan	OK	5190	Spinning at normal speed
Rear Tray Third fan	OK	5190	Spinning at normal speed
Rear Tray Fourth fan	OK	5190	Spinning at normal speed
Rear Tray Fifth fan	OK	3420	Spinning at normal speed
Rear Tray Sixth fan	OK	3420	Spinning at normal speed

```

Rear Tray Seventh fan    OK      3420    Spinning at normal speed
Rear Tray Bottom fan     OK      3420    Spinning at normal speed

```

```
lcc2-re0:
```

```

-----
Item                Status  RPM    Measurement
Top Left Front fan   OK      3420    Spinning at normal speed
Top Left Middle fan  OK      3420    Spinning at normal speed
Top Left Rear fan    OK      3450    Spinning at normal speed
Top Right Front fan   OK      3420    Spinning at normal speed
Top Right Middle fan  OK      3450    Spinning at normal speed
Top Right Rear fan    OK      3360    Spinning at normal speed
Bottom Left Front fan OK      3420    Spinning at normal speed
Bottom Left Middle fan OK     3480    Spinning at normal speed
Bottom Left Rear fan  OK      3420    Spinning at normal speed
Bottom Right Front fan OK     3420    Spinning at normal speed
Bottom Right Middle fan OK     3390    Spinning at normal speed
Bottom Right Rear fan OK      3420    Spinning at normal speed
Rear Tray Top fan     OK      3420    Spinning at normal speed
Rear Tray Second fan  OK      3420    Spinning at normal speed
Rear Tray Third fan   OK      3420    Spinning at normal speed
Rear Tray Fourth fan  OK      3420    Spinning at normal speed
Rear Tray Fifth fan   OK      3420    Spinning at normal speed
Rear Tray Sixth fan   OK      3420    Spinning at normal speed
Rear Tray Seventh fan OK      3420    Spinning at normal speed
Rear Tray Bottom fan  OK      3420    Spinning at normal speed

```

#### show chassis fan (TX Matrix Plus Router)

```

user@host> show chassis fan
sfc0-re0:

```

```

-----
Item                Status  RPM    Measurement
Fan Tray 0 Fan 1     OK      4350    Spinning at normal speed
Fan Tray 0 Fan 2     OK      4380    Spinning at normal speed
Fan Tray 0 Fan 3     OK      4410    Spinning at normal speed
Fan Tray 0 Fan 4     OK      4380    Spinning at normal speed
Fan Tray 0 Fan 5     OK      4350    Spinning at normal speed
Fan Tray 0 Fan 6     OK      4380    Spinning at normal speed
Fan Tray 1 Fan 1     OK      4410    Spinning at normal speed
Fan Tray 1 Fan 2     OK      4380    Spinning at normal speed
Fan Tray 1 Fan 3     OK      4410    Spinning at normal speed
Fan Tray 1 Fan 4     OK      4380    Spinning at normal speed
Fan Tray 1 Fan 5     OK      4410    Spinning at normal speed
Fan Tray 1 Fan 6     OK      4410    Spinning at normal speed
Fan Tray 2 Fan 1     OK      4380    Spinning at normal speed
Fan Tray 2 Fan 2     OK      4380    Spinning at normal speed
Fan Tray 2 Fan 3     OK      4380    Spinning at normal speed
Fan Tray 2 Fan 4     OK      4410    Spinning at normal speed
Fan Tray 2 Fan 5     OK      4380    Spinning at normal speed
Fan Tray 2 Fan 6     OK      4410    Spinning at normal speed
Fan Tray 2 Fan 7     OK      4410    Spinning at normal speed
Fan Tray 2 Fan 8     OK      4380    Spinning at normal speed
Fan Tray 2 Fan 9     OK      4380    Spinning at normal speed
Fan Tray 3 Fan 1     OK      4350    Spinning at normal speed
Fan Tray 3 Fan 2     OK      4380    Spinning at normal speed
Fan Tray 3 Fan 3     OK      4410    Spinning at normal speed
Fan Tray 3 Fan 4     OK      4440    Spinning at normal speed
Fan Tray 3 Fan 5     OK      4380    Spinning at normal speed
Fan Tray 3 Fan 6     OK      4410    Spinning at normal speed
Fan Tray 3 Fan 7     OK      4410    Spinning at normal speed

```

Fan Tray 3 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 3 Fan 9	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 1	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 2	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 4 Fan 4	OK	4380	Spinning at normal speed
Fan Tray 4 Fan 5	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 8	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 9	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 1	OK	4350	Spinning at normal speed
Fan Tray 5 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 4	OK	4350	Spinning at normal speed
Fan Tray 5 Fan 5	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 9	OK	4410	Spinning at normal speed

lcc0-re0:

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3450	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3420	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	7050	Spinning at normal speed
Rear Tray Second fan	OK	7050	Spinning at normal speed
Rear Tray Third fan	OK	7050	Spinning at normal speed
Rear Tray Fourth fan	OK	7050	Spinning at normal speed
Rear Tray Fifth fan	OK	7050	Spinning at normal speed
Rear Tray Sixth fan	OK	7050	Spinning at normal speed
Rear Tray Seventh fan	OK	7050	Spinning at normal speed
Rear Tray Bottom fan	OK	7050	Spinning at normal speed

#### show chassis fan (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fan
sfc0-re0:
```

Item	Status	RPM	Measurement
Fan Tray 0 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 2	OK	4860	Spinning at normal speed
Fan Tray 0 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 4	OK	4800	Spinning at normal speed
Fan Tray 0 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 6	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 1	OK	4800	Spinning at normal speed
Fan Tray 1 Fan 2	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 3	OK	4800	Spinning at normal speed
Fan Tray 1 Fan 4	OK	4770	Spinning at normal speed

Fan Tray 1 Fan 5	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 6	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 1	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 2	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 6	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 7	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 8	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 9	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 1	OK	4860	Spinning at normal speed
Fan Tray 3 Fan 2	OK	4860	Spinning at normal speed
Fan Tray 3 Fan 3	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 6	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 7	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 8	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 9	OK	4800	Spinning at normal speed
Fan Tray 4 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 2	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 6	OK	4860	Spinning at normal speed
Fan Tray 4 Fan 7	OK	4800	Spinning at normal speed
Fan Tray 4 Fan 8	OK	4860	Spinning at normal speed
Fan Tray 4 Fan 9	OK	4770	Spinning at normal speed
Fan Tray 5 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 2	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 4	OK	4800	Spinning at normal speed
Fan Tray 5 Fan 5	OK	4800	Spinning at normal speed
Fan Tray 5 Fan 6	OK	4800	Spinning at normal speed
Fan Tray 5 Fan 7	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 8	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 9	Check	2010	

1cc0-re0:

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3390	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3390	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray fan 1 (Top)	OK	7740	Spinning at normal speed
Rear Tray fan 2	OK	7740	Spinning at normal speed
Rear Tray fan 3	OK	7740	Spinning at normal speed
Rear Tray fan 4	OK	7740	Spinning at normal speed
Rear Tray fan 5	OK	7740	Spinning at normal speed
Rear Tray fan 6	OK	7740	Spinning at normal speed
Rear Tray fan 7	OK	7740	Spinning at normal speed

Rear Tray fan 8	OK	7740	Spinning at normal speed
Rear Tray fan 9	OK	7740	Spinning at normal speed
Rear Tray fan 10	OK	7740	Spinning at normal speed
Rear Tray fan 11	OK	7740	Spinning at normal speed
Rear Tray fan 12	OK	7740	Spinning at normal speed
Rear Tray fan 13	OK	7740	Spinning at normal speed
Rear Tray fan 14	OK	7740	Spinning at normal speed
Rear Tray fan 15	OK	7740	Spinning at normal speed
Rear Tray fan 16 (Bottom)	OK	7740	Spinning at normal speed

```
1cc2-re0:
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3390	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray fan 1 (Top)	OK	7740	Spinning at normal speed
Rear Tray fan 2	OK	7740	Spinning at normal speed
Rear Tray fan 3	OK	7740	Spinning at normal speed
Rear Tray fan 4	OK	7740	Spinning at normal speed
Rear Tray fan 5	OK	7740	Spinning at normal speed
Rear Tray fan 6	OK	7740	Spinning at normal speed
Rear Tray fan 7	OK	7740	Spinning at normal speed
Rear Tray fan 8	OK	7740	Spinning at normal speed
Rear Tray fan 9	OK	7740	Spinning at normal speed
Rear Tray fan 10	OK	7740	Spinning at normal speed
Rear Tray fan 11	OK	7740	Spinning at normal speed
Rear Tray fan 12	OK	7740	Spinning at normal speed
Rear Tray fan 13	OK	7740	Spinning at normal speed
Rear Tray fan 14	OK	7740	Spinning at normal speed
Rear Tray fan 15	OK	7740	Spinning at normal speed
Rear Tray fan 16 (Bottom)	OK	7740	Spinning at normal speed

### show chassis fan (PTX5000 Packet Transport Router)

```
user@host> show chassis fan
user@host> show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	29%	2700 RPM
Fan Tray 0 Fan 2	OK	29%	2700 RPM
Fan Tray 0 Fan 3	OK	29%	2742 RPM
Fan Tray 0 Fan 4	OK	29%	2700 RPM
Fan Tray 0 Fan 5	OK	30%	2828 RPM
Fan Tray 0 Fan 6	OK	30%	2828 RPM
Fan Tray 0 Fan 7	OK	29%	2700 RPM
Fan Tray 0 Fan 8	OK	30%	2785 RPM
Fan Tray 0 Fan 9	OK	30%	2828 RPM
Fan Tray 0 Fan 10	OK	30%	2828 RPM
Fan Tray 0 Fan 11	OK	30%	2785 RPM
Fan Tray 0 Fan 12	OK	30%	2828 RPM
Fan Tray 0 Fan 13	OK	31%	2871 RPM
Fan Tray 0 Fan 14	OK	30%	2828 RPM



Fan Tray 1 Fan 1	OK	42%	3033 RPM
Fan Tray 1 Fan 2	OK	42%	3066 RPM
Fan Tray 1 Fan 3	OK	43%	3099 RPM
Fan Tray 1 Fan 4	OK	43%	3166 RPM
Fan Tray 1 Fan 5	OK	45%	3266 RPM
Fan Tray 1 Fan 6	OK	43%	3133 RPM
Fan Tray 2 Fan 1	OK	29%	2099 RPM
Fan Tray 2 Fan 2	OK	30%	2199 RPM
Fan Tray 2 Fan 3	OK	30%	2166 RPM
Fan Tray 2 Fan 4	OK	33%	2399 RPM
Fan Tray 2 Fan 5	OK	29%	2133 RPM
Fan Tray 2 Fan 6	OK	32%	2366 RPM

### show chassis fan (MX104 Router)

```
user@host > show chassis fan
```

Item	Status	RPM	Measurement
Fan 1	OK	5640	Spinning at normal speed
Fan 2	OK	5640	Spinning at normal speed
Fan 3	OK	5760	Spinning at normal speed
Fan 4	OK	5640	Spinning at normal speed
Fan 5	OK	5640	Spinning at normal speed

### show chassis fan (MX2010 Router)

```
user@host > show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	37%	3360 RPM
Fan Tray 0 Fan 2	OK	38%	3480 RPM
Fan Tray 0 Fan 3	OK	37%	3360 RPM
Fan Tray 0 Fan 4	OK	37%	3360 RPM
Fan Tray 0 Fan 5	OK	38%	3480 RPM
Fan Tray 0 Fan 6	OK	37%	3360 RPM
Fan Tray 1 Fan 1	OK	38%	3480 RPM
Fan Tray 1 Fan 2	OK	40%	3600 RPM
Fan Tray 1 Fan 3	OK	38%	3480 RPM
Fan Tray 1 Fan 4	OK	38%	3480 RPM
Fan Tray 1 Fan 5	OK	38%	3480 RPM
Fan Tray 1 Fan 6	OK	38%	3480 RPM
Fan Tray 2 Fan 1	OK	38%	3480 RPM
Fan Tray 2 Fan 2	OK	41%	3720 RPM
Fan Tray 2 Fan 3	OK	38%	3480 RPM
Fan Tray 2 Fan 4	OK	38%	3480 RPM
Fan Tray 2 Fan 5	OK	38%	3480 RPM
Fan Tray 2 Fan 6	OK	38%	3480 RPM
Fan Tray 3 Fan 1	OK	38%	3480 RPM
Fan Tray 3 Fan 2	OK	40%	3600 RPM
Fan Tray 3 Fan 3	OK	40%	3600 RPM
Fan Tray 3 Fan 4	OK	40%	3600 RPM
Fan Tray 3 Fan 5	OK	40%	3600 RPM
Fan Tray 3 Fan 6	OK	38%	3480 RPM

### show chassis fan (MX2020 Router)

```
user@host > show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	37%	3360 RPM
Fan Tray 0 Fan 2	OK	37%	3360 RPM
Fan Tray 0 Fan 3	OK	36%	3240 RPM
Fan Tray 0 Fan 4	OK	37%	3360 RPM
Fan Tray 0 Fan 5	OK	37%	3360 RPM
Fan Tray 0 Fan 6	OK	37%	3360 RPM

Fan Tray 1 Fan 1	OK	37%	3360 RPM
Fan Tray 1 Fan 2	OK	37%	3360 RPM
Fan Tray 1 Fan 3	OK	37%	3360 RPM
Fan Tray 1 Fan 4	OK	37%	3360 RPM
Fan Tray 1 Fan 5	OK	37%	3360 RPM
Fan Tray 1 Fan 6	OK	36%	3240 RPM
Fan Tray 2 Fan 1	OK	37%	3360 RPM
Fan Tray 2 Fan 2	OK	37%	3360 RPM
Fan Tray 2 Fan 3	OK	37%	3360 RPM
Fan Tray 2 Fan 4	OK	37%	3360 RPM
Fan Tray 2 Fan 5	OK	37%	3360 RPM
Fan Tray 2 Fan 6	OK	38%	3480 RPM
Fan Tray 3 Fan 1	OK	38%	3480 RPM
Fan Tray 3 Fan 2	OK	38%	3480 RPM
Fan Tray 3 Fan 3	OK	38%	3480 RPM
Fan Tray 3 Fan 4	OK	37%	3360 RPM
Fan Tray 3 Fan 5	OK	37%	3360 RPM
Fan Tray 3 Fan 6	OK	37%	3360 RPM

#### show chassis fan (ACX4000 Router)

```
user@host > show chassis fan
```

Item	Status	RPM	Measurement
Fan 1	OK	4140	Spinning at normal speed
Fan 2	OK	4200	Spinning at normal speed

#### show chassis fan (QFX5100 Switch)

```
user@switch > show chassis fan
```

Item	Status	RPM	Measurement
FPC 0 Tray 0 Fan 0	OK	6428	Spinning at normal speed
FPC 0 Tray 0 Fan 1	OK	5515	Spinning at normal speed
FPC 0 Tray 1 Fan 0	OK	6360	Spinning at normal speed
FPC 0 Tray 1 Fan 1	OK	5532	Spinning at normal speed

## show chassis firmware

<b>List of Syntax</b>	<a href="#">Syntax on page 163</a> <a href="#">Syntax (TX Matrix Routers) on page 163</a> <a href="#">Syntax (TX Matrix Plus Routers) on page 163</a> <a href="#">Syntax (MX Series Routers) on page 163</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 163</a> <a href="#">Syntax (QFX Series) on page 163</a> <a href="#">Syntax (ACX Series Universal Access Routers) on page 163</a> <a href="#">Syntax (EX Series Switches) on page 163</a>
<b>Syntax</b>	show chassis firmware
<b>Syntax (TX Matrix Routers)</b>	show chassis firmware <fcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Routers)</b>	show chassis firmware <fcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Routers)</b>	show chassis firmware <all-members> <local> <member <i>member-id</i> >
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)</b>	show chassis firmware
<b>Syntax (QFX Series)</b>	show chassis firmware interconnect-device <i>name</i> node-device <i>name</i>
<b>Syntax (ACX Series Universal Access Routers)</b>	show chassis firmware
<b>Syntax (EX Series Switches)</b>	show chassis firmware <detail>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.4 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced for EX8200 switches in Junos OS Release 10.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for ACX4000 Universal Access Routers.</p> <p>Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.</p>

**Description** On routers and switches, display the version levels of the firmware running on the System Control Board (SCB), Switching and Forwarding Module (SFM), System and Switch Board (SSB), Forwarding Engine Board (FEB), Flexible PIC Concentrators (FPCs), and Routing Engines. On a TX Matrix Plus router, display the version levels of the firmware running on the FPCs and the Switch Processor Mezzanine Board (SPMBs).

On EX2200, EX3200, and EX4200 switches, and the QFX Series, display the version levels of the firmware running on the switch. On an EX8208 switch, display the version levels of the firmware running on the Switch Fabric and Routing Engine (SRE) modules and on the line cards (shown as FPCs). On an EX8216 switch, display the version levels of the firmware running on the Routing Engine (RE) modules and on the line cards (shown as FPCs).

**Options** **none**—Display the version levels of the firmware running. For an EX4200 switch that is a member of a Virtual Chassis, display version levels for all members. For a TX Matrix router, display version levels for the firmware on the TX Matrix router and on all the T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, display version levels for the firmware on the TX Matrix Plus router and on all the routers connected to the TX Matrix Plus router.

**all-members**—(MX Series routers only) (Optional) Display the version levels of the firmware running for all members of the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems) (Optional) Display the version levels of the firmware running on the Interconnect device.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display version levels for the firmware on a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the version levels for the firmware on a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display the version levels of the firmware running for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display the version levels of the firmware running for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**node-device**—(QFabric systems only) (Optional) Display the version levels of the firmware running on the Node device.

**scc**—(TX Matrix router only) (Optional) Display version levels for the firmware on the TX Matrix router (switch-card chassis).

**sfc number**—(TX Matrix Plus router only) (Optional) Display version levels for the firmware on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with **0**.

**detail**—(EX3200, EX3300, EX4200, and EX4500 standalone and Virtual Chassis member switches only) (Optional) Display version levels of the firmware running on the switch for its programmable hardware components.

**Required Privilege Level** view

**Related Documentation** • *Upgrading the HSM Firmware*

**List of Sample Output**

- [show chassis firmware \(M10 Router\) on page 166](#)
- [show chassis firmware \(M20 Router\) on page 166](#)
- [show chassis firmware \(M40 Router\) on page 167](#)
- [show chassis firmware \(M120 Router\) on page 167](#)
- [show chassis firmware \(M160 Router\) on page 167](#)
- [show chassis firmware \(MX104 Router\) on page 167](#)
- [show chassis firmware \(MX240 Router\) on page 167](#)
- [show chassis firmware \(MX480 Router\) on page 168](#)
- [show chassis firmware \(MX960 Router\) on page 168](#)
- [show chassis firmware \(MX2010 Router\) on page 168](#)
- [show chassis firmware \(MX2020 Router\) on page 168](#)
- [show chassis firmware \(MX240, MX480, MX960 Router with Application Services Modular Line Card\) on page 169](#)
- [show chassis firmware \(EX4200 Switch\) on page 169](#)
- [show chassis firmware \(EX8200 Switch\) on page 169](#)
- [show chassis firmware \(EX9200 Switch\) on page 170](#)
- [show chassis firmware lcc \(TX Matrix Router\) on page 170](#)
- [show chassis firmware scc \(TX Matrix Router\) on page 170](#)
- [show chassis firmware \(TX Matrix Plus Router\) on page 170](#)
- [show chassis firmware lcc \(TX Matrix Plus Router\) on page 172](#)
- [show chassis firmware sfc \(TX Matrix Plus Router\) on page 172](#)
- [show chassis firmware \(QFX Series\) on page 172](#)
- [show chassis firmware interconnect-device \(QFabric System\) on page 173](#)
- [show chassis firmware \(ACX2000 Universal Access Router\) on page 173](#)
- [show chassis firmware detail \(EX3300 Switch\) on page 173](#)
- [show chassis firmware \(MX Routers with Media Services Blade \[MSB\]\) on page 173](#)

**Output Fields** [Table 12 on page 166](#) lists the output fields for the **show chassis firmware** command. Output fields are listed in the approximate order in which they appear.

Table 12: show chassis firmware Output Fields

Field Name	Field Description
<b>Part</b>	(MX Series, MX2010, and MX2020 routers) Chassis part name.
<b>Type</b>	(MX Series, MX2010, and MX2020 routers) Type of firmware: On routers: <b>ROM</b> or <b>O/S</b> . On switches: <b>uboot</b> or <b>loader</b> .
<b>Version</b>	(MX Series, MX2010, and MX2020 routers) Version of firmware running on the chassis part.
<b>FPC</b>	( <i>detail</i> option only) Number of FPC. For a standalone switch, the value is 0. For a Virtual Chassis configuration, value in the range of 0-9; refers to the member ID assigned to the switch.
<b>AFEB</b>	(MX104 routers) Version of the compact Forwarding Engine Board.
<b>Boot</b>	( <i>detail</i> option only) Version of the SYSPLD.
<b>PoE</b>	( <i>detail</i> option only) Version of the PoE firmware.
<b>PFE-&lt;number&gt;</b>	( <i>detail</i> option only) Version of the PFE used in the switch.
<b>PHY-</b>	( <i>detail</i> option only) Version of the physical layer device (PHY) used in the switch.
<b>microcode</b>	( <i>detail</i> option only) Microcode of the physical layer devices (PHY) used in the switch.
<b>uboot</b>	( <i>detail</i> option only) Version of the u-boot used in the switch.
<b>loader</b>	( <i>detail</i> option only) Version of the loader used in the switch.

## Sample Output

### show chassis firmware (M10 Router)

```

user@host> show chassis firmware
Part          Type      Version
Forwarding engine board  ROM      Juniper ROM Monitor Version 4.1b2
                                O/S      Version 4.1I1 by tlim on 2000-04-24 11:27

```

### show chassis firmware (M20 Router)

```

user@host> show chassis firmware
Part          Type      Version
System switch board  ROM      Juniper ROM Monitor Version 3.4b26
                                O/S      Version 3.4I16 by smackie on 2000-02-29 2
FPC 1          ROM      Juniper ROM Monitor Version 3.0b1
                                O/S      Version 3.4I4 by smackie on 2000-02-25 21
FPC 2          ROM      Juniper ROM Monitor Version 3.0b1
                                O/S      Version 3.4I4 by smackie on 2000-02-25 21

```

**show chassis firmware (M40 Router)**

```

user@host> show chassis firmware
Part                Type      Version
System control board ROM       Juniper ROM Monitor Version 2.0i126Copyri
                  O/S       Version 2.0i1 by root on Thu Jul 23 00:51
FPC 5               ROM       Juniper ROM Monitor Version 2.0i49Copyrig
                  O/S       Version 2.0i1 by root on Thu Jul 23 00:59

```

**show chassis firmware (M120 Router)**

```

user@host> show chassis firmware
FPC 2               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:2
FPC 3               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:2
FPC 4               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:2
FEB 3               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:1
FEB 4               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:1

```

**show chassis firmware (M160 Router)**

```

user@host> show chassis firmware
Part                Type      Version
SFM 0               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:50
SFM 1               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:50
FPC 0               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:56
FPC 1               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:56
FPC 2               ROM       Juniper ROM Monitor Version 4.0b3
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:56

```

**show chassis firmware (MX104 Router)**

```

user@host > show chassis firmware
Part                Type      Version
FPC 0               ROM       Juniper ROM Monitor Version 13.1b24
                  O/S       Version 13.2-20130514.1 by builder on 2013-
FPC 1               ROM       Juniper ROM Monitor Version 13.1b24
                  O/S       Version 13.2-20130514.1 by builder on 2013-
FPC 2               ROM       Juniper ROM Monitor Version 13.1b24
                  O/S       Version 13.2-20130514.1 by builder on 2013-
AFEB                ROM       Juniper ROM Monitor Version 13.1b24
                  O/S       Version 13.2-20130514.1 by builder on 2013-

```

**show chassis firmware (MX240 Router)**

```

user@host> show chassis firmware
Part                Type      Version
FPC 1               ROM       Juniper ROM Monitor Version 8.3b1
                  O/S       Version 9.0-20080103.0 by builder on 2008-0
FPC 2               ROM       Juniper ROM Monitor Version 8.3b1
                  O/S       Version 9.0-20080103.0 by builder on 2008-0

```

### show chassis firmware (MX480 Router)

```
user@host> show chassis firmware
Part      Type      Version
FPC 1     ROM       Juniper ROM Monitor Version 8.3b1
           O/S      Version 9.0-20070916.3 by builder on 2007-0
```

### show chassis firmware (MX960 Router)

```
user@host> show chassis firmware
Part      Type      Version
FPC 4     ROM       Juniper ROM Monitor Version 8.0b8
           O/S      Version 8.2I59 by artem on 2006-10-31 19:22
FPC 7     ROM       Juniper ROM Monitor Version 8.2b1
           O/S      Version 8.2-20061026.1 by builder on 2006-1
```

### show chassis firmware (MX2010 Router)

```
user@host> show chassis firmware
Part      Type      Version
FPC 0     ROM       Juniper ROM Monitor Version 12.3b1
           O/S      Version 12.3-20121220.0 by builder on 2012-
FPC 1     ROM       Juniper ROM Monitor Version 10.1b3
           O/S      Version 12.3-20121220.0 by builder on 2012-
FPC 2     ROM       Juniper ROM Monitor Version 10.1b3
           O/S      Version 12.3-20121220.0 by builder on 2012-
FPC 3     ROM       Juniper ROM Monitor Version 10.1b3
           O/S      Version 12.3-20121220.0 by builder on 2012-
FPC 4     ROM       Juniper ROM Monitor Version 10.0b39
           O/S      Version 12.3-20121220.0 by builder on 2012-
FPC 5     ROM       Juniper ROM Monitor Version 10.0b39
           O/S      Version 12.3-20121220.0 by builder on 2012-
FPC 6     ROM       Juniper ROM Monitor Version 10.4b1
           O/S      Version 12.3-20121220.0 by builder on 2012-
FPC 7     ROM       Juniper ROM Monitor Version 10.1b3
           O/S      Version 12.3-20121220.0 by builder on 2012-
FPC 8     ROM       Juniper ROM Monitor Version 10.4b1
           O/S      Version 12.3-20121220.0 by builder on 2012-
FPC 9     ROM       Juniper ROM Monitor Version 10.4b1
           O/S      Version 12.3-20121220.0 by builder on 2012-
SPMB 0    ROM       Juniper ROM Monitor Version 12.1b1
           O/S      Version 12.3-20121220.0 by builder on 2012-
SPMB 1    ROM       Juniper ROM Monitor Version 12.1b1
           O/S      Version 12.3-20121220.0 by builder on 2012-
```

### show chassis firmware (MX2020 Router)

```
user@host> show chassis firmware
Part      Type      Version
FPC 0     ROM       Juniper ROM Monitor Version 10.0b39
           O/S      Version 12.3-20130415.0 by builder on 2013-
FPC 1     ROM       Juniper ROM Monitor Version 10.0b39
           O/S      Version 12.3-20130415.0 by builder on 2013-
FPC 2     ROM       Juniper ROM Monitor Version 10.0b39
           O/S      Version 12.3-20130415.0 by builder on 2013-
FPC 3     ROM       Juniper ROM Monitor Version 10.0b39
           O/S      Version 12.3-20130415.0 by builder on 2013-
FPC 4     ROM       Juniper ROM Monitor Version 10.0b39
           O/S      Version 12.3-20130415.0 by builder on 2013-
FPC 5     ROM       Juniper ROM Monitor Version 10.0b39
           O/S      Version 12.3-20130415.0 by builder on 2013-
```



FPC 6	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 7	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 8	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 9	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 10	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 11	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 12	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 13	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 14	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 15	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 16	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 17	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 18	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 19	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
SPMB 0	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.3-20130415.0 by builder on 2013-
SPMB 1	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.3-20130415.0 by builder on 2013-

#### show chassis firmware (MX240, MX480, MX960 Router with Application Services Modular Line Card)

```
user@host> show chassis firmware
```

Part	Type	Version
FPC 1	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.2I21 by manish on 2012-06-19 17:

#### show chassis firmware (EX4200 Switch)

```
user@switch> show chassis firmware
```

Part	Type	Version
FPC 0	uboot	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42)
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.1
FPC 1	uboot	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42)
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.1
FPC 2	uboot	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42)
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.1

#### show chassis firmware (EX8200 Switch)

```
user@switch> show chassis firmware
```

Part	Type	Version
FPC 0	U-Boot	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 3	U-Boot	U-Boot 1.1.6 (Dec 4 2009 - 13:17:34) 3.1.0
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.2

FPC 5	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 7	U-Boot loader	U-Boot 1.1.6 (Feb 6 2009 - 05:31:46) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
Routing Engine 0	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
Routing Engine 1	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2

**show chassis firmware (EX9200 Switch)**

```
user@switch> show chassis firmware
Part      Type      Version
FPC 2     ROM       Juniper ROM Monitor Version 11.4b2
           O/S      Version 14.1I20140312_0741_bavig by bavig o
FPC 3     ROM       Juniper ROM Monitor Version 10.4b1
           O/S      Version 14.1I20140312_0741_bavig by bavig o
```

**show chassis firmware lcc (TX Matrix Router)**

```
user@host> show chassis firmware lcc 0
lcc0-re0:
-----
Part      Type      Version
FPC 1     ROM       Juniper ROM Monitor Version 6.4b18
           O/S      Version 7.0-20040804.0 by builder on 2004-0
FPC 2     ROM       Juniper ROM Monitor Version 6.4b20
           O/S      Version 7.0-20040804.0 by builder on 2004-0
SPMB 0    ROM       Juniper ROM Monitor Version 6.4b18
           O/S      Version 7.0-20040804.0 by builder on 2004-0
```

**show chassis firmware scc (TX Matrix Router)**

```
user@host> show chassis firmware scc
scc-re0:
-----
Part      Type      Version
SPMB 0    ROM       Juniper ROM Monitor Version 6.4b18
           O/S      Version 7.0-20040804.0 by builder on 2004-0
```

**show chassis firmware (TX Matrix Plus Router)**

```
user@host> show chassis firmware
sfc0-re0:
-----
Part      Type      Version
Global FPC 4
Global FPC 6
Global FPC 7
Global FPC 12
Global FPC 14
Global FPC 15
Global FPC 20
Global FPC 21
Global FPC 22
Global FPC 23
Global FPC 24
Global FPC 25
Global FPC 26
Global FPC 28
```

```

Global FPC 29
Global FPC 31
SPMB 0          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0

```

#### lcc0-re1:

```

-----
Part          Type      Version
FPC 4          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 6          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 7          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 0          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0

```

#### lcc1-re1:

```

-----
Part          Type      Version
FPC 4          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 6          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 7          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 0          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0

```

#### lcc2-re1:

```

-----
Part          Type      Version
FPC 4          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 5          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 6          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 7          ROM      Juniper ROM Monitor Version 7.5b4
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 0          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0

```

#### lcc3-re1:

```

-----
Part          Type      Version
FPC 0          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 1          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 2          ROM      Juniper ROM Monitor Version 9.0b2
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 4          ROM      Juniper ROM Monitor Version 7.5b4

```

	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 5	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 7	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

#### show chassis firmware lcc (TX Matrix Plus Router)

```
user@host> show chassis firmware lcc 0
lcc0-re1:
```

Part	Type	Version
FPC 4	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 6	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 7	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

#### show chassis firmware sfc (TX Matrix Plus Router)

```
user@host> show chassis firmware sfc 0
sfc0-re0:
```

Part	Type	Version
Global FPC 4		
Global FPC 6		
Global FPC 7		
Global FPC 12		
Global FPC 14		
Global FPC 15		
Global FPC 20		
Global FPC 21		
Global FPC 22		
Global FPC 23		
Global FPC 24		
Global FPC 25		
Global FPC 26		
Global FPC 28		
Global FPC 29		
Global FPC 31		
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

#### show chassis firmware (QFX Series)

```
user@switch> show chassis firmware
```

Part	Type	Version
FPC 0		
Routing Engine 0	U-Boot loader	U-Boot 1.1.6 (Sep 15 2010 - 02:11:11) 1.0.5 FreeBSD/MIPS U-Boot bootstrap loader 0.1

**show chassis firmware interconnect-device (QFabric System)**

```

user@switch> show chassis firmware interconnect-device interconnect1
Part                Type      Version
Routing Engine 0    U-Boot    U-Boot 1.1.6 (May 10 2011 - 04:52:59) 1.1.1
                  loader    FreeBSD/MIPS U-Boot bootstrap loader 0.1
Routing Engine 1    U-Boot    U-Boot 1.1.6 (May 10 2011 - 04:52:59) 1.1.1
                  loader    FreeBSD/MIPS U-Boot bootstrap loader 0.1

```

**show chassis firmware (ACX2000 Universal Access Router)**

```

user@switch> show chassis firmware
Part      Type      Version
FPC       O/S       Version 12.2I13 by jisjoy on 2012-05-29 06:
FEB       O/S       Version 12.2I13 by jisjoy on 2012-05-29 06:

```

**show chassis firmware detail (EX3300 Switch)**

```

user@switch> show chassis firmware detail
FPC 0
  Boot SYSPLD          3
  PoE firmware         4.1.6
  PFE-0                3
  PFE-1                3
  PHY
    microcode          0x514
  Boot Firmware
    uboot              U-Boot 1.1.6 (Aug 21 2011 - 01:45:26) 1.0.0
    loader              FreeBSD/arm U-Boot loader 1.0

```

**show chassis firmware (MX Routers with Media Services Blade [MSB])**

```

user@switch> show chassis firmware
Part      Type      Version
FPC 1     ROM      Juniper ROM Monitor Version 12.1b1
          O/S      Version 12.2I21 by manish on 2012-06-19 17:

```

## show chassis lcd

---

<b>List of Syntax</b>	<a href="#">show chassis lcd (EX Series) on page 174</a> <a href="#">show chassis lcd (QFX Series and QFabric Systems) on page 174</a>
<b>show chassis lcd (EX Series)</b>	<code>show chassis lcd</code> <code>&lt;fpc-slot <i>fpc-slot-number</i>&gt;</code> <code>&lt;menu &lt;(all-members   local   member <i>member-id</i>)&gt;&gt;</code>
<b>show chassis lcd (QFX Series and QFabric Systems)</b>	<code>show chassis lcd</code> <code>&lt;fpc-slot <i>fpc-slot-number</i>&gt;</code> <code>&lt;interconnect-device <i>device-id</i>&gt;</code> <code>&lt;node-device <i>device-id</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. <code>menu</code> option introduced in Junos OS Release 10.2 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 13.1 for QFabric systems.
<b>Description</b>	Display the information that appears on the LCD panel of EX3200, EX3300, EX4200, EX4500, EX6200, and EX8200 switches, XRE200 External Routing Engines, QFX Series standalone switches, and Interconnect devices and Node devices within a QFabric system. Display the status of the currently selected port parameter of the Status LED for each network port on the device.
<b>Options</b>	<b>none</b> —Display the information that appears on the LCD panel (for any EX Series member switch in a Virtual Chassis or for XRE200 External Routing Engines, display the information for all Virtual Chassis members). Display the status of the currently selected port parameter of the Status LED for each network port.  <b>fpc-slot &lt;<i>fpc-slot-number</i>&gt;</b> —(Optional) Display the information as follows: <ul style="list-style-type: none"><li>(EX3200, EX3300, EX4200, and EX4500 switches, or the QFX Series) Display the information that appears on the LCD panel for either an FPC slot with no <i>fpc-slot-number</i> value specified or for the FPC slot specified by <b>fpc-slot 0</b>. <i>fpc-slot</i> refers to the switch itself and <b>0</b> is the only valid value for <i>fpc-slot-number</i>. Output for these options is the same as for the <b>none</b> option.  Also display the status of the currently selected port parameter of the Status LED for each network port.</li><li>(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) If no <i>fpc-slot-number</i> value is specified, display the information that appears on the LCD panel for all members of the Virtual Chassis. Output for this option is the same as for the <b>none</b> option. If the <i>fpc-slot-number</i> value is specified (it equals the <i>member-id</i> value), display the information for the specified member.  Also display the status of the currently selected port parameter of the Status LED for each network port.</li><li>(EX6200 or EX8200 switches)—Display the information that appears on the LCD panel for the line card in the line-card slot specified by the <i>fpc-slot-number</i> value.</li></ul>

Also display the status of the currently selected port parameter of the Status LED for each network port.

**interconnect-device *device-id***—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Interconnect device.

**menu**—(Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel.

**menu all-members**—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for all Virtual Chassis members.

**menu local**—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the Virtual Chassis member from which you issued the command.

**menu member *member-id***—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the specified Virtual Chassis member.

**node-device *device-id***—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Node device.

**Required Privilege Level**

view

**Related Documentation**

- [LCD Panel in EX3200 Switches](#)
- [LCD Panel in EX4200 Switches](#)
- [LCD Panel in EX4500 Switches](#)
- [LCD Panel in an EX8200 Switch](#)
- [LCD Panel in an XRE200 External Routing Engine](#)
- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\) on page 33](#)
- [set chassis display message on page 146](#)

**List of Sample Output**

[show chassis lcd \(Two-Member EX4200 Virtual Chassis\) on page 176](#)  
[show chassis lcd fpc-slot 1 \(EX4200 Virtual Chassis\) on page 178](#)  
[show chassis lcd \(EX8200 Switch\) on page 178](#)  
[show chassis lcd fpc-slot 2 \(EX8200 Switch\) on page 180](#)  
[show chassis lcd menu \(EX4200 Switch\) on page 180](#)  
[show chassis lcd menu \(EX8200 Switch\) on page 180](#)  
[show chassis lcd \(QFX3500 Switches\) on page 181](#)  
[show chassis lcd \(XRE200 External Routing Engine in EX8200 Virtual Chassis\) on page 181](#)  
[show chassis lcd interconnect-device \(QFabric Systems\) on page 184](#)

[show chassis lcd node-device \(QFabric Systems\) on page 186](#)

**Output Fields** Table 13 on page 176 lists the output fields for the **show chassis lcd** command. Output fields are listed in the approximate order in which they appear.

**Table 13: show chassis lcd Output Fields**

Field Name	Field Description
<b>membernumber</b> (XRE200 External Routing Engine)	Member ID of the device whose content is being displayed.
<b>Front panel contents for slot</b>	FPC slot number of the switch whose content is being displayed. The number is always <b>0</b> , except for EX4200 switches in a Virtual Chassis, where it is the member ID value.
<b>Front panel contents</b> (EX6200, EX8200 switch, XRE200 External Routing Engine, and QFX Series)	<p>On EX6200 switches, EX8200 switches, and XRE200 External Routing Engines, no slot number is displayed.</p> <p>On XRE200 External Routing Engines, this field appears under the <b>member number</b> field for each member device in the EX8200 Virtual Chassis.</p>
<b>LCD screen</b>	<p>The first line displays the hostname (for Virtual Chassis members, displays the member ID, the current role, and hostname; for EX8200 switches, displays <b>RE</b> and the hostname). The second line displays the currently selected port parameter of the Status LED and the alarms counter. The Status LED port parameters are:</p> <ul style="list-style-type: none"> <li>• <b>ADM</b>—Administrative</li> <li>• <b>SPD</b>—Speed</li> <li>• <b>DPX</b>—Duplex</li> <li>• <b>POE</b>—Power over Ethernet (EX3200 and EX4200 switches only)</li> </ul>
<b>LEDs status</b>	Current state of the Alarms, System, and Master LEDs (chassis status LEDs).
<b>Interface</b>	Names of the interfaces on the switch.
<b>LED (ADM/SPD/DPX/POE)</b>	<p>State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are:</p> <p><b>NOTE:</b> The XRE200 External Routing Engine always displays the <b>NA</b> parameter. The QFX Series products do not have any of the port parameters listed below.</p> <ul style="list-style-type: none"> <li>• <b>ADM</b>—Administrative</li> <li>• <b>SPD</b>—Speed</li> <li>• <b>DPX</b>—Duplex</li> <li>• <b>NA</b>—Not applicable.</li> <li>• <b>POE</b>—Power over Ethernet</li> </ul>
<b>fpcx</b>	On standalone EX Series and QFX Series switches, always <b>0</b> . On EX Series Virtual Chassis member switches, member ID of the Virtual Chassis member whose LCD menu is displayed.

## Sample Output

### show chassis lcd (Two-Member EX4200 Virtual Chassis)

```
user@switch> show chassis lcd
```



## Front panel contents for slot: 0

```

-----
LCD screen:
  00:BK switch1
  LED:SPD ALARM 00
LEDs status:
  Alarms LED: Off
  System LED: Green
  Master LED: Off
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      Off
ge-0/0/1      Off
ge-0/0/2      Off
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Off
ge-0/0/6      Off
ge-0/0/7      Off
ge-0/0/8      Off
ge-0/0/9      Off
ge-0/0/10     Off
ge-0/0/11     Off
ge-0/0/12     Off
ge-0/0/13     Off
ge-0/0/14     Off
ge-0/0/15     Off
ge-0/0/16     Off
ge-0/0/17     Off
ge-0/0/18     Off
ge-0/0/19     Off
ge-0/0/20     Off
ge-0/0/21     Off
ge-0/0/22     Off
ge-0/0/23     Off

```

## Front panel contents for slot: 1

```

-----
LCD screen:
  01:RE switch2
  LED:SPD ALARM 01
LEDs status:
  Alarms LED: Yellow
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-1/0/0      Off
ge-1/0/1      Off
ge-1/0/2      Off
ge-1/0/3      Off
ge-1/0/4      Off
ge-1/0/5      Off
ge-1/0/6      Off
ge-1/0/7      Off
ge-1/0/8      Off
ge-1/0/9      Off
ge-1/0/10     Off
ge-1/0/11     Off
ge-1/0/12     Off
ge-1/0/13     Off
ge-1/0/14     Off

```

ge-1/0/15	Off
ge-1/0/16	Off
ge-1/0/17	Off
ge-1/0/18	Off
ge-1/0/19	Off
ge-1/0/20	Off
ge-1/0/21	Off
ge-1/0/22	Off
ge-1/0/23	Off

The output for the **show chassis lcd fpc-slot** command is the same as the output for the **show chassis lcd** command.

#### show chassis lcd fpc-slot 1 (EX4200 Virtual Chassis)

```
user@switch> show chassis lcd fpc-slot 1
Front panel contents for slot: 1
-----
LCD screen:
  01:RE switch2
  LED:SPD ALARM 01
LEDs status:
  Alarms LED: Yellow
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-1/0/0       Off
ge-1/0/1       Off
ge-1/0/2       Off
ge-1/0/3       Off
ge-1/0/4       Off
ge-1/0/5       Off
ge-1/0/6       Off
ge-1/0/7       Off
ge-1/0/8       Off
ge-1/0/9       Off
ge-1/0/10      Off
ge-1/0/11      Off
ge-1/0/12      Off
ge-1/0/13      Off
ge-1/0/14      Off
ge-1/0/15      Off
ge-1/0/16      Off
ge-1/0/17      Off
ge-1/0/18      Off
ge-1/0/19      Off
ge-1/0/20      Off
ge-1/0/21      Off
ge-1/0/22      Off
ge-1/0/23      Off
```

#### show chassis lcd (EX8200 Switch)

```
user@switch> show chassis lcd
Front panel contents:
-----
LCD screen:
  RE st-8200-r
  LED:ADM ALARM 01
```

## LEDs status:

Alarms LED: Yellow

System LED: Yellow

Master LED: Green

Interface LED(ADM/SPD/DPX)

```

-----
ge-0/0/0      Off
ge-0/0/1      Off
ge-0/0/2      Off
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Off
ge-0/0/6      Off
ge-0/0/7      Off
ge-0/0/8      Off
ge-0/0/9      Off
ge-0/0/10     Off
ge-0/0/11     Off
ge-0/0/12     Off
ge-0/0/13     Off
ge-0/0/14     Off
ge-0/0/15     Off
ge-0/0/16     Off
ge-0/0/17     Off
ge-0/0/18     Off
ge-0/0/19     Off
ge-0/0/20     Off
ge-0/0/21     Off
ge-0/0/22     Off
ge-0/0/23     Off
ge-0/0/24     Off
ge-0/0/25     Off
ge-0/0/26     Off
ge-0/0/27     Off
ge-0/0/28     Off
ge-0/0/29     Off
ge-0/0/30     Off
ge-0/0/31     Off
ge-0/0/32     Off
ge-0/0/33     Off
ge-0/0/34     Off
ge-0/0/35     Off
ge-0/0/36     Off
ge-0/0/37     Off
ge-0/0/38     Off
ge-0/0/39     Off
ge-0/0/40     Off
ge-0/0/41     Off
ge-0/0/42     Off
ge-0/0/43     Off
ge-0/0/44     Off
ge-0/0/45     Off
ge-0/0/46     Off
ge-0/0/47     Off
xe-2/0/0      Off
xe-2/0/1      Off
xe-2/0/2      Off
xe-2/0/3      Off
xe-2/0/4      Off
xe-2/0/5      Off
xe-2/0/6      Off

```

xe-2/0/7	Off
xe-3/0/0	Off
xe-3/0/1	Off
xe-3/0/2	Off
xe-3/0/3	Off
xe-3/0/4	Off
xe-3/0/5	Off
xe-3/0/6	Off
xe-3/0/7	Off
xe-5/0/0	Off
xe-5/0/1	Off
xe-5/0/2	Off
xe-5/0/3	Off
xe-5/0/4	Off
xe-5/0/5	Off
xe-5/0/6	On
xe-5/0/7	On
xe-7/0/5	Off

#### show chassis lcd fpc-slot 2 (EX8200 Switch)

show chassis lcd fpc-slot 2

Interface	LED (ADM/SPD/DPX)
xe-2/0/0	Off
xe-2/0/1	Off
xe-2/0/2	Off
xe-2/0/3	Off
xe-2/0/4	Off
xe-2/0/5	Off
xe-2/0/6	Off
xe-2/0/7	Off

#### show chassis lcd menu (EX4200 Switch)

```
user@switch> show chassis lcd menu
fpc0:
```

```
-----
status-menu
status-menu vcp-status
status-menu power-status
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu vc-uplink-config
maintenance-menu factory-default
```

On an EX4200 switch in a Virtual Chassis, the output for the **show chassis lcd menu** **all-members** command is the same as the output for the **show chassis lcd menu** command.

#### show chassis lcd menu (EX8200 Switch)

```
user@switch> show chassis lcd menu
status-menu
status-menu sf-status1-menu
status-menu sf-status2-menu
status-menu psu-status1-menu
```

```

status-menu psu-status2-menu
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu factory-default

```

### show chassis lcd (QFX3500 Switches)

```

user@switch> show chassis lcd
Front panel contents for slot: 0
-----
LCD screen:
00:RE switch
ALARM 01
LEDs status:
Status/Beacon LED: Yellow Blinking
Interface STATUS LED ACTIVITY LED
-----
fte-0/1/0 Off Off

```

### show chassis lcd (XRE200 External Routing Engine in EX8200 Virtual Chassis)

```

user@external-routing-engine> show chassis lcd
member0:
-----
Front panel contents:
-----
LCD screen:
  RE ex8200-member0
  LED:ADM ALARM 04
LEDs status:
  Alarms LED: Red
  System LED: Yellow
  Master LED: Green

member1:
-----

member8:
-----
Front panel contents:
-----
LCD screen:
  BACKUP

member9:
-----
Front panel contents:
-----
LCD screen:
  09:RE xre200-member9
  LED: NA ALARM 01
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      On
ge-0/0/1      On
ge-0/0/2      On
ge-0/0/3      On

```

ge-0/0/4	Off
ge-0/0/5	Off
ge-0/0/6	Off
ge-0/0/7	Off
ge-0/0/8	Off
ge-0/0/9	Off
ge-0/0/10	On
ge-0/0/11	Off
ge-0/0/12	Off
ge-0/0/13	Off
ge-0/0/14	Off
ge-0/0/15	Off
ge-0/0/16	Off
ge-0/0/17	Off
ge-0/0/18	Off
ge-0/0/19	Off
ge-0/0/20	Off
ge-0/0/21	Off
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Off
ge-0/0/25	Off
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Off
ge-0/0/29	Off
ge-0/0/30	Off
ge-0/0/31	Off
ge-0/0/32	Off
ge-0/0/33	Off
ge-0/0/34	Off
ge-0/0/35	Off
ge-0/0/36	Off
ge-0/0/37	Off
ge-0/0/38	Off
ge-0/0/39	Off
ge-0/0/40	On
ge-0/0/41	On
ge-0/0/42	On
ge-0/0/43	On
ge-0/0/44	On
ge-0/0/45	On
ge-0/0/46	On
ge-0/0/47	On
ge-16/0/0	On
ge-16/0/1	Off
ge-16/0/2	On
ge-16/0/3	Off
ge-16/0/4	On
ge-16/0/5	Off
ge-16/0/6	On
ge-16/0/7	Off
ge-16/0/8	Off
ge-16/0/9	Off
ge-16/0/10	Off
ge-16/0/11	Off
ge-16/0/12	Off
ge-16/0/13	On
ge-16/0/14	Off
ge-16/0/15	On
ge-16/0/16	Off

ge-16/0/17	On
ge-16/0/18	On
ge-16/0/19	On
ge-16/0/20	On
ge-16/0/21	Off
ge-16/0/22	On
ge-16/0/23	Off
ge-16/0/24	Off
ge-16/0/25	Off
ge-16/0/26	On
ge-16/0/27	Off
ge-16/0/28	Off
ge-16/0/29	Off
ge-16/0/30	On
ge-16/0/31	Off
ge-16/0/32	On
ge-16/0/33	On
ge-16/0/34	On
ge-16/0/35	Off
ge-16/0/36	On
ge-16/0/37	Off
ge-16/0/38	Off
ge-16/0/39	Off
ge-16/0/40	Off
ge-16/0/41	Off
ge-16/0/42	On
ge-16/0/43	Off
ge-16/0/44	Off
ge-16/0/45	Off
ge-16/0/46	Off
ge-16/0/47	Off
xe-19/0/0	Off
xe-19/0/1	On
xe-19/0/2	On
xe-19/0/3	On
xe-19/0/4	On
xe-19/0/5	On
ge-22/0/0	Off
ge-22/0/1	Off
ge-22/0/2	On
ge-22/0/3	Off
ge-22/0/4	On
ge-22/0/5	On
ge-22/0/6	On
ge-22/0/7	On
ge-22/0/8	Off
ge-22/0/9	Off
ge-22/0/10	Off
ge-22/0/11	Off
ge-22/0/12	Off
ge-22/0/13	Off
ge-22/0/14	Off
ge-22/0/15	Off
ge-22/0/16	On
ge-22/0/17	Off
ge-22/0/18	On
ge-22/0/19	Off
ge-22/0/20	On
ge-22/0/21	Off
ge-22/0/22	On
ge-22/0/23	Off

```
ge-22/0/24    On
ge-22/0/25    Off
ge-22/0/26    Off
ge-22/0/27    Off
ge-22/0/28    Off
ge-22/0/29    Off
ge-22/0/30    Off
ge-22/0/31    Off
ge-22/0/32    On
ge-22/0/33    Off
ge-22/0/34    On
ge-22/0/35    Off
ge-22/0/36    Off
ge-22/0/37    Off
ge-22/0/38    Off
ge-22/0/39    Off
ge-22/0/40    Off
ge-22/0/41    Off
ge-22/0/42    Off
ge-22/0/43    Off
ge-22/0/44    Off
ge-22/0/45    Off
ge-22/0/46    Off
ge-22/0/47    Off
```

#### show chassis lcd interconnect-device (QFabric Systems)

```
show chassis lcd interconnect-device IC-F1012
      Front Panel Module Information
      -----
      LCD screen:
      IC-F1012          3 Alarms active

LEDs status:
  Status LED: Green
  Power LED : Green
  Major Alarm LED: off
  Minor Alarm LED: Yellow
  Fan 0 LED : Green
  Fan 1 LED : Green
  Fan 2 LED : Green
  Fan 3 LED : Green
  Fan 4 LED : Green
  Fan 5 LED : Green
  Fan 6 LED : Green
  Fan 7 LED : Green
  Fan 8 LED : Green
  Fan 9 LED : Green
  PEM 0 LED : Green
  PEM 1 LED : Green
  PEM 2 LED : Green
  PEM 3 LED : off
  PEM 4 LED : off
  PEM 5 LED : off

      LED info for: CB - 0
      -----

LEDs status:
  Status LED: Green
  Mastership LED: Green

Interface          STATUS LED    LINK/ACTIVITY LED
```



```

-----
IC-F1012:pme0 :      Green      N/A
IC-F1012:pme1 :      Green      N/A
IC-F1012:pme2 :      off        N/A
IC-F1012:pme3 :      off        N/A

```

```

LED info for: CB - 1
-----

```

```

LEDs status:
  Status LED: Green
  Mastership LED: Amber

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:pme0 :	Green	N/A
IC-F1012:pme1 :	Green	N/A
IC-F1012:pme2 :	off	N/A
IC-F1012:pme3 :	off	N/A

```

LED info for: FC 0 FPC - 0
-----

```

```

LEDs status:
  Status LED: Green

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-0/0/0	Green	N/A
IC-F1012:fte-0/0/1	Green	N/A
IC-F1012:fte-0/0/2	Green	N/A
IC-F1012:fte-0/0/3	Green	N/A
IC-F1012:fte-0/0/4	Green	N/A

```

LED info for: FC 1 FPC - 1
-----

```

```

LEDs status:
  Status LED: Green

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-1/0/0	Green	N/A
IC-F1012:fte-1/0/1	Green	N/A
IC-F1012:fte-1/0/2	Green	N/A
IC-F1012:fte-1/0/3	Green	N/A
IC-F1012:fte-1/0/4	Green	N/A

```

LED info for: RC 0 FPC - 8
-----

```

```

LEDs status:
  Status LED: Green

```

```

LED info for: RC 1 FPC - 9
-----

```

```

LEDs status:
  Status LED: Green

```

```

LED info for: RC 2 FPC - 10
-----

```

```

LEDs status:
  Status LED: Green

```

```

LED info for: RC 3 FPC - 11

```

```

-----
LEDs status:
  Status LED: Green

      LED info for: RC 4 FPC - 12
-----
LEDs status:
  Status LED: Green

      LED info for: RC 5 FPC - 13
-----
LEDs status:
  Status LED: Green

      LED info for: RC 6 FPC - 14
-----
LEDs status:
  Status LED: Green

      LED info for: RC 7 FPC - 15
-----
LEDs status:
  Status LED: Green

```

#### show chassis lcd node-device (QFabric Systems)

```

show chassis lcd node-device P3774-C
  Front panel contents for: P3774-C
-----
  LCD screen:
  P3774-C
LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	STATUS LED	LINK/ACTIVITY LED
P3774-C:xe-0/0/6	Green	Green
P3774-C:xe-0/0/7	Green	Green
P3774-C:ge-0/0/10	Green	Green
P3774-C:ge-0/0/11	Green	Green Blinking
P3774-C:ge-0/0/12	Green	Off
P3774-C:ge-0/0/13	Green	Green Blinking
P3774-C:ge-0/0/20	Green	Green
P3774-C:ge-0/0/21	Green	Green
P3774-C:ge-0/0/22	Green	Green Blinking
P3774-C:ge-0/0/23	Green	Off
P3774-C:ge-0/0/30	Green	Green
P3774-C:ge-0/0/31	Green	Green
P3774-C:ge-0/0/32	Green	Green Blinking
P3774-C:ge-0/0/33	Green	Green Blinking
P3774-C:fte-0/1/0	Green	Green
P3774-C:fte-0/1/1	Green	Green Blinking
P3774-C:fte-0/1/2	Green	Green Blinking
P3774-C:fte-0/1/3	Green	Green

## show configuration

---

<b>Syntax</b>	<code>show configuration</code> <code>&lt;statement-path&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Display the configuration that currently is running on the router or switch, which is the last committed configuration.
<b>Options</b>	<p><b>none</b>—Display the entire configuration.</p> <p><b>statement-path</b>—(Optional) Display one of the following hierarchies in a configuration. (Each <b>statement-path</b> option has additional suboptions not described here. See the appropriate configuration guide or EX Series switch documentation for more information.)</p> <ul style="list-style-type: none"> <li>• <b>access</b>—Network access configuration.</li> <li>• <b>access-profile</b>—Access profile configuration.</li> <li>• <b>accounting-options</b>—Accounting data configuration.</li> <li>• <b>applications</b>—Applications defined by protocol characteristics.</li> <li>• <b>apply-groups</b>—Groups from which configuration data is inherited.</li> <li>• <b>chassis</b>—Chassis configuration.</li> <li>• <b>chassis network-services</b>—Current running mode.</li> <li>• <b>class-of-service</b>—Class-of-service configuration.</li> <li>• <b>diameter</b>—Diameter base protocol layer configuration.</li> <li>• <b>ethernet-switching-options</b>—(EX Series switch only) Ethernet switching configuration.</li> <li>• <b>event-options</b>—Event processing configuration.</li> <li>• <b>firewall</b>—Firewall configuration.</li> <li>• <b>forwarding-options</b>—Options that control packet sampling.</li> <li>• <b>groups</b>—Configuration groups.</li> <li>• <b>interfaces</b>—Interface configuration.</li> <li>• <b>jsrc</b>—JSRC partition configuration.</li> <li>• <b>jsrc-partition</b>—JSRC partition configuration.</li> <li>• <b>logical-systems</b>—Logical system configuration.</li> <li>• <b>poe</b>—(EX Series switch only) Power over Ethernet configuration.</li> <li>• <b>policy-options</b>—Routing policy option configuration.</li> <li>• <b>protocols</b>—Routing protocol configuration.</li> </ul>

- **routing-instances**—Routing instance configuration.
- **routing-options**—Protocol-independent routing option configuration.
- **security**—Security configuration.
- **services**—Service PIC applications configuration.
- **snmp**—Simple Network Management Protocol configuration.
- **system**—System parameters configuration.
- **virtual-chassis**—(EX Series switch only) Virtual Chassis configuration.
- **vlan**—(EX Series switch only) VLAN configuration.

**Additional Information** The portions of the configuration that you can view depend on the user class that you belong to and the corresponding permissions. If you do not have permission to view a portion of the configuration, the text **ACCESS-DENIED** is substituted for that portion of the configuration. If you do not have permission to view authentication keys and passwords in the configuration, because the secret permission bit is not set for your user account, the text **SECRET-DATA** is substituted for that portion of the configuration. If an identifier in the configuration contains a space, the identifier is displayed in quotation marks.

Likewise, when you issue the **show configuration** command with the **| display set** pipe option to view the configuration as **set** commands, those portions of the configuration that you do not have permissions to view are substituted with the text **ACCESS-DENIED**.

**Required Privilege Level**

view

**Related Documentation**

- *Displaying the Current Junos OS Configuration*
- *Overview of Junos OS CLI Operational Mode Commands*

**List of Sample Output**

[show configuration on page 188](#)  
[show configuration policy-options on page 189](#)

**Output Fields**

This command displays information about the current running configuration.

## Sample Output

**show configuration**

```
user@host> show configuration
## Last commit: 2006-10-31 14:13:00 PST by alant version "8.2IO [builder]"; ##
last changed: 2006-10-31 14:05:53 PST
system {
    host-name nestor;
    domain-name east.net;
    backup-router 192.1.1.254;
    time-zone America/Los_Angeles;
    default-address-selection;
    name-server {
        192.154.169.254;
        192.154.169.249;
```

```

        192.154.169.176;
    }
    services {
        telnet;
    }
    tacplus-server {
        1.2.3.4 {
            secret /* SECRET-DATA */;
            ...
        }
    }
}
interfaces {
    ...
}
protocols {
    isis {
        export "direct routes";
    }
}
policy-options {
    policy-statement "direct routes" {
        from protocol direct;
        then accept;
    }
}

```

#### show configuration policy-options

```

user@host> show configuration policy-options
policy-options {
    policy-statement "direct routes" {
        from protocol direct;
        then accept;
    }
}

```

## show host

---

<b>Syntax</b>	<code>show host <i>hostname</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display Domain Name System (DNS) hostname information.
<b>Options</b>	<i>hostname</i> —Hostname or address.
<b>Additional Information</b>	The <b>show host</b> command displays the raw data received from the DNS server.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show host on page 190</a>

## Sample Output

### show host

```
user@host> show host snark
snark.boojum.net has address 192.168.1.254

user@host> show host 192.168.1.254
Name: snark.boojum.net
Address: 192.168.1.254
Aliases:
```

## show ntp associations

<b>Syntax</b>	<code>show ntp associations</code> <code>&lt;no-resolve&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display Network Time Protocol (NTP) peers and their state.
<b>Options</b>	<b>none</b> —Display NTP peers and their state.  <b>no-resolve</b> —(Optional) Suppress symbolic addressing.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ntp status on page 193</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ntp associations on page 192</a>
<b>Output Fields</b>	<a href="#">Table 14 on page 191</a> describes the output fields for the <b>show ntp associations</b> command. Output fields are listed in the approximate order in which they appear.

**Table 14: show ntp associations Output Fields**

Field Name	Field Description
<b>remote</b>	Address or name of the remote NTP peer.
<b>refid</b>	Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of <b>0.0.0.0</b> .
<b>st</b>	Stratum of the remote peer.
<b>t</b>	Type of peer: <b>b</b> (broadcast), <b>l</b> (local), <b>m</b> (multicast), or <b>u</b> (unicast).
<b>when</b>	When the last packet from the peer was received.
<b>poll</b>	Polling interval, in seconds.
<b>reach</b>	Reachability register, in octal.
<b>delay</b>	Current estimated delay of the peer, in milliseconds.
<b>offset</b>	Current estimated offset of the peer, in milliseconds.
<b>disp</b>	Current estimated dispersion of the peer, in milliseconds.

Table 14: show ntp associations Output Fields (*continued*)

Field Name	Field Description
<i>peer-name</i>	<p>Peer name and status of the peer in the clock selection process:</p> <ul style="list-style-type: none"> <li>• space—Discarded because of a high stratum value or failed sanity checks.</li> <li>• x—Designated "falseticker" by the intersection algorithm.</li> <li>• .—Culled from the end of the candidate list.</li> <li>• — —Discarded by the clustering algorithm.</li> <li>• +—Included in the final selection set.</li> <li>• #—Selected for synchronization, but the distance exceeds the maximum.</li> <li>• *—Selected for synchronization.</li> <li>• o—Selected for synchronization, but the packets-per-second (pps) signal is in use.</li> </ul>

## Sample Output

### show ntp associations

```

user@host> show ntp associations
      remote          refid      st t when poll reach  delay  offset  disp
=====
*wolfe-gw.junipe tick.ucla.edu    2 u  43   64  377   1.86   0.319   0.08

```



## show ntp status

<b>Syntax</b>	<code>show ntp status</code> <code>&lt;no-resolve&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the values of internal variables returned by Network Time Protocol (NTP) peers.
<b>Options</b>	<b>none</b> —Display the values of internal variables returned by NTP peers.  <b>no-resolve</b> —(Optional) Suppress symbolic addressing.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ntp associations on page 191</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ntp status on page 194</a>
<b>Output Fields</b>	<a href="#">Table 15 on page 193</a> describes the output fields for the <b>show ntp status</b> command. Output fields are listed in the approximate order in which they appear.

**Table 15: show ntp status Output Fields**

Field Name	Field Description
<b>status</b>	System status word, a code representing the status items listed.
<b>leap_none</b>	Indicates a normal synchronized state with no leap seconds imminent. Other options could be <b>leap_add_sec</b> , <b>leap_del_sec</b> , or <b>leap_alarm</b> , indicating a leap second will be added, deleted, or a leap second requirement is upcoming.
<b>sync_ntp</b>	Indicates the current synchronization source, in this case, an NTP server. Other options include <b>sync_alarm</b> and <b>sync_unspec</b> , both indicating that the router has not been synched.
<b>x events</b>	Indicates the number of events that have occurred since that last code change. An event is often the receipt of an NTP polling message.
<b>event_peer/strat_chg</b>	Describes the most recent event, in this case, the stratum of the peer server changed.
<b>version</b>	A detailed description of the version of NTP being used.
<b>processor</b>	Indicates the current hardware platform and version of the processor.
<b>system</b>	Detailed description of the name and version of the operating system in use.
<b>leap</b>	The number of leap seconds in use.

Table 15: show ntp status Output Fields (*continued*)

Field Name	Field Description
<b>stratum</b>	The stratum of the peer server. Anything greater than 1 is a secondary reference source, and the number roughly represents the number of hops away from the stratum 1 server.. Stratum 1 is a primary reference, such as an atomic clock.
<b>precision</b>	The precision of the peer clock, how precisely the frequency and time can be maintained with this particular timekeeping system.
<b>rootdelay</b>	The total roundtrip delay to the primary reference source, in seconds.
<b>rootdispersion</b>	The maximum error relative to the primary reference source, in seconds.
<b>peer</b>	An identification number of the peer in use.
<b>refid</b>	Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of 0.0.0.0.
<b>reftime</b>	The local time, in timestamp format, when the local clock was last updated. If the local clock has never been synchronized, the value is zero.
<b>poll</b>	The NTP broadcast message polling interval, in seconds.
<b>clock</b>	The current time on the local router clock.
<b>state</b>	The current mode of NTP operation, where 1 is symmetric active, 2 is symmetric passive, 3 is client, 4 is server, and 5 is broadcast.
<b>offset</b>	Current estimated offset of the peer, in milliseconds. Indicates the time difference between the reference clock and the local clock.
<b>frequency</b>	The frequency of the clock.
<b>jitter</b>	Indicates the magnitude of jitter, in milliseconds, between several time queries.
<b>stability</b>	A measure of how well this clock can maintain a constant frequency.

## Sample Output

### show ntp status

```

user@host> show ntp status
assID=0 status=0544 leap_none, sync_local_proto, 4 events, event_peer/strat_chg,
version="ntpd 4.2.2p1@1.1570-o Tue May 19 13:57:55 UTC 2009 (1)",
processor="x86_64", system="Linux/2.6.18-164.el5", leap=00, stratum=4,
precision=-10, rootdelay=0.000, rootdispersion=11.974, peer=59475,
refid=LOCAL(0),
reftime=d495c32c.0e71eaf2 Mon, Jan 7 2013 13:57:00.056, poll=10,
clock=d495c32c.cebd43bd Mon, Jan 7 2013 13:57:00.807, state=4,
offset=0.000, frequency=0.000, jitter=0.977, noise=0.977,
stability=0.000, tai=0

```



## show system firmware

**Syntax** show system firmware  
<compatibility>

**Release Information** Command introduced in Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.4 for EX Series switches.

**Description** (J Series routers and EX8200 switches only) Display firmware information.



**NOTE:** On SRX100, SRX210, SRX240, and SRX 650 devices, the **show system firmware** command now displays all the installed firmware versions, even if the installed firmware versions are earlier than the currently installed firmware version.

**Options** **compatibility**—(Optional) Display firmware compatibility information.

**Required Privilege Level** view

**List of Sample Output** [show system firmware on page 197](#)  
[show system firmware compatibility on page 197](#)

**Output Fields** [Table 16 on page 196](#) lists the output fields for the show system firmware command. Output fields are listed in the approximate order in which they appear.

**Table 16: show system firmware Output Fields**

Field Name	Field Description
<b>Part</b>	Physical part on the router or switch affected by the firmware.
<b>Type</b>	Type of firmware on the router or switch.
<b>Tag</b>	Location of the firmware on the interface.
<b>Current version</b>	Firmware version on the affected router or switch parts.
<b>Available version</b>	New versions of firmware for upgrading or downgrading.
<b>Status</b>	Firmware condition on the router or switch.
<b>Action</b>	Whether you can upgrade or downgrade, or if no action is available ( <b>none</b> ).

## Sample Output

### show system firmware

```
user@host> show system firmware
```

Part	Type	Tag	Current version	Available version	Status
FPC 0	ROM Monitor	0	0	6.4.10	OK
Routing Engine 0	RE BIOS	0	0		OK

### show system firmware compatibility

```
user@host> show system firmware compatibility
```

Part	Type	Tag	Current version	Available version	Action
FPC 0	ROM Monitor	0	0	6.4.10	None
Routing Engine 0	RE BIOS	0	0		None

## show system reboot

---

<b>List of Syntax</b>	<a href="#">Syntax on page 198</a> <a href="#">Syntax (EX Series Switches) on page 198</a> <a href="#">Syntax (TX Matrix Router) on page 198</a> <a href="#">Syntax (TX Matrix Plus Router) on page 198</a> <a href="#">Syntax (MX Series Router) on page 198</a> <a href="#">Syntax (QFX Series) on page 198</a>
<b>Syntax</b>	show system reboot <both-routing-engines>
<b>Syntax (EX Series Switches)</b>	show system reboot <all-members> <both-routing-engines> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	show system reboot <all-chassis   all-lcc   lcc <i>number</i>   scc> <both-routing-engines>
<b>Syntax (TX Matrix Plus Router)</b>	show system reboot <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> > <both-routing-engines>
<b>Syntax (MX Series Router)</b>	show system reboot <all-members> <both-routing-engines> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show system reboot <both-routing-engines> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-device <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display pending system reboots or halts.
<b>Options</b>	<b>none</b> —Display pending reboots or halts on the active Routing Engine.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display halt or reboot request information for all the T640 routers in the chassis that are connected to the TX Matrix router. On a TX Matrix router, display halt or reboot request information for all the T1600 or T4000 routers in the chassis that are connected to the TX Matrix Plus router.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for all members of the Virtual Chassis configuration.

**all-lcc**—(TX Matrix routers and TX Matrix Plus router only) (Optional) On a TX Matrix router, display system halt or reboot request information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display halt or reboot request information for all connected T1600 or T4000 LCCs.

**both-routing-engines**—(Systems with multiple Routing Engines) (Optional) Display halt or reboot request information on both Routing Engines.

**infrastructure *name***—(QFabric systems only) (Optional) Display reboot request information on the fabric manager Routing Engines and fabric control Routing Engines.

**interconnect-device *name***—(QFabric systems only) (Optional) Display reboot request information on the Interconnect device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display halt or reboot request information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display halt or reboot request information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display reboot request information on the Node group.

**scc**—(TX Matrix router only) (Optional) Display halt or reboot request information for the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus router only) (Optional) Display halt or reboot request information for the TX Matrix Plus router.

**Additional Information** By default, when you issue the **show system reboot** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) backup Routing Engines that are connected to it.

**Required Privilege Level** maintenance

**Related Documentation**

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [show system reboot on page 200](#)  
[show system reboot all-lcc \(TX Matrix Router\) on page 200](#)  
[show system reboot sfc \(TX Matrix Plus Router\) on page 200](#)  
[show system reboot \(QFX3500 Switch\) on page 200](#)

## Sample Output

### [show system reboot](#)

```
user@host> show system reboot
reboot requested by root at Wed Feb 10 17:40:46 1999
[process id 17885]
```

### [show system reboot all-lcc \(TX Matrix Router\)](#)

```
user@host> show system reboot all-lcc
lcc0-re0:
```

```
-----
No shutdown/reboot scheduled.
```

```
lcc2-re0:
```

```
-----
No shutdown/reboot scheduled.
```

### [show system reboot sfc \(TX Matrix Plus Router\)](#)

```
user@host> show system sfc 0
No shutdown/reboot scheduled.
```

### [show system reboot \(QFX3500 Switch\)](#)

```
user@switch> show system reboot
No shutdown/reboot scheduled.
```



## show system software

<b>List of Syntax</b>	<a href="#">Syntax on page 201</a> <a href="#">Syntax (EX Series Switches) on page 201</a> <a href="#">Syntax (TX Matrix Router) on page 201</a> <a href="#">Syntax (TX Matrix Plus Router) on page 201</a> <a href="#">Syntax (J Series Routers) on page 201</a> <a href="#">Syntax (QFX Series) on page 201</a>
<b>Syntax</b>	show system software <detail>
<b>Syntax (EX Series Switches)</b>	show system software <all-members> <detail> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	show system software <all-chassis   all-lcc   lcc <i>number</i>   scc> <detail>
<b>Syntax (TX Matrix Plus Router)</b>	show system software <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> > <detail>
<b>Syntax (J Series Routers)</b>	show system software <backup> <detail>
<b>Syntax (QFX Series)</b>	show system software <detail> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the Junos OS extensions loaded on your router or switch.
<b>Options</b>	<b>none</b> —Display standard information about all loaded Junos OS extensions.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system software information for all the T640 routers (TX Matrix Router) or all the routers (TX Matrix Plus Router) in the chassis.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system software information for all T640 routers connected to the

TX Matrix router. On a TX Matrix Plus router, display system software information for all connected T1600 or T4000 LCCs.

**all-members**—(EX4200 switches only) (Optional) Display the system software running on all members of the Virtual Chassis configuration.

**backup**—(J Series routers only) (Optional) Display the status of old system software packages only.

**detail**—(Optional) Display detailed information about available Junos OS extensions.

**infrastructure name**—(QFabric systems only) (Optional) Display the system software running on the fabric control Routing Engine and the fabric manager Routing Engine.

**interconnect-device name**—(QFabric systems only) (Optional) Display the system software running on the Interconnect device.

**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system software information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system software information for a specific router that is connected to the TX Matrix Plus router. Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches only) (Optional) Display the system software running on the local Virtual Chassis member.

**member member-id**—(EX4200 switches only) (Optional) Display the system software running on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

**node-group name**—(QFabric systems only) (Optional) Display the system software running on the Node group.

**scc**—(Routing matrix only) (Optional) Display the system software running on a TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display system software information for the TX Matrix Plus router.

**Required Privilege Level** maintenance

Related Documentation	<ul style="list-style-type: none"> <li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li> </ul>
List of Sample Output	<a href="#">show system software on page 203</a> <a href="#">show system software (TX Matrix Plus Router) on page 203</a> <a href="#">show system software (QFX Series) on page 207</a>
Output Fields	When you enter this command, you are provided a list of Junos OS packages installed on the router and their corresponding Junos OS release number.

## Sample Output

### show system software

```

user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]

Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]
Information for jdocs:

Comment:
JUNOS Online Documentation [7.2R1.7]

Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

Information for jpfe:

Comment:
JUNOS Packet Forwarding Engine Support (M20/M40) [7.2R1.7]

Information for jroute:

Comment:
JUNOS Routing Software Suite [7.2R1.7]

Information for junos:

Comment:
JUNOS Base OS boot [7.2R1.7]
```

### show system software (TX Matrix Plus Router)

```

user@host> show system software
sfc0-re0:
-----
Information for jbase:
```

Comment:  
JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:

Comment:  
JUNOS Crypto Software Suite [9.6-20090515.0]

Information for jdocs:

Comment:  
JUNOS Online Documentation [9.6-20090515.0]  
Information for jkernel:

Comment:  
JUNOS Kernel Software Suite [9.6-20090515.0]

Information for jpfe:

Comment:  
JUNOS Packet Forwarding Engine Support (T-Series) [9.6-20090515.0]

Information for jpfe-common:

Comment:  
JUNOS Packet Forwarding Engine Support (M/T Common) [9.6-20090515.0]

Information for jroute:Comment:  
JUNOS Routing Software Suite [9.6-20090515.0]

Information for jservices-aacl:

Comment:  
JUNOS Services ACL Container package [9.6-20090515.0]

Information for jservices-appid:

Comment:  
JUNOS AppId Services [9.6-20090515.0]

Information for jservices-bgf:

Comment:  
JUNOS Border Gateway Function package [9.6-20090515.0]

Information for jservices-idp:

Comment:  
JUNOS IDP Services [9.6-20090515.0]

Information for jservices-llpdf:

Comment:  
JUNOS Services LL-PDF Container package [9.6-20090515.0]

Information for jservices-sfw:

Comment:  
JUNOS Services Stateful Firewall [9.6-20090515.0]  
Information for jservices-voice:

Comment:  
JUNOS Voice Services Container package [9.6-20090515.0]

Information for junos:

Comment:  
JUNOS Base OS boot [9.6-20090515.0]  
...  
lcc0-re0:

-----  
Information for jbase:

Comment:  
JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:

Comment:  
JUNOS Crypto Software Suite [9.6-20090515.0]

Information for jdocs:

Comment:  
JUNOS Online Documentation [9.6-20090515.0]

Information for jkernel:

Comment:  
JUNOS Kernel Software Suite [9.6-20090515.0]

Information for jpfe:

Comment:  
JUNOS Packet Forwarding Engine Support (T-Series) [9.6-20090515.0]

Information for jpfe-common:

Comment:  
JUNOS Packet Forwarding Engine Support (M/T Common) [9.6-20090515.0]

Information for jroute:

Comment:  
JUNOS Routing Software Suite [9.6-20090515.0]

Information for jservices-aacl:

Comment:  
JUNOS Services ACL Container package [9.6-20090515.0]

Information for jservices-appid:

Comment:  
JUNOS AppId Services [9.6-20090515.0]

Information for jservices-bgf:

Comment:  
JUNOS Border Gateway Function package [9.6-20090515.0]

Information for jservices-idp:

Comment:  
JUNOS IDP Services [9.6-20090515.0]

Information for jservices-llpdf:

Comment:  
JUNOS Services LL-PDF Container package [9.6-20090515.0]

Information for jservices-sfw:

Comment:  
JUNOS Services Stateful Firewall [9.6-20090515.0]

Information for jservices-voice:

Comment:

JUNOS Voice Services Container package [9.6-20090515.0]

Information for junos:

Comment:

JUNOS Base OS boot [9.6-20090515.0]

lcc1-re0:

-----  
Information for jbase:

Comment:

JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [9.6-20090515.0]

...

### show system software (QFX Series)

user@switch> **show system software**

Information for jbase:

Comment:

JUNOS Base OS Software Suite [11.3-20110730.0]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [11.3-20110730.0]

Information for jdocs:

Comment:

JUNOS Online Documentation [11.3-20110730.0]

Information for jkernel:

Comment:

JUNOS Kernel Software Suite [11.3-20110730.0]

Information for jpfe:

Comment:

JUNOS Packet Forwarding Engine Support (QFX) [11.3-20110730.0]

Information for jroute:

Comment:

JUNOS Routing Software Suite [11.3-20110730.0]

Information for jswitch:

Comment:

JUNOS Enterprise Software Suite [11.3-20110730.0]

Information for junos:

Comment:

JUNOS Base OS boot [11.3-20110730.0]

Information for jweb:

Comment:

JUNOS Web Management [11.3-20110730.0]



## show system storage

<b>List of Syntax</b>	<a href="#">Syntax on page 209</a> <a href="#">Syntax (EX Series Switches) on page 209</a> <a href="#">Syntax (MX Series Router) on page 209</a> <a href="#">Syntax (QFX Series) on page 209</a> <a href="#">Syntax (SRX Series) on page 209</a> <a href="#">Syntax (TX Matrix Router) on page 209</a> <a href="#">Syntax (TX Matrix Plus Router and TX Matrix Plus Router with 3D SIBs) on page 209</a>
<b>Syntax</b>	<pre>show system storage &lt;detail&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show system storage &lt;detail&gt; &lt;all-members&gt; &lt;local&gt; &lt;member member-id&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (MX Series Router)</b>	<pre>show system storage &lt;detail&gt; &lt;all-members&gt; &lt;local&gt; &lt;member member-id&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show system storage &lt;detail&gt; &lt;infrastructure name&gt; &lt;interconnect-device name&gt; &lt;node-group name&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (SRX Series)</b>	<pre>show system storage &lt;detail&gt; &lt;partitions&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>show system storage &lt;detail&gt; &lt;all-chassis   all-lcc   lcc number   scc&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (TX Matrix Plus Router and TX Matrix Plus Router with 3D SIBs)</b>	<pre>show system storage &lt;detail&gt; &lt;all-chassis   all-lcc   lcc number   sfc number&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>

**sfc** option introduced for the TX Matrix Plus router in JUNOS Release 9.6.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Option **invoke-on (all-routing-engines | other-routing-engine)** introduced in Junos OS Release 14.1

<b>Description</b>	Display statistics about the amount of free disk space in the router's or switch's file systems.
<b>Options</b>	<p><b>none</b>—Display standard information about the amount of free disk space in the router's or switch's file systems.</p> <p><b>detail</b>—(Optional) Display detailed output.</p> <p><b>invoke-on all-routing-engines</b>—(Optional) Display the system storage information on all master and backup Routing Engines on a routing matrix based on the TX Matrix or TX Matrix Plus router or on a router that has dual Routing Engines.</p> <p><b>invoke-on other-routing-engines</b>—(Optional) Display the system storage information on the other Routing Engine. For example, if you issue this command on the master Routing Engine on an M320 router, the JUNOS Software displays the system storage information on the backup Routing Engine. On a routing matrix based on the TX Matrix or TX Matrix Plus router, if you issue this command on the TX Matrix or TX Matrix Plus router's master Routing Engine, the JUNOS Software displays all the system storage information on all the backup Routing Engines.</p> <p><b>all-chassis</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system storage statistics for all the routers in the chassis.</p> <p><b>all-lcc</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system storage statistics for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system storage statistics for all routers connected to the TX Matrix Plus router.</p> <p><b>all-members</b>—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for all members of the Virtual Chassis configuration.</p> <p><b>infrastructure <i>name</i></b>—(QFabric systems only) (Optional) Display system storage statistics for the fabric control Routing Engines or fabric manager Routing Engines.</p> <p><b>interconnect-device <i>name</i></b>—(QFabric systems only) (Optional) Display system storage statistics for the Interconnect device.</p> <p><b>lcc <i>number</i></b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system storage statistics for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system storage statistics for a specific router that is connected to the TX Matrix Plus router.</p>

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display system storage statistics for the Node group.

**scc**—(TX Matrix routers only) (Optional) Display system storage statistics for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display system storage statistics for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system storage** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation**

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)
- [show system storage partitions \(View SRX Series\)](#)

**List of Sample Output**

- [show system storage on page 212](#)
- [show system storage \(TX Matrix Plus Router\) on page 212](#)
- [show system storage \(QFX3500 Switch\) on page 214](#)
- [show system storage invoke-on all-routing-engines on page 215](#)
- [show system storage invoke-on other-routing-engine on page 216](#)

**Output Fields** Table 17 on page 212 describes the output fields for the **show system storage** command. Output fields are listed in the approximate order in which they appear.

**Table 17: show system storage Output Fields**

Field Name	Field Description
<b>Filesystem</b>	Name of the filesystem.
<b>Size</b>	Size of the filesystem.
<b>Used</b>	Amount of space used in the filesystem.
<b>Avail</b>	Amount of space available in the filesystem.
<b>Capacity</b>	Percentage of the filesystem space that is being used.
<b>Mounted on</b>	Directory in which the filesystem is mounted.

## Sample Output

### show system storage

```

user@host> show system storage
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a      77M       37M       34M      52%      /
devfs           16K       16K        0B     100%    /dev/
/dev/vn0         12M       12M        0B     100%  /packages/mnt/jbase
/dev/vn1         39M       39M        0B     100%
/packages/mnt/jkernel-7.2R1.7
/dev/vn2         12M       12M        0B     100%
/packages/mnt/jpfe-M40-7.2R1.7
/dev/vn3         2.3M      2.3M        0B     100%
/packages/mnt/jdocs-7.2R1.7
/dev/vn4         14M       14M        0B     100%
/packages/mnt/jroute-7.2R1.7
/dev/vn5         4.5M      4.5M        0B     100%
/packages/mnt/jcrypto-7.2R1.7
mfs:172         1.5G      4.0K      1.3G       0%    /tmp
/dev/ad0s1e      12M       20K        11M       0%    /config
procfs          4.0K      4.0K        0B     100%    /proc
/dev/ad1s1f      9.4G      4.9G      3.7G      57%    /var

```

### show system storage (TX Matrix Plus Router)

```

user@host> show system storage
sfc0-re0:
-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a      3.4G      178M      2.9G       6%      /
devfs           1.0K      1.0K        0B     100%    /dev
devfs           1.0K      1.0K        0B     100%    /dev/
/dev/md0         33M       33M        0B     100%  /packages/mnt/jbase
/dev/md1        216M      216M        0B     100%
/packages/mnt/jkernel-9.6-20090519.0
/dev/md2         66M       66M        0B     100%
/packages/mnt/jpfe-T-9.6-20090519.0

```

/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					
/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	1.0M	1.8G	0%	/mfs
/dev/ad0s1e	383M	82K	352M	0%	/config
procfs	4.0K	4.0K	0B	100%	/proc
/dev/ad1s1f	52G	7.5G	40G	16%	/var

lcc0-re0:

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/ad0s1a	3.4G	178M	2.9G	6%	/
devfs	1.0K	1.0K	0B	100%	/dev
devfs	1.0K	1.0K	0B	100%	/dev/
/dev/md0	33M	33M	0B	100%	/packages/mnt/jbase
/dev/md1	216M	216M	0B	100%	
/packages/mnt/jkernel-9.6-20090519.0					
/dev/md2	66M	66M	0B	100%	
/packages/mnt/jpfe-T-9.6-20090519.0					
/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					
/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	540K	1.8G	0%	/mfs
/dev/ad0s1e	383M	88K	352M	0%	/config
procfs	4.0K	4.0K	0B	100%	/proc
/dev/ad1s1f	52G	6.3G	41G	13%	/var

lcc1-re0:

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/ad0s1a	3.4G	178M	2.9G	6%	/
devfs	1.0K	1.0K	0B	100%	/dev
devfs	1.0K	1.0K	0B	100%	/dev/
/dev/md0	33M	33M	0B	100%	/packages/mnt/jbase
/dev/md1	216M	216M	0B	100%	
/packages/mnt/jkernel-9.6-20090519.0					
/dev/md2	66M	66M	0B	100%	
/packages/mnt/jpfe-T-9.6-20090519.0					
/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					
/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	540K	1.8G	0%	/mfs
/dev/ad0s1e	383M	88K	352M	0%	/config

```

procfs                4.0K      4.0K      0B      100% /proc
/dev/ad1s1f           23G      13G      7.7G      64% /var

lcc2-re0:
-----
Filesystem            Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a           3.4G      178M      2.9G        6% /
devfs                 1.0K      1.0K      0B      100% /dev
devfs                 1.0K      1.0K      0B      100% /dev/
/dev/md0              33M      33M      0B      100% /packages/mnt/jbase
/dev/md1             216M     216M      0B      100%
/packages/mnt/jkernel-9.6-20090519.0
/dev/md2              66M      66M      0B      100%
/packages/mnt/jpfe-T-9.6-20090519.0
/dev/md3             4.1M      4.1M      0B      100%
/packages/mnt/jdocs-9.6-20090519.0
/dev/md4              57M      57M      0B      100%
/packages/mnt/jroute-9.6-20090519.0
/dev/md5              15M      15M      0B      100%
/packages/mnt/jcrypto-9.6-20090519.0
/dev/md6              34M      34M      0B      100%
/packages/mnt/jpfe-common-9.6-20090519.0
/dev/md7              2.0G     10.0K      1.8G        0% /tmp
/dev/md8              2.0G      540K      1.8G        0% /mfs
/dev/ad0s1e           383M      64K      352M        0% /config
procfs                4.0K      4.0K      0B      100% /proc
/dev/ad1s1f           23G      3.7G     17G      18% /var

lcc3-re0:
-----
Filesystem            Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a           3.4G      178M      2.9G        6% /
devfs                 1.0K      1.0K      0B      100% /dev
devfs                 1.0K      1.0K      0B      100% /dev/
/dev/md0              33M      33M      0B      100% /packages/mnt/jbase
/dev/md1             216M     216M      0B      100%
/packages/mnt/jkernel-9.6-20090519.0
/dev/md2              66M      66M      0B      100%
/packages/mnt/jpfe-T-9.6-20090519.0
/dev/md3             4.1M      4.1M      0B      100%
/packages/mnt/jdocs-9.6-20090519.0
/dev/md4              57M      57M      0B      100%
/packages/mnt/jroute-9.6-20090519.0
/dev/md5              15M      15M      0B      100%
/packages/mnt/jcrypto-9.6-20090519.0
/dev/md6              34M      34M      0B      100%
/packages/mnt/jpfe-common-9.6-20090519.0
/dev/md7              2.0G     10.0K      1.8G        0% /tmp
/dev/md8              2.0G      540K      1.8G        0% /mfs
/dev/ad0s1e           383M      34K      352M        0% /config
procfs                4.0K      4.0K      0B      100% /proc
/dev/ad1s1f           23G      18G      3.5G      84% /var

```

#### show system storage (QFX3500 Switch)

```

user@switch> show system storage
Filesystem            Size      Used      Avail  Capacity  Mounted on
/dev/da0s2a           343M     192M     123M      61% /
devfs                 1.0K      1.0K      0B      100% /dev
/dev/md0             119M     119M      0B      100% /packages/mnt/jbase
/dev/md1             513M     513M      0B      100%

```

```

/packages/mnt/jkernel-qfx-11.1R1.5
/dev/md2          37M          37M          0B          100%
/packages/mnt/jpfe-qfx-e9xxx-11.1R1.5
/dev/md3          6.0M          6.0M          0B          100%
/packages/mnt/jdocs-qfx-11.1R1.5
/dev/md4          216M          216M          0B          100%
/packages/mnt/jroute-qfx-11.1R1.5
/dev/md5          59M          59M          0B          100%
/packages/mnt/jcrypto-qfx-11.1R1.5
/dev/md6          85M          85M          0B          100%
/packages/mnt/jswitch-qfx-11.1R1.5
/dev/md7          63M          8.0K          58M          0% /tmp
/dev/da0s2f       228M          14M          196M          7% /var
/dev/da0s3d       590M          3.0M          540M          1% /var/tmp
/dev/da0s3e       104M          162K          95M          0% /config
procfs           4.0K          4.0K          0B          100% /proc

```

### show system storage invoke-on all-routing-engines

```
user@host> show system storage invoke-on all-routing-engines
```

```
re0:
```

```

-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a     3.3G      440M      2.6G      14%      /
devfs           1.0K      1.0K      0B        100%     /dev
/dev/md0        76M       76M       0B        100%     /packages/mnt/jbase
/dev/md1        40M       40M       0B        100%
/packages/mnt/jkernel64-14.1-20140407.1
/dev/md2        219M      219M       0B        100%
/packages/mnt/jpfe-T-14.1-20140407.1
/dev/md3        5.4M      5.4M       0B        100%
/packages/mnt/jdocs-14.1-20140407.1
/dev/md4        116M      116M       0B        100%
/packages/mnt/jroute-14.1-20140407.1
/dev/md5        44M       44M       0B        100%
/packages/mnt/jcrypto64-14.1-20140407.1
/dev/md6        70M       70M       0B        100%
/packages/mnt/jpfe-common-14.1-20140407.1
/dev/md7        182K      182K       0B        100%
/packages/mnt/jplatform-14.1-20140407.1
/dev/md8        499M      499M       0B        100%
/packages/mnt/jruntime-14.1-20140407.1
/dev/md9        41M       41M       0B        100%
/packages/mnt/jruntime64-14.1-20140407.1
/dev/md10       12M       12M       0B        100%
/packages/mnt/py-base-i386-14.1-20140407.1
/dev/md11       3.2G      8.0K      2.9G        0% /tmp
/dev/md12       3.2G      1.1M      2.9G        0% /mfs
/dev/ad0s1e     376M      220K      346M        0% /config
procfs         4.0K      4.0K       0B        100% /proc
/dev/ad1s1f     50G       43G      3.2G       93% /var

```

```
re1:
```

```

-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a     3.3G      440M      2.6G      14%      /
devfs           1.0K      1.0K      0B        100%     /dev
/dev/md0        76M       76M       0B        100%     /packages/mnt/jbase
/dev/md1        40M       40M       0B        100%
/packages/mnt/jkernel64-14.1-20140407.1
/dev/md2        219M      219M       0B        100%

```

```

/packages/mnt/jpfe-T-14.1-20140407.1
/dev/md3          5.4M      5.4M      0B      100%
/packages/mnt/jdocs-14.1-20140407.1
/dev/md4          116M     116M      0B      100%
/packages/mnt/jroute-14.1-20140407.1
/dev/md5          44M      44M      0B      100%
/packages/mnt/jcrypto64-14.1-20140407.1
/dev/md6          70M      70M      0B      100%
/packages/mnt/jpfe-common-14.1-20140407.1
/dev/md7          182K     182K      0B      100%
/packages/mnt/jplatform-14.1-20140407.1
/dev/md8          499M     499M      0B      100%
/packages/mnt/jruntime-14.1-20140407.1
/dev/md9          41M      41M      0B      100%
/packages/mnt/jruntime64-14.1-20140407.1
/dev/md10         12M      12M      0B      100%
/packages/mnt/py-base-i386-14.1-20140407.1
/dev/md11         3.2G      8.0K      2.9G      0% /tmp
/dev/md12         3.2G     662K      2.9G      0% /mfs
/dev/ad0s1e       375M     230K     344M      0% /config
procfs           4.0K      4.0K      0B      100% /proc
/dev/ad1s1f       52G      46G      2.2G     95% /var

```

#### show system storage invoke-on other-routing-engine

```

user@host> show system storage invoke-on other-routing-engine
rel:

```

```

-----
Filesystem          Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a         3.3G      440M      2.6G      14%      /
devfs               1.0K      1.0K      0B      100%     /dev
/dev/md0            76M      76M      0B      100%     /packages/mnt/jbase
/dev/md1            40M      40M      0B      100%
/packages/mnt/jkernel64-14.1-20140407.1
/dev/md2            219M     219M      0B      100%
/packages/mnt/jpfe-T-14.1-20140407.1
/dev/md3            5.4M      5.4M      0B      100%
/packages/mnt/jdocs-14.1-20140407.1
/dev/md4            116M     116M      0B      100%
/packages/mnt/jroute-14.1-20140407.1
/dev/md5            44M      44M      0B      100%
/packages/mnt/jcrypto64-14.1-20140407.1
/dev/md6            70M      70M      0B      100%
/packages/mnt/jpfe-common-14.1-20140407.1
/dev/md7            182K     182K      0B      100%
/packages/mnt/jplatform-14.1-20140407.1
/dev/md8            499M     499M      0B      100%
/packages/mnt/jruntime-14.1-20140407.1
/dev/md9            41M      41M      0B      100%
/packages/mnt/jruntime64-14.1-20140407.1
/dev/md10           12M      12M      0B      100%
/packages/mnt/py-base-i386-14.1-20140407.1
/dev/md11           3.2G      8.0K      2.9G      0% /tmp
/dev/md12           3.2G     662K      2.9G      0% /mfs
/dev/ad0s1e         375M     230K     344M      0% /config
procfs              4.0K      4.0K      0B      100% /proc
/dev/ad1s1f         52G      46G      2.2G     95% /var

```



## show system switchover

<b>List of Syntax</b>	<a href="#">Syntax on page 217</a> <a href="#">Syntax (TX Matrix Router) on page 217</a> <a href="#">Syntax (TX Matrix Plus Router) on page 217</a> <a href="#">Syntax (MX Series Router) on page 217</a>
<b>Syntax</b>	show system switchover
<b>Syntax (TX Matrix Router)</b>	show system switchover <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system switchover <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system switchover <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.
<b>Description</b>	Display whether graceful Routing Engine switchover is configured, the state of the kernel replication (ready or synchronizing), any replication errors, and whether the primary and standby Routing Engines are using compatible versions of the kernel database.



**NOTE:** Issue the **show system switchover** command *only* on the backup Routing Engine. This command is *not* supported on the master Routing Engine, because the kernel-replication process daemon does not run on the master Routing Engine. This process runs only on the backup Routing Engine.

Beginning Junos OS Release 9.6, the **show system switchover** command has been deprecated on the master Routing Engine on all routers other than a TX Matrix (switch-card chassis) or a TX Matrix Plus (switch-fabric chassis) router.

However, in a routing matrix, if you issue the **show system switchover** command on the master Routing Engine of the TX Matrix router (or switch-card chassis), the CLI displays graceful switchover information for the master Routing Engine of the T640 routers (or line-card chassis) in the routing matrix. Likewise, if you issue the **show system switchover** command on the master Routing Engine of a TX Matrix Plus router (or switch-fabric chassis), the CLI displays output for the master Routing Engine of T1600 or T4000 routers in the routing matrix.

**Options** **all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all Routing Engines on the TX Matrix router and the T640 routers configured in the routing matrix. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all Routing Engines on the TX Matrix Plus router and the T1600 or T4000 routers configured in the routing matrix.

**all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all connected T1600 or T4000 LCCs.

Note that in this instance, packets get dropped. The LCCs perform GRES on their own chassis (GRES cannot be handled by one particular chassis for the entire router) and synchronization is not possible as the LCC plane bringup time varies for each LCC. Therefore, when there is traffic on these planes, there may be a traffic drop.

**all-members**—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on all members of the Virtual Chassis configuration.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for a specific T640 router connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for a specific router connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display graceful Routing Engines switchover information for all Routing Engines on the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**scc**—(TX Matrix router only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix Plus router.

<b>Additional Information</b>	<p>If you issue the <b>show system switchover</b> command on a TX Matrix backup Routing Engine, the command is broadcast to all the T640 backup Routing Engines that are connected to it.</p> <p>Likewise, if you issue the <b>show system switchover</b> command on a TX Matrix Plus backup Routing Engine, the command is broadcast to all the T1600 or T4000 backup Routing Engines that are connected to it.</p> <p>If you issue the <b>show system switchover</b> command on the active Routing Engine in the master router of an MX Series Virtual Chassis, the router displays an error message that graceful Routing Engine switchover (GRES) is not enabled on this member.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show system switchover (Backup Routing Engine) on page 220</a></p> <p><a href="#">show system switchover all-lcc (Routing Matrix) on page 220</a></p>
<b>Output Fields</b>	Table 18 on page 219 describes the output fields for the <b>show system switchover</b> command. Output fields are listed in the approximate order in which they appear.

Table 18: show system switchover Output Fields

Field Name	Field Description
<b>Graceful switchover</b>	<p>Display graceful Routing Engine switchover status:</p> <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates <b>graceful-switchover</b> is specified for the <b>routing-options</b> configuration command.</li> <li>• <b>Off</b>—Indicates <b>graceful-switchover</b> is not specified for the <b>routing-options</b> configuration command.</li> </ul>
<b>Configuration database</b>	<p>State of the configuration database:</p> <ul style="list-style-type: none"> <li>• <b>Ready</b>—Configuration database has synchronized.</li> <li>• <b>Synchronizing</b>—Configuration database is synchronizing. Displayed when there are updates within the last 5 seconds.</li> <li>• <b>Synchronize failed</b>—Configuration database synchronize process failed.</li> </ul>
<b>Kernel database</b>	<p>State of the kernel database:</p> <ul style="list-style-type: none"> <li>• <b>Ready</b>—Kernel database has synchronized.</li> <li>• <b>Synchronizing</b>—Kernel database is synchronizing. Displayed when there are updates within the last 5 seconds.</li> <li>• <b>Version incompatible</b>—The primary and standby Routing Engines are running incompatible kernel database versions.</li> <li>• <b>Replication error</b>—An error occurred when the state was replicated from the primary Routing Engine. Inspect <b>Steady State</b> for possible causes, or notify Juniper Networks customer support.</li> </ul>
<b>Peer state</b>	<p>Routing Engine peer state:</p> <ul style="list-style-type: none"> <li>• <b>Steady State</b>—Peer completed switchover transition.</li> <li>• <b>Peer Connected</b>—Peer in switchover transition.</li> </ul>

## Sample Output

### show system switchover (Backup Routing Engine)

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

### show system switchover all-lcc (Routing Matrix)

```
user@host> show system switchover all-lcc
```

```
lcc0-re0:
```

```
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

```
lcc2-re0:
```

```
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

## show system uptime

<b>List of Syntax</b>	<a href="#">Syntax on page 221</a> <a href="#">Syntax (EX Series Switches) on page 221</a> <a href="#">Syntax (QFX Series) on page 221</a> <a href="#">Syntax (TX Matrix Router) on page 221</a> <a href="#">Syntax (TX Matrix Plus Router) on page 221</a> <a href="#">Syntax (MX Series Router) on page 221</a>
<b>Syntax</b>	show system uptime
<b>Syntax (EX Series Switches)</b>	show system uptime <all-members> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show system uptime <director-group <i>name</i> > <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
<b>Syntax (TX Matrix Router)</b>	show system uptime <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system uptime <detail> <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system uptime <all-members> <invoke-on> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the current time and information about how long the router or switch, router or switch software, and routing protocols have been running.
<b>Options</b>	<b>none</b> —Show time since the system rebooted and processes started.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started on all the routers in the chassis.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus

router, show time since the system rebooted and processes started for all connected T1600 or T4000 LCCs.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on all members of the Virtual Chassis configuration.

**director-group *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Director group.

**infrastructure *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the fabric control Routing Engine and fabric manager Routing Engine.

**interconnect-device *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Interconnect device.

**invoke-on**—(MX Series routers only) (Optional) Display the time since the system rebooted and processes started on the master Routing Engine, backup Routing Engine, or both, on a router with two Routing Engines.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show time since the system rebooted and processes started for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Node group.

**scc**—(TX Matrix routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix router (or switch-card chassis).

**sfc number**—(TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system uptime** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation**

- [Monitoring System Process Information](#)
- [Monitoring System Properties](#)
- [10-Gigabit Ethernet LAN/WAN PIC with XFP \(T640 Router\)](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

- [show system uptime on page 224](#)
- [show system uptime all-lcc \(TX Matrix Router\) on page 224](#)
- [show system uptime all-lcc \(TX Matrix Plus Router\) on page 224](#)
- [show system uptime \(EX Series\) on page 225](#)
- [show system uptime \(QFX Series\) on page 225](#)

**Output Fields** [Table 19 on page 223](#) describes the output fields for the **show system uptime** command. Output fields are listed in the approximate order in which they appear.

**Table 19: show system uptime Output Fields**

Field Name	Field Description
<b>Current time</b>	Current system time in UTC.
<b>System booted</b>	Date and time when the Routing Engine on the router or switch was last booted and how long it has been running.
<b>Protocols started</b>	Date and time when the routing protocols were last started and how long they have been running.
<b>Last configured</b>	Date and time when a configuration was last committed. Also shows the name of the user who issued the last <b>commit</b> command.
<b>time and up</b>	Current time, in the local time zone, and how long the router or switch has been operational.
<b>users</b>	Number of users logged in to the router or switch.
<b>load averages</b>	Load averages for the last 1 minute, 5 minutes, and 15 minutes.

## Sample Output

### show system uptime

```
user@host> show system uptime
Current time:      1998-10-13 19:45:47 UTC
System booted:     1998-10-12 20:51:41 UTC (22:54:06 ago)
Protocols started: 1998-10-13 19:33:45 UTC (00:12:02 ago)
Last configured:   1998-10-13 19:33:45 UTC (00:12:02 ago) by abc
12:45PM up 22:54, 2 users, load averages: 0.07, 0.02, 0.01
```

### show system uptime all-lcc (TX Matrix Router)

```
user@host> show system uptime all-lcc
lcc0-re0:
-----
Current time: 2004-09-13 09:55:35 PDT
System booted: 2004-09-13 03:13:55 PDT (06:41:40 ago)
Last configured: 2004-09-13 03:17:48 PDT (06:37:47 ago) by root
9:55AM PDT up 6:42, 1 user, load averages: 0.02, 0.03, 0.00
lcc2-re0:
-----
Current time: 2004-09-13 09:55:35 PDT
System booted: 2004-09-12 03:23:43 PDT (1d 06:31 ago)
Last configured: 2004-09-13 03:05:36 PDT (06:49:59 ago) by root
9:55AM PDT up 1 day, 6:32, 1 user, load averages: 0.02, 0.01, 0.00
```

### show system uptime all-lcc (TX Matrix Plus Router)

```
user@host> show system uptime all-lcc
sfc0-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:33 PDT (17:44:57 ago)
Protocols started: 2009-05-24 06:40:30 PDT (17:44:00 ago)
Last configured: 2009-05-24 06:33:27 PDT (17:51:03 ago) by gregdo
12:24AM up 17:45, 2 users, load averages: 0.07, 0.05, 0.01

lcc0-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:46 PDT (17:44:44 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:47 PDT (17:43:43 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

lcc1-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:38 PDT (17:44:52 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:18 PDT (17:44:12 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

lcc2-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:48 PDT (17:44:42 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:44 PDT (17:43:46 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00
```



lcc3-re0:

-----  
Current time: 2009-05-25 00:24:30 PDT  
System booted: 2009-05-24 06:39:44 PDT (17:44:46 ago)  
error: the routing subsystem is not running  
Last configured: 2009-05-24 06:40:08 PDT (17:44:22 ago) by root  
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

#### show system uptime (EX Series)

```
user@switch> show system uptime
Current time: 2014-03-12 16:39:56 UTC
System booted: 2014-03-12 14:58:05 UTC (01:41:51 ago)
Protocols started: 2014-03-12 14:59:48 UTC (01:40:08 ago)
Last configured: 2014-03-12 14:58:58 UTC (01:40:58 ago) by root
4:39PM up 1:42, 4 users, load averages: 0.02, 0.02, 0.00
```

#### show system uptime (QFX Series)

```
user@switch> show system uptime
Current time: 2010-08-27 03:12:30 PDT
System booted: 2010-08-13 17:11:54 PDT (1w6d 10:00 ago)
Protocols started: 2010-08-13 17:13:56 PDT (1w6d 09:58 ago)
Last configured: 2010-08-26 05:54:00 PDT (21:18:30 ago) by regress
3:12AM up 13 days, 10:01, 3 users, load averages: 0.00, 0.00, 0.00
```

## show system users

<b>List of Syntax</b>	<a href="#">Syntax on page 226</a> <a href="#">Syntax (TX Matrix Router) on page 226</a> <a href="#">Syntax (TX Matrix Plus Router) on page 226</a> <a href="#">Syntax (MX Series Router) on page 226</a>
<b>Syntax</b>	show system users <no-resolve>
<b>Syntax (TX Matrix Router)</b>	show system users <all-chassis   all-lcc   lccnumber   scc> <no-resolve>
<b>Syntax (TX Matrix Plus Router)</b>	show system users <detail> <all-chassis   all-lcc   lcc number   sfc number> <no-resolve>
<b>Syntax (MX Series Router)</b>	show system users <all-members> <local> <member member-id> <no-resolve>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in JUNOS OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	List information about the users who are currently logged in to the router or switch.



**NOTE:** The **show system users** command lists the information about administrative users that are logged in to a router or switch using the CLI, J-Web, or an SSH client. The output does not list information about web users or automated users that are logged in from a remote client application using Junos XML APIs, such as NETCONF.

- Options**    **none**—List information about the users who are currently logged in to the router or switch.
- all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show users currently logged in to all the routers in the chassis.
- all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show users currently logged in to all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, show users currently logged in to all connected T1600 or T4000 LCCs.
- all-members**—(MX Series routers only) (Optional) Display users currently logged in to all members of the Virtual Chassis configuration.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show users currently logged in to a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show users currently logged in to a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display users currently logged in to the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display users currently logged in to the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**no-resolve**—(Optional) Do not attempt to resolve IP addresses to hostnames.

**scc**—(TX Matrix routers only) (Optional) Show users currently logged in to the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Show users currently logged in to the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system users** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation** • [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [show system users on page 228](#)  
[show system users lcc no-resolve \(TX Matrix, TX Matrix Plus Router\) on page 228](#)  
[show system users \(TX Matrix Plus Router\) on page 228](#)  
[show system users \(QFX Series\) on page 229](#)  
[show system users no-resolve \(QFX Series\) on page 229](#)

**Output Fields** Table 20 on page 228 describes the output fields for the **show system users** command. Output fields are listed in the approximate order in which they appear.

**Table 20: show system users Output Fields**

Field Name	Field Description
<b>time and up</b>	Current time, in the local time zone, and how long the router or switch has been operational.
<b>users</b>	Number of users logged in to the router or switch.
<b>load averages</b>	Load averages for the last 1 minute, 5 minutes, and 15 minutes.
<b>USER</b>	Username.
<b>TTY</b>	Terminal through which the user is logged in.
<b>FROM</b>	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.
<b>LOGIN@</b>	Time when the user logged in.
<b>IDLE</b>	How long the user has been idle.
<b>WHAT</b>	Processes that the user is running.

## Sample Output

### show system users

```
user@host> show system users
 7:30PM up 4 days, 2:26, 2 users, load averages: 0.07, 0.02, 0.01
USER   TTY FROM          LOGIN@  IDLE WHAT
root   d0  -             Fri05PM 4days -csh (csh)
blue   p0  level5.company.net 7:30PM  - cli
```

### show system users lcc no-resolve (TX Matrix, TX Matrix Plus Router)

```
user@host> show system users lcc 2 no-resolve

lcc2-re0:
-----
10:34AM PDT up 1 day, 7:11, 5 users, load averages: 0.03, 0.01, 0.00
USER   TTY FROM          LOGIN@  IDLE WHAT
root   d0  -             3:21AM  7:12 /bin/csh
user1  p0  scc-re0       10:15AM - telnet hostA
user1  p1  scc-re0       10:16AM - telnet hostA
user1  p2  scc-re0       10:19AM - telnet hostA
user1  p3  scc-re0       10:24AM - telnet hostA
```

### show system users (TX Matrix Plus Router)

```
user@host> show system users
sfc0-re0:
-----
1:41AM up 26 mins, 3 users, load averages: 0.08, 0.04, 0.03
```

```

USER      TTY      FROM                                LOGIN@  IDLE WHAT
user2     p0       10.209.208.123                     1:18AM  21 cli
user2     p1       172.17.29.207                      1:37AM   2 cli
user2     p2       172.17.28.19                       1:40AM   - cli

lcc0-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.00, 0.00, 0.03

lcc1-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.00, 0.02, 0.03

lcc2-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.16, 0.06, 0.02

lcc3-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.12, 0.04, 0.04

user3@aj> show system users
sfc0-re0:
-----
1:42AM up 28 mins, 4 users, load averages: 0.02, 0.03, 0.02
USER      TTY      FROM                                LOGIN@  IDLE WHAT
user3     p0       pssraj-t61.jnpr.net                1:18AM  22 cli
user3     p1       eng-shell14.juniper.net             1:37AM   - cli
user3     p2       bigpink.juniper.net                 1:40AM   - cli
user3     p3       sv-cutty-01.englab.juniper.net       1:42AM   - csh (csh)

lcc0-re0:
-----
1:42AM up 28 mins, 0 users, load averages: 0.02, 0.01, 0.03

lcc1-re0:
-----
1:42AM up 28 mins, 0 users, load averages: 0.07, 0.04, 0.03

lcc2-re0:
-----
1:42AM up 27 mins, 0 users, load averages: 0.07, 0.06, 0.02

lcc3-re0:
-----
1:42AM up 28 mins, 0 users, load averages: 0.05, 0.04, 0.04

```

#### show system users (QFX Series)

```

user@switch> show system users
USER      TTY      FROM                                LOGIN@  IDLE WHAT
tlewis    p0       172.22.18.117                      2:54AM  39 -cli (cli)
tlewis    p1       172.22.18.117                      3:01AM   - -cli (cli)
tcheng    p2       172.22.17.197                      3:08AM  11 -cli (cli)

```

#### show system users no-resolve (QFX Series)

```

user@switch> show system users no-resolve
USER      TTY      FROM                                LOGIN@  IDLE WHAT
tlewis    p0       172.22.18.117                      2:54AM  39 -cli (cli)

```

tLewis	p1	172.22.18.117	3:01AM	- -cli (cli)
tcheng	p2	172.22.17.197	3:08AM	11 -cli (cli)

## show system virtual-memory

<b>List of Syntax</b>	<a href="#">Syntax on page 231</a> <a href="#">Syntax (EX Series) on page 231</a> <a href="#">Syntax (TX Matrix Router) on page 231</a> <a href="#">Syntax (TX Matrix Plus Router) on page 231</a> <a href="#">Syntax (MX Series Router) on page 231</a> <a href="#">Syntax (QFX Series) on page 231</a>
<b>Syntax</b>	show system virtual-memory
<b>Syntax (EX Series)</b>	show system virtual-memory <all-members> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	show system virtual-memory <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system virtual-memory <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system virtual-memory <all-members> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show system virtual-memory <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the usage of Junos OS kernel memory listed first by size of allocation and then by type of usage. Use the <b>show system virtual-memory</b> command for troubleshooting with Juniper Networks Customer Support.
<b>Options</b>	<b>none</b> —Display kernel dynamic memory usage information.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display kernel dynamic memory usage information for all chassis.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display kernel dynamic memory usage information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display kernel dynamic memory usage information for all connected T1600 or T4000 LCCs.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for all members of the Virtual Chassis configuration.

**infrastructure *name***—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the fabric control Routing Engine and fabric manager Routing Engine.

**interconnect-device *name***—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the Interconnect device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display kernel dynamic memory usage information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display kernel dynamic memory usage information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the Node group.

**scc**—(TX Matrix routers only) (Optional) Display kernel dynamic memory usage information for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display kernel dynamic memory usage information for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system virtual-memory** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix



or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.



**NOTE:** The `show system virtual-memory` command with the `| display XML` pipe option now displays XML output for the command in the parent tags: `<vmstat-memstat-malloc>`, `<vmstat-memstat-zone>`, `<vmstat-sumstat>`, `<vmstat-intr>`, and `<vmstat-kernel-state>` with each child element as a separate XML tag. In Junos OS Releases 10.1 and earlier, the `| display XML` option for this command does not have an XML API element and the entire output is displayed in a single `<output>` tag element.

Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li> </ul>
List of Sample Output	<a href="#">show system virtual-memory on page 235</a> <a href="#">show system virtual-memory scc (TX Matrix Router) on page 239</a> <a href="#">show system virtual-memory sfc (TX Matrix Plus Router) on page 240</a> <a href="#">show system virtual-memory   display xml on page 243</a> <a href="#">show system virtual-memory (QFX Series) on page 266</a>
Output Fields	<p><a href="#">Table 21 on page 234</a> lists the output fields for the <code>show system virtual-memory</code> command. Output fields are listed in the approximate order in which they appear.</p>

Table 21: show system virtual-memory Output Fields

Field Name	Field Description
<b>Memory statistics by bucket size</b>	
<b>Size</b>	Memory block size (bytes). The kernel memory allocator appropriates blocks of memory whose size is exactly a power of 2.
<b>In Use</b>	Number of memory blocks of this size that are in use (bytes).
<b>Free</b>	Number of memory blocks of this size that are free (bytes).
<b>Requests</b>	Number of memory allocation requests made.
<b>HighWater</b>	Maximum value the free list can have. Once the system starts reclaiming physical memory, it continues until the free list is increased to this value.
<b>Couldfree</b>	Total number of times that the free elements for a bucket size exceed the high-water mark for that bucket size.
<b>Memory usage type by bucket size</b>	
<b>Size</b>	Memory block size (bytes).
<b>Type(s)</b>	Kernel modules that are using these memory blocks. For a definition of each type, refer to a FreeBSD book.
<b>Memory statistics by type</b>	
<b>Type</b>	Kernel module that is using dynamic memory.
<b>InUse</b>	Number of memory blocks used by this type. The number is rounded up.
<b>MemUse</b>	Amount of memory in use, in kilobytes (KB).
<b>HighUse</b>	Maximum memory ever used by this type.
<b>Limit</b>	Maximum memory that can be allocated to this type.
<b>Requests</b>	Total number of dynamic memory allocation requests this type has made.
<b>Type Limit</b>	Number of times requests were blocked for reaching the maximum limit.
<b>Kern Limit</b>	Number of times requests were blocked for the kernel map.
<b>Size(s)</b>	Memory block sizes this type is using.
<b>Memory Totals</b>	
<b>In Use</b>	Total kernel dynamic memory in use (bytes, rounded up).
<b>Free</b>	Total kernel dynamic memory free (bytes, rounded up).

Table 21: show system virtual-memory Output Fields (*continued*)

Field Name	Field Description
<b>Requests</b>	Total number of memory allocation requests.
<b>ITEM</b>	Kernel module that is using memory.
<b>Size</b>	Memory block size (bytes).
<b>Limit</b>	Maximum memory that can be allocated to this type.
<b>Used</b>	Number of memory blocks used by this type. The number is rounded up.
<b>Free</b>	Number of memory blocks available to this type.
<b>Requests</b>	Total number of memory allocation requests this type has made.
<b>interrupt</b>	Timer events and scheduling interruptions.
<b>total</b>	Total number of interruptions for each type.
<b>rate</b>	Interruption rate.
<b>Total</b>	Total for all interruptions.

## Sample Output

### show system virtual-memory

```

user@host> show system virtual-memory
Memory statistics by bucket size
Size    In Use    Free    Requests  HighWater  Couldfree
16      906      118     154876    1280       0
32      455      313     209956    640        0
64      4412     260     75380     320        20
128     3200     32      19361     160        81
256     1510     10      8844      80         4
512     446      2        5085      40         0
1K      18       2        5901      20         0
2K     1128     2        4445      10        1368
4K      185      1         456       5          0
8K       5       1        2653       5          0
16K     181      0         233       5          0
32K       2      0        1848       5          0
64K      20      0         22        5          0
128K      5      0          5        5          0
256K      2      0          2        5          0
512K      1      0          1        5          0

Memory usage type by bucket size
Size    Type(s)
16      uc_devlist, nexusdev, iftable, temp, devbuf, atexit, COS, BPF,
        DEVFS mount, DEVFS node, vnodes, mount, pcb, soname, proc-args, kld,
        MD disk, rman, ATA generic, bus, sysctl, ippool, pfestat, ifstate,

```

```

pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode
32 atkbddev, dirrem, mkdir, diradd, freefile, freefrag, indirdep,
bmsafemap, newblk, temp, devbuf, COS, vnodes, cluster_save buffer,
pcb, soname, proc-args, sigio, kld, Gzip trees, taskqueue, SWAP,
eventhandler, bus, sysctl, uidinfo, subproc, pgrp, pfestat, itable32,
ifstate, pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode, rtnexthop
64 isadev, iftable, MFS node, allocindir, allocdirect, pagedep, temp,
devbuf, lockf, COS, NULLFS hash, DEVFS name, vnodes,
cluster_save buffer, vfscache, pcb, soname, proc-args, file,
AR driver, AD driver, Gzip trees, rman, eventhandler, bus, sysctl,
subproc, pfestat, pic, ifstate, pfe_ipc, mkey, ifaddr, rtable, ipfw
128 ZONE, freeblks, inodedep, temp, devbuf, zombie, COS, DEVFS node,
vnodes, mount, vfscache, pcb, soname, proc-args, ttys, dev_t,
timecounter, kld, Gzip trees, ISOFS node, bus, uidinfo, cred,
session, pic, itable16, ifstate, pfe_ipc, rtable, ifstat, metrics,
rtnexthop, iffamily
256 iflogical, iftable, MFS node, FFS node, newblk, temp, devbuf,
NFS daemon, vnodes, proc-args, kqueue, file desc, Gzip trees, bus,
subproc, itable16, ifstate, pfe_ipc, sysctl, rtnexthop
512 UFS mount, temp, devbuf, mount, BIO buffer, ptys, ttys, AR driver,
Gzip trees, ISOFS mount, msg, ioctlops, ATA generic, bus, proc,
pfestat, lr, ifstate, pfe_ipc, rtable, ipfw, ifstat, rtnexthop
1K iftable, temp, devbuf, NQNFS Lease, kqueue, kld, AD driver,
Gzip trees, sem, MD disk, bus, ifstate, pfe_ipc, ipfw
2K uc_devlist, UFS mount, temp, devbuf, BIO buffer, pcb, AR driver,
Gzip trees, ioctlops, bus, ipfw, ifstat, rcache
4K memdesc, iftable, UFS mount, temp, devbuf, kld, Gzip trees, sem, msg
8K temp, devbuf, syncache, Gzip trees
16K indirdep, temp, devbuf, shm, msg
32K pagedep, kld, Gzip trees
64K VM pgdata, devbuf, MSDOSFS mount
128K UFS ihash, inodedep, NFS hash, kld, ISOFS mount
256K mbuf, vfscache
512K SWAP

```

Memory statistics by type					Type	Kern	
Type	InUse	MemUse	HighUse	Limit	Requests	Limit	Size(s)
isadev	13	1K	1K127753K	13	0	0	64
atkbddev	2	1K	1K127753K	2	0	0	32
uc_devlist	24	3K	3K127753K	24	0	0	16,2K
nexusdev	3	1K	1K127753K	3	0	0	16
memdesc	1	4K	4K127753K	1	0	0	4K
mbuf	1	152K	152K127753K	1	0	0	256K
iflogical	6	2K	2K127753K	6	0	0	256
iftable	17	9K	9K127753K	18	0	0	16,64,256,1K,4K
ZONE	15	2K	2K127753K	15	0	0	128
VM pgdata	1	64K	64K127753K	1	0	0	64K
UFS mount	12	26K	26K127753K	12	0	0	512,2K,4K
UFS ihash	1	128K	128K127753K	1	0	0	128K
MFS node	6	2K	3K127753K	35	0	0	64,256
FFS node	906	227K	227K127753K	1352	0	0	256
dirrem	0	0K	4K127753K	500	0	0	32
mkdir	0	0K	1K127753K	38	0	0	32
diradd	0	0K	6K127753K	521	0	0	32
freefile	0	0K	4K127753K	374	0	0	32
freeblks	0	0K	8K127753K	219	0	0	128
freefrag	0	0K	1K127753K	193	0	0	32
allocindir	0	0K	25K127753K	1518	0	0	64
indirdep	0	0K	17K127753K	76	0	0	32,16K
allocdirect	0	0K	10K127753K	760	0	0	64
bmsafemap	0	0K	1K127753K	72	0	0	32

newblk	1	1K	1K127753K	2279	0	0	32,256
inodedep	1	128K	175K127753K	2367	0	0	128,128K
pagedep	1	32K	33K127753K	47	0	0	64,32K
temp	1239	92K	96K127753K	8364	0	0	16,32,64K
devbuf	1413	5527K	5527K127753K	1535	0	0	16,32,64,128,256
lockf	38	3K	3K127753K	2906	0	0	64
atexit	1	1K	1K127753K	1	0	0	16
zombie	0	0K	2K127753K	3850	0	0	128
NFS hash	1	128K	128K127753K	1	0	0	128K
NQNFS Lease	1	1K	1K127753K	1	0	0	1K
NFS daemon	1	1K	1K127753K	1	0	0	256
syncache	1	8K	8K127753K	1	0	0	8K
COS	353	44K	44K127753K	353	0	0	16,32,64,128
BPF	189	3K	3K127753K	189	0	0	16
MSDOSFS mount	1	64K	64K127753K	1	0	0	64K
NULLFS hash	1	1K	1K127753K	1	0	0	64
DEVFS mount	2	1K	1K127753K	2	0	0	16
DEVFS name	487	31K	31K127753K	487	0	0	64
DEVFS node	471	58K	58K127753K	479	0	0	16,128
vnodes	28	7K	7K127753K	429	0	0	16,32,64,128,256
mount	15	8K	8K127753K	18	0	0	16,128,512
cluster_save buffer	0	0K	1K127753K	55	0	0	32,64
vfscache	1898	376K	376K127753K	3228	0	0	64,128,256K
BIO buffer	49	98K	398K127753K	495	0	0	512,2K
pcb	159	16K	17K127753K	399	0	0	16,32,64,128,2K
soname	82	10K	10K127753K	42847	0	0	16,32,64,128
proc-args	57	2K	3K127753K	2105	0	0	16,32,64,128,256
ptys	32	16K	16K127753K	32	0	0	512
ttys	254	33K	33K127753K	522	0	0	128,512
kqueue	5	3K	4K127753K	23	0	0	256,1K
sigio	1	1K	1K127753K	27	0	0	32
file	383	24K	24K127753K	16060	0	0	64
file desc	76	19K	20K127753K	3968	0	0	256
shm	1	12K	12K127753K	1	0	0	16K
dev_t	286	36K	36K127753K	286	0	0	128
timecounter	10	2K	2K127753K	10	0	0	128
kld	11	117K	122K127753K	34	0	0	16,32,128,1K,4K
AR driver	1	1K	3K127753K	5	0	0	64,512,2K
AD driver	2	2K	3K127753K	2755	0	0	64,1K
Gzip trees	0	0K	46K127753K	133848	0	0	32,64,128,256
ISOFS node	1136	142K	142K127753K	1189	0	0	128
ISOFS mount	9	132K	132K127753K	10	0	0	512,128K
sem	3	6K	6K127753K	3	0	0	1K,4K
MD disk	2	2K	2K127753K	2	0	0	16,1K
msg	4	25K	25K127753K	4	0	0	512,4K,16K
rman	59	4K	4K127753K	461	0	0	16,64
ioctlops	0	0K	2K127753K	992	0	0	512,2K
taskqueue	2	1K	1K127753K	2	0	0	32
SWAP	2	413K	413K127753K	2	0	0	32,512K
ATA generic	6	3K	3K127753K	6	0	0	16,512
eventhandler	17	1K	1K127753K	17	0	0	32,64
bus	340	30K	31K127753K	794	0	0	16,32,64,128,256
sysctl	0	0K	1K127753K	130262	0	0	16,32,64
uidinfo	4	1K	1K127753K	10	0	0	32,128
cred	22	3K	3K127753K	3450	0	0	128
subproc	156	10K	10K127753K	7882	0	0	32,64,256
proc	2	1K	1K127753K	2	0	0	512
session	12	2K	2K127753K	34	0	0	128
pgrp	16	1K	1K127753K	45	0	0	32
ippool	1	1K	1K127753K	1	0	0	16
pfestat	0	0K	1K127753K	47349	0	0	16,32,64,512

pic	5	1K	1K127753K	5	0	0	64,128
lr	1	1K	1K127753K	1	0	0	512
itable32	110	4K	4K127753K	110	0	0	32
itable16	161	26K	26K127753K	161	0	0	128,256
ifstate	694	159K	160K127753K	1735	0	0	16,32,64,128,1K
pfe_ipc	0	0K	1K127753K	56218	0	0	16,32,64,128,1K
mkey	250	4K	4K127753K	824	0	0	16,32,64
ifaddr	9	1K	1K127753K	9	0	0	64
sysctl	0	0K	1K127753K	30	0	0	256
rtable	49	6K	6K127753K	307	0	0	16,32,64,128,512
ifmaddr	22	1K	1K127753K	22	0	0	16,32
ipfw	23	10K	10K127753K	48	0	0	16,32,64,512,2K
ifstat	698	805K	805K127753K	698	0	0	128,512,2K
rcache	4	8K	8K127753K	4	0	0	2K
rnode	27	1K	1K127753K	285	0	0	16,32
metrics	1	1K	1K127753K	3	0	0	128
rtnexthop	57	9K	9K127753K	312	0	0	32,128,256,512
iffamily	12	2K	2K127753K	12	0	0	128

Memory Totals: In Use      Free      Requests  
                   9311K        54K        489068

ITEM	SIZE	LIMIT	USED	FREE	REQUESTS
PIPE:	192,	0,	4,	81,	4422
SWAPMETA:	160,	95814,	0,	0,	0
unpcb:	160,	0,	114,	36,	279
ripcb:	192,	25330,	5,	37,	5
syncache:	128,	15359,	0,	64,	5
tcpcb:	576,	25330,	23,	12,	32
udpcb:	192,	25330,	14,	28,	255
socket:	256,	25330,	246,	26,	819
KNOTE:	96,	0,	27,	57,	71
NFSNODE:	352,	0,	0,	0,	0
NFSMOUNT:	544,	0,	0,	0,	0
VNODE:	224,	0,	2778,	43,	2778
NAMEI:	1024,	0,	0,	8,	40725
VMSPACE:	192,	0,	57,	71,	3906
PROC:	448,	0,	73,	17,	3923
DP fakepg:	64,	0,	0,	0,	0
PV ENTRY:	28,	499566,	44530,	152053,	1525141
MAP ENTRY:	48,	0,	1439,	134,	351075
KMAP ENTRY:	48,	35645,	179,	119,	10904
MAP:	108,	0,	7,	3,	7
VM OBJECT:	92,	0,	2575,	109,	66912

792644 cpu context switches  
 9863474 device interrupts  
 286510 software interrupts  
 390851 traps  
 3596829 system calls  
   16 kernel threads created  
   3880 fork() calls  
   27 vfork() calls  
   0 rfork() calls  
   0 swap pager pageins  
   0 swap pager pages paged in  
   0 swap pager pageouts  
   0 swap pager pages paged out  
 380 vnode pager pageins  
 395 vnode pager pages paged in  
 122 vnode pager pageouts

```

1476 vnode pager pages paged out
    0 page daemon wakeups
    0 pages examined by the page daemon
101 pages reactivated
161722 copy-on-write faults
    0 copy-on-write optimized faults
84623 zero fill pages zeroed
83063 zero fill pages prezeroed
    7 intransit blocking page faults
535606 total VM faults taken
    0 pages affected by kernel thread creation
238254 pages affected by fork()
    2535 pages affected by vfork()
    0 pages affected by rfork()
283379 pages freed
    0 pages freed by daemon
190091 pages freed by exiting processes
17458 pages active
29166 pages inactive
    0 pages in VM cache
10395 pages wired down
134610 pages free
    4096 bytes per page
183419 total name lookups
    cache hits (90% pos + 7% neg) system 0% per-directory
    deletions 0%, falsehits 0%, toolong 0%

interrupt          total          rate
ata0 irq14         113338           3
mux irq7           727643          21
fxp1 irq10        1178671          34
sio0 irq4           833             0
clk irq0          3439769          99
rtc irq8          4403221         127
Total             9863475         286

Kernel direct memory map:
    4423 pages used
    4057340 pages maximum

```

*Note:* Kernel direct memory map only displays for 64 bit platform.

### show system virtual-memory scc (TX Matrix Router)

```
user@host> show system virtual-memory scc
```

```

Memory statistics by bucket size
Size  In Use  Free  Requests  HighWater  Couldfree
16    898    126   749493    1280       0
32    2018   1310  980643    640       632
64    3490   13342 935420    320       5365
...

Memory usage type by bucket size
Size  Type(s)
16    uc_devlist, COS, BPF, DEVFS mount, DEVFS node, vnodes, mount, pcb,
      soname, rman, bus, sysctl, ifstate, pfe_ipc, mkey, socket, rtable,
      ifmaddr, ipfw, rnode, iftable, temp, devbuf, atexit, proc-args, kld,
      MD disk
32    atkbddev, Gzip trees, dirrem, mkdir, diradd, freefile, freefrag,
      indirdep, bmsafemap, newblk, tseg_qent, COS, vnodes,

```

...

```

Memory statistics by type
      Type InUse MemUse HighUse Limit Requests Limit Limit Size(s)
      isadev 12 1K 1K166400K 12 0 0 64
      atkbdddev 2 1K 1K166400K 2 0 0 32
      uc_devlist 24 3K 3K166400K 24 0 0 16,2K
      ....

```

```

Memory Totals: In Use Free Requests
                6091K 1554K 2897122

```

### show system virtual-memory sfc (TX Matrix Plus Router)

```

user@host> show system virtual-memory sfc 0
sfc0-re0:

```

```

-----
      Type InUse MemUse HighUse Requests Size(s)
CAM dev queue 1 1K - 1 64
  entropy 1024 64K - 1024 64
  linker 487 6272K - 1163 16,32,64,4096,32768,131072
  USB 127 10K - 127 16,32,64,128,256,1024,2048
  lockf 46 3K - 98418 64
  USBdev 10 2K - 34 16,128,2048,16384
ifstateSLLNode 0 0K - 1096 16
  devbuf 21243 15683K - 21810
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768,65536,131072
  temp 1283 151K - 2483472
16,32,64,128,256,512,2048,4096,8192,16384,32768,65536,131072
  ip6ndp 0 0K - 4 64
  in6ifmulti 1 1K - 1 64
  in6grentry 1 1K - 1 64
  iflogical 20 5K - 29 2048
  iffamilly 45 6K - 69 32,1024,2048
  rtnexthop 266 46K - 608013 32,256,512,1024,2048,4096
  metrics 31 4K - 54 256
  rnode 212 4K - 607848 16,32
  rcache 4 8K - 4 65536
  iflist 0 0K - 6 16,64
  ifdevice 11 8K - 17 16,32768
  ifstat 424 472K - 427 512,16384,65536
  ipfw 42 23K - 145
16,32,64,128,256,512,1024,16384,32768,65536,131072
  ifmaddr 415 11K - 415 16,32
  rtable 329 28K - 608066 16,32,64,128,1024,16384
  sysctl 0 0K - 887976 16,32,64,4096,16384,32768
  ifaddr 64 5K - 70 32,64,128
  mkey 331 6K - 12528 16,128
  pfe_ipc 0 0K - 7299115
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768,65536,131072
  ifstate 1245054 70088K - 3040437
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768
  idxbucket 1 1K - 1 16
  itable16 5069 1250K - 5103 1024,4096
  itable32 157 10K - 157 64
  itable64 2 1K - 2 128
  lr 1 1K - 4 16384
  pic 37 6K - 37 64,16384
  pfestat 0 0K - 6220 32,64,128,256,131072
  gencfg 1486 424K - 2614 16,32,64,256,512,16384,32768,65536

```



```

        jsr      2      1K      -      22  16
        idl      1      4K      -      165
32, 64, 128, 256, 512, 1024, 2048, 8192, 16384, 32768, 65536, 131072
        rtmsg    0      0K      -      16  131072
        module  250     16K      -      250  64, 128
        mtx_pool 1       8K      -       1  64, 128
        DEVFS3   113    13K      -      114  256
        DEVFS1   106    24K      -      106  2048
        pgrp     15     1K      -      8600 64
        session  11     2K      -      2829 512
        proc     2      1K      -       2  16384
        subproc  296    572K     -     24689 2048, 131072
        cred     38     5K      -     619244 256
        plimit   18     4K      -     21311 2048
        uidinfo   3     1K      -       10  32, 512
        sysctl0id 2701   82K      -     2701  16, 32, 64
        sysctltmp 0      0K      -     15572  16, 32, 64, 1024
        umtx     171    11K      -      171  64
        SWAP      2    277K      -       2  64
        bus      779   125K      -     3072  16, 32, 64, 128, 32768
        bus-sc    67    62K      -     1477
16, 32, 64, 512, 1024, 2048, 8192, 16384, 65536, 131072
        devstat   8     17K      -       8  16, 131072
        eventhandler 46    2K      -       47  32, 128
        kobj      93   186K      -      111  65536
        DEVFS      8     1K      -       9  16, 64
        rman     106    7K      -      490  16, 32, 64
        sbuf       0     0K      -     28234  16, 32, 32768, 131072

```

...

lcc0-re0:

```

-----
      Type InUse MemUse HighUse Requests Size(s)
CAM dev queue    1     1K      -       1  64
      entropy  1024    64K      -     1024  64
      linker   487   6272K      -     1163  16, 32, 64, 4096, 32768, 131072
      USB     127    10K      -      127  16, 32, 64, 128, 256, 1024, 2048
      lockf    23     2K      -    169585  64
      USBdev   10     2K      -       34  16, 128, 2048, 16384
      devbuf   5128  10760K      -     5310
16, 32, 64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384, 32768, 65536, 131072
      temp    1285    151K      -     10770
16, 32, 64, 128, 256, 512, 2048, 4096, 8192, 16384, 32768, 65536, 131072
      ip6ndp    0     0K      -       4  64
      iflogical 20     5K      -      29  2048
      iffamilly 45     6K      -      69  32, 1024, 2048
      rtnexthop 189    29K      -    1211988  32, 256, 512, 1024, 2048, 4096
      metrics   11     2K      -      16  256
      rnode    135     3K      -    606391  16, 32
      rcache     4     8K      -       4  65536
      iflist     0     0K      -       6  16, 64
      ifdevice  11     8K      -      17  16, 32768
      ifstat   412   471K      -     415  512, 16384, 65536
      ipfw      42    23K      -      91
16, 32, 64, 128, 256, 512, 1024, 16384, 32768, 65536, 131072
      ifmaddr   415    11K      -      415  16, 32
      rtable    225    20K      -    606584  16, 32, 64, 128, 1024, 16384
      sysctl     0     0K      -    2302479  16, 32, 64
      ifaddr    53     4K      -      69  32, 64, 128
      mkey     133     3K      -     8974  16, 128
      pfe_ipc    0     0K      -    19035108
16, 32, 64, 128, 512, 1024, 2048, 8192, 16384, 32768, 65536, 131072

```

```

        ifstate 710270 42176K      - 9583703
16,32,64,128,256,512,1024,2048,8192,16384,32768
        idxbucket 1 1K      - 1 16
        itable16 5045 1245K      - 1825178 1024,4096
        itable32 157 10K      - 157 64
        itable64 2 1K      - 2 128
        lr 1 1K      - 4 16384
        pic 37 6K      - 37 64,16384
        pfestat 0 0K      - 1682 32,64,128,256,131072
        gencfg 1486 424K      - 2812 16,32,64,256,512,16384,32768,65536
        jsr 0 0K      - 22 16
        idl 0 0K      - 4 32768,131072
        rtsmsg 0 0K      - 3 131072
        module 250 16K      - 250 64,128
        mtx_pool 1 8K      - 1 64,128
        DEVFS3 108 12K      - 109 256
        DEVFS1 101 23K      - 101 2048
        pgrp 5 1K      - 917 64
        session 5 1K      - 917 512
        proc 2 1K      - 2 16384
        subproc 217 441K      - 4867 2048,131072
        cred 21 3K      - 48719 256
        plimit 9 2K      - 5255 2048
        uidinfo 2 1K      - 2 32,512
        sysctluid 2786 85K      - 2786 16,32,64
        sysctltmp 0 0K      - 1833 16,32,64,1024
        umtx 126 8K      - 126 64
        SWAP 2 277K      - 2 64
        bus 780 125K      - 2734 16,32,64,128,32768
        bus-sc 69 69K      - 1194
16,32,64,512,1024,2048,8192,16384,65536,131072
        devstat 8 17K      - 8 16,131072
        eventhandler 45 2K      - 46 32,128
        kobj 93 186K      - 111 65536
        DEVFS 8 1K      - 9 16,64
        rman 94 6K      - 477 16,32,64
        sbuf 0 0K      - 532 16,32,32768,131072
        NULLFS hash 1 1K      - 1 64
        taskqueue 5 1K      - 5 64
        turnstiles 127 8K      - 127 64
        Unitno 6 1K      - 44 16,64
        ioctlops 0 0K      - 1771718 16,32,64,128,8192,16384,65536,131072

        iov 0 0K      - 79425 16,64,128,256,512,1024,2048,131072
        msg 4 25K      - 4 32768,131072
        sem 4 7K      - 4 16384,32768,131072
        shm 2 13K      - 4 32768
        ttys 93 16K      - 195 512,32768
        soname 31 3K      - 389284 16,32,64,256
        pcb 101 16K      - 4374
16,32,64,128,1024,2048,4096,16384,65536
        BIO buffer 40 80K      - 750 65536
        vfscache 1 512K      - 1 65536
        cluster_save buffer 0 OK      - 55 32,64
        VFS hash 1 256K      - 1 32,64
        vnodes 1 1K      - 1 512
        mount 266 21K      - 481 16,32,64,128,256,4096,32768
        vnodemarker 0 0K      - 2497 16384
        pfs_nodes 25 3K      - 25 128
        pfs_vncache 144 5K      - 386 32
        STP 1 1K      - 1 64

```

```

      GEOM      173      15K      -      1068
16,32,64,128,256,512,2048,16384,32768,131072
      syncache      1      8K      -      1
16,32,64,128,256,512,2048,16384,32768,131072
      tlv_stat      0      0K      -      223
16,32,64,128,256,512,2048,16384,32768,131072
      NFS daemon      1      8K      -      1
16,32,64,128,256,512,2048,16384,32768,131072
      p1003.1b      1      1K      -      1 16
      MD disk      9      18K      -      9 65536
      ata_generic      2      2K      -      25 16,16384,32768
      ISOFS mount      7      1K      -      13 512
      ISOFS node 1439      135K      -      1453 128
      CAM SIM      1      1K      -      1 64
      CAM XPT      6      1K      -      9 16,64,16384
      CAM periph      1      1K      -      1 128
      ad_driver      2      1K      -      2 256
      pagedep      1      64K      -      105 64
      inodedep      1      256K      -      552 256
      newblk      1      1K      -      327 64,4096
      bmsafemap      0      0K      -      19 64
      allocdirect      0      0K      -      326 128
      freefrag      0      0K      -      31 32
      freeblks      0      0K      -      103 2048
      freefile      0      0K      -      175 32
      diradd      0      0K      -      590 64
      mkdir      0      0K      -      166 32
      dirrem      0      0K      -      382 32
      savedino      0      0K      -      283 512
      UFS mount      15      36K      -      15 2048,65536,131072
      ata_dma      6      1K      -      6 256
      UMAHash      1      4K      -      5 4096,16384,32768,65536,131072
      cdev      26      3K      -      26 256
      file desc 111      25K      -      5199 16,1024,2048,16384
      VM pgdata      2      65K      -      2 64
      sigio      1      1K      -      27 32
      kenv      30      5K      -      33 16,32,64,131072
      atkbddev      2      1K      -      2 32
      kqueue      0      0K      -      88 1024,4096,32768
      proc-args      28      2K      -      3970 32,64,128,256,512,1024
      isadev      23      2K      -      23 64
      zombie      1      1K      -      4651 128
      ithread      92      7K      -      92 16,64,256
      legacydrv      3      1K      -      3 16
      memdesc      1      4K      -      1 131072
      nexusdev      2      1K      -      2 16
      CAM queue      3      1K      -      3 16
      KTRACE      100      10K      -      100 128
      kbdmux      5      9K      -      5 128,2048,65536,131072
ITEM      SIZE      LIMIT      USED      FREE      REQUESTS
UMA Kegs:      136,      0,      71,      1,      71
...
```

### show system virtual-memory | display xml

```

user@host> show system virtual-memory | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.2R1/junos">
  <system-virtual-memory-information>
    <vmstat-memstat-malloc>
      <memstat-name>CAM dev queue</memstat-name>
      <inuse>1</inuse>
    </vmstat-memstat-malloc>
  </system-virtual-memory-information>
</rpc-reply>
```

```
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>entropy</memstat-name>
<inuse>1024</inuse>
<memuse>64</memuse>
<high-use>--</high-use>
<memstat-req>1024</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>linker</memstat-name>
<inuse>481</inuse>
<memuse>1871</memuse>
<high-use>--</high-use>
<memstat-req>1145</memstat-req>
<memstat-size>16,32,64,4096,32768,131072</memstat-size>
<memstat-name>lockf</memstat-name>
<inuse>56</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>5998</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>devbuf</memstat-name>
<inuse>2094</inuse>
<memuse>3877</memuse>
<high-use>--</high-use>
<memstat-req>2099</memstat-req>

<memstat-size>16,32,64,128,512,1024,4096,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>temp</memstat-name>
<inuse>21</inuse>
<memuse>66</memuse>
<high-use>--</high-use>
<memstat-req>3127</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,4096,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>ip6ndp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>in6ifmulti</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>in6grenty</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>iflogical</memstat-name>
<inuse>13</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
```

```

<memstat-size>64,2048</memstat-size>
<memstat-name>iffamily</memstat-name>
<inuse>28</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>28</memstat-req>
<memstat-size>32,1024,2048</memstat-size>
<memstat-name>rtnexthop</memstat-name>
<inuse>127</inuse>
<memuse>18</memuse>
<high-use>--</high-use>
<memstat-req>129</memstat-req>
<memstat-size>32,256,512,1024,2048,4096</memstat-size>
<memstat-name>metrics</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>inifmulti</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>ingrentry</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>rnode</memstat-name>
<inuse>68</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>76</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rcache</memstat-name>
<inuse>4</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ifdevice</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>ifstat</memstat-name>
<inuse>40</inuse>
<memuse>22</memuse>
<high-use>--</high-use>
<memstat-req>40</memstat-req>
<memstat-size>512,16384,32768</memstat-size>
<memstat-name>ipfw</memstat-name>
<inuse>42</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>91</memstat-req>

```

```
<memstat-size>16,32,64,128,256,512,1024,16384,32768,65536,131072</memstat-size>
  <memstat-name>ifmaddr</memstat-name>
  <inuse>103</inuse>
  <memuse>3</memuse>
  <high-use>--</high-use>
  <memstat-req>103</memstat-req>
  <memstat-size>16,32</memstat-size>
  <memstat-name>rtable</memstat-name>
  <inuse>129</inuse>
  <memuse>14</memuse>
  <high-use>--</high-use>
  <memstat-req>139</memstat-req>
  <memstat-size>16,32,64,128,1024,16384</memstat-size>
  <memstat-name>sysctl</memstat-name>
  <inuse>0</inuse>
  <memuse>0</memuse>
  <high-use>--</high-use>
  <memstat-req>14847</memstat-req>
  <memstat-size>16,32,64,4096,16384,32768</memstat-size>
  <memstat-name>ifaddr</memstat-name>
  <inuse>29</inuse>
  <memuse>3</memuse>
  <high-use>--</high-use>
  <memstat-req>29</memstat-req>
  <memstat-size>64,128</memstat-size>
  <memstat-name>mkey</memstat-name>
  <inuse>345</inuse>
  <memuse>6</memuse>
  <high-use>--</high-use>
  <memstat-req>2527</memstat-req>
  <memstat-size>16,128</memstat-size>
  <memstat-name>pfe_ipc</memstat-name>
  <inuse>0</inuse>
  <memuse>0</memuse>
  <high-use>--</high-use>
  <memstat-req>1422</memstat-req>

<memstat-size>16,32,64,128,512,1024,2048,8192,16384,32768,65536,131072</memstat-size>

  <memstat-name>ifstate</memstat-name>
  <inuse>594</inuse>
  <memuse>51</memuse>
  <high-use>--</high-use>
  <memstat-req>655</memstat-req>

<memstat-size>16,32,64,128,256,1024,2048,4096,16384,32768</memstat-size>
  <memstat-name>itable16</memstat-name>
  <inuse>276</inuse>
  <memuse>52</memuse>
  <high-use>--</high-use>
  <memstat-req>294</memstat-req>
  <memstat-size>1024,4096</memstat-size>
  <memstat-name>itable32</memstat-name>
  <inuse>160</inuse>
  <memuse>10</memuse>
  <high-use>--</high-use>
  <memstat-req>160</memstat-req>
  <memstat-size>64</memstat-size>
  <memstat-name>itable64</memstat-name>
  <inuse>2</inuse>
  <memuse>1</memuse>
```

```

<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>lr</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pic</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64,512</memstat-size>
<memstat-name>pfestat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>162</memstat-req>
<memstat-size>16,32,128,256,16384</memstat-size>
<memstat-name>gencfg</memstat-name>
<inuse>224</inuse>
<memuse>56</memuse>
<high-use>--</high-use>
<memstat-req>540</memstat-req>
<memstat-size>16,32,64,256,512,32768,65536</memstat-size>
<memstat-name>jsr</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>idl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>16,32,64,128,256,4096,16384,32768,131072</memstat-size>

<memstat-name>rtsmsg</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>module</memstat-name>
<inuse>249</inuse>
<memuse>16</memuse>
<high-use>--</high-use>
<memstat-req>249</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mtx_pool</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>DEVFS3</memstat-name>
<inuse>109</inuse>
<memuse>12</memuse>

```

```
<high-use>--</high-use>
<memstat-req>117</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>DEVFS1</memstat-name>
<inuse>102</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>109</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>pgrp</memstat-name>
<inuse>12</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>session</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>proc</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>subproc</memstat-name>
<inuse>244</inuse>
<memuse>496</memuse>
<high-use>--</high-use>
<memstat-req>1522</memstat-req>
<memstat-size>2048,131072</memstat-size>
<memstat-name>cred</memstat-name>
<inuse>30</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>11409</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>plimit</memstat-name>
<inuse>17</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>133</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>uidinfo</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>32,512</memstat-size>
<memstat-name>sysctlpid</memstat-name>
<inuse>1117</inuse>
<memuse>34</memuse>
<high-use>--</high-use>
<memstat-req>1117</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sysctltmp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
```



```

<memstat-req>743</memstat-req>
<memstat-size>16,32,64,1024</memstat-size>
<memstat-name>umtx</memstat-name>
<inuse>144</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>144</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>SWAP</memstat-name>
<inuse>2</inuse>
<memuse>209</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>bus</memstat-name>
<inuse>496</inuse>
<memuse>55</memuse>
<high-use>--</high-use>
<memstat-req>1196</memstat-req>
<memstat-size>16,32,64,128,32768</memstat-size>
<memstat-name>bus-sc</memstat-name>
<inuse>23</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>335</memstat-req>

<memstat-size>16,32,64,512,1024,2048,8192,16384,65536,131072</memstat-size>
<memstat-name>devstat</memstat-name>
<inuse>10</inuse>
<memuse>21</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>16,131072</memstat-size>
<memstat-name>eventhandler</memstat-name>
<inuse>35</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>36</memstat-req>
<memstat-size>32,128</memstat-size>
<memstat-name>kobj</memstat-name>
<inuse>93</inuse>
<memuse>186</memuse>
<high-use>--</high-use>
<memstat-req>111</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>DEVFS</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>rman</memstat-name>
<inuse>71</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>433</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sbuf</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>

```

```
<memstat-req>522</memstat-req>
<memstat-size>16,32,32768,131072</memstat-size>
<memstat-name>NULLFS hash</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>taskqueue</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>turnstiles</memstat-name>
<inuse>145</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>145</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>Unitno</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>44</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>iocltops</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>27622</memstat-req>
<memstat-size>16,64,8192,16384,131072</memstat-size>
<memstat-name>iov</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>18578</memstat-req>
<memstat-size>16,64,128,256,512,1024,2048,131072</memstat-size>
<memstat-name>msg</memstat-name>
<inuse>4</inuse>
<memuse>25</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32768,131072</memstat-size>
<memstat-name>sem</memstat-name>
<inuse>4</inuse>
<memuse>7</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16384,32768,131072</memstat-size>
<memstat-name>shm</memstat-name>
<inuse>9</inuse>
<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>14</memstat-req>
<memstat-size>32768</memstat-size>
<memstat-name>ttys</memstat-name>
<inuse>321</inuse>
<memuse>61</memuse>
<high-use>--</high-use>
<memstat-req>528</memstat-req>
```

```

<memstat-size>512,32768</memstat-size>
<memstat-name>ptys</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>mbuf_tag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>23383</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>soname</memstat-name>
<inuse>115</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>24712</memstat-req>
<memstat-size>16,32,64,256</memstat-size>
<memstat-name>pcb</memstat-name>
<inuse>216</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>484</memstat-req>

<memstat-size>16,32,64,128,1024,2048,4096,16384,32768,65536</memstat-size>
<memstat-name>BIO buffer</memstat-name>
<inuse>43</inuse>
<memuse>86</memuse>
<high-use>--</high-use>
<memstat-req>405</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>vfscache</memstat-name>
<inuse>1</inuse>
<memuse>256</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>cluster_save buffer</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>VFS hash</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>vnodes</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>mount</memstat-name>
<inuse>290</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>535</memstat-req>

```

```
<memstat-size>16,32,64,128,256,4096,32768</memstat-size>
<memstat-name>vnodemarker</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>498</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pfs_nodes</memstat-name>
<inuse>25</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>25</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>pfs_vncache</memstat-name>
<inuse>27</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>53</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>STP</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>GEOM</memstat-name>
<inuse>146</inuse>
<memuse>11</memuse>
<high-use>--</high-use>
<memstat-req>1042</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>syncache</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>tlv_stat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>8</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>NFS_daemon</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>p1003.1b</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>MD_disk</memstat-name>
<inuse>10</inuse>
```

```

<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ata_generic</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>16,16384,32768</memstat-size>
<memstat-name>ISOFs mount</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>ISOFs node</memstat-name>
<inuse>1440</inuse>
<memuse>135</memuse>
<high-use>--</high-use>
<memstat-req>1457</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>CAM SIM</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>CAM XPT</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64,16384</memstat-size>
<memstat-name>CAM periph</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ad_driver</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>pagedep</memstat-name>
<inuse>1</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>106</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>inodedep</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>464</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>newblk</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>-</high-use>
<memstat-req>336</memstat-req>
<memstat-size>64,4096</memstat-size>
<memstat-name>bmsafemap</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>63</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>allocdirect</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>320</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>indirdep</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>17</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>allocindir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>15</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>freefrag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>12</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>freeblks</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>40</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>freefile</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>101</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>diradd</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>465</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>mkdir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>136</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>dirrem</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
```

```

<memstat-req>168</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>newdirblk</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>saveditino</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>157</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>UFS mount</memstat-name>
<inuse>15</inuse>
<memuse>36</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>2048,65536,131072</memstat-size>
<memstat-name>ata_dma</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>UMAHash</memstat-name>
<inuse>1</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>4096,16384,32768,65536</memstat-size>
<memstat-name>cdev</memstat-name>
<inuse>22</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>22</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>file desc</memstat-name>
<inuse>141</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>1583</memstat-req>
<memstat-size>16,1024,2048,16384</memstat-size>
<memstat-name>VM pgdata</memstat-name>
<inuse>2</inuse>
<memuse>65</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>sigio</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>20</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kenv</memstat-name>
<inuse>24</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>27</memstat-req>

```

```
<memstat-size>16,32,64,131072</memstat-size>
<memstat-name>atkbddev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kqueue</memstat-name>
<inuse>15</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>19</memstat-req>
<memstat-size>1024,4096,32768</memstat-size>
<memstat-name>proc-args</memstat-name>
<inuse>57</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>1001</memstat-req>
<memstat-size>16,32,64,128,256,512,1024</memstat-size>
<memstat-name>isadev</memstat-name>
<inuse>21</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>zombie</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1278</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ithread</memstat-name>
<inuse>69</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>69</memstat-req>
<memstat-size>16,64,256</memstat-size>
<memstat-name>legacydrv</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>memdesc</memstat-name>
<inuse>1</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>nexusdev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>CAM queue</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>16</memstat-size>
```



```

    <memstat-name>$PIR</memstat-name>
    <inuse>4</inuse>
    <memuse>1</memuse>
    <high-use>--</high-use>
    <memstat-req>4</memstat-req>
    <memstat-size>32</memstat-size>
    <memstat-name>KTRACE</memstat-name>
    <inuse>100</inuse>
    <memuse>10</memuse>
    <high-use>--</high-use>
    <memstat-req>100</memstat-req>
    <memstat-size>128</memstat-size>
    <memstat-name>kbdmux</memstat-name>
    <inuse>5</inuse>
    <memuse>9</memuse>
    <high-use>--</high-use>
    <memstat-req>5</memstat-req>
    <memstat-size>128,2048,65536,131072</memstat-size>
  </vmstat-memstat-malloc>
  <vmstat-memstat-zone>
    <zone-name>UMA Kegs:</zone-name>
    <zone-size>136</zone-size>
    <count-limit>0</count-limit>
    <used>71</used>
    <free>1</free>
    <zone-req>71</zone-req>
    <zone-name>UMA Zones:</zone-name>
    <zone-size>120</zone-size>
    <count-limit>0</count-limit>
    <used>71</used>
    <free>19</free>
    <zone-req>71</zone-req>
    <zone-name>UMA Slabs:</zone-name>
    <zone-size>64</zone-size>
    <count-limit>0</count-limit>
    <used>490</used>
    <free>41</free>
    <zone-req>579</zone-req>
    <zone-name>UMA RCntSlabs:</zone-name>
    <zone-size>104</zone-size>
    <count-limit>0</count-limit>
    <used>276</used>
    <free>20</free>
    <zone-req>276</zone-req>
    <zone-name>UMA Hash:</zone-name>
    <zone-size>128</zone-size>
    <count-limit>0</count-limit>
    <used>4</used>
    <free>26</free>
    <zone-req>5</zone-req>
    <zone-name>16 Bucket:</zone-name>
    <zone-size>76</zone-size>
    <count-limit>0</count-limit>
    <used>30</used>
    <free>20</free>
    <zone-req>30</zone-req>
    <zone-name>32 Bucket:</zone-name>
    <zone-size>140</zone-size>
    <count-limit>0</count-limit>
    <used>33</used>
    <free>23</free>

```

```
<zone-req>33</zone-req>
<zone-name>64 Bucket:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>9</free>
<zone-req>33</zone-req>
<zone-name>128 Bucket:</zone-name>
<zone-size>524</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>0</free>
<zone-req>49</zone-req>
<zone-name>VM OBJECT:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>2111</used>
<free>79</free>
<zone-req>25214</zone-req>
<zone-name>MAP:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>7</used>
<free>41</free>
<zone-req>7</zone-req>
<zone-name>KMAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>35336</count-limit>
<used>19</used>
<free>149</free>
<zone-req>2397</zone-req>
<zone-name>MAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2031</used>
<free>153</free>
<zone-req>62417</zone-req>
<zone-name>PV ENTRY:</zone-name>
<zone-size>24</zone-size>
<count-limit>509095</count-limit>
<used>57177</used>
<free>6333</free>
<zone-req>1033683</zone-req>
<zone-name>DP fakepg:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mt_zone:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>238</used>
<free>57</free>
<zone-req>238</zone-req>
<zone-name>16:</zone-name>
<zone-size>16</zone-size>
<count-limit>0</count-limit>
<used>2114</used>
<free>119</free>
<zone-req>80515</zone-req>
```

```
<zone-name>32:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>1335</used>
<free>134</free>
<zone-req>10259</zone-req>
<zone-name>64:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>3529</used>
<free>129</free>
<zone-req>29110</zone-req>
<zone-name>96:</zone-name>
<zone-size>96</zone-size>
<count-limit>0</count-limit>
<used>2062</used>
<free>58</free>
<zone-req>4365</zone-req>
<zone-name>112:</zone-name>
<zone-size>112</zone-size>
<count-limit>0</count-limit>
<used>361</used>
<free>164</free>
<zone-req>24613</zone-req>
<zone-name>128:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>359</used>
<free>61</free>
<zone-req>942</zone-req>
<zone-name>160:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>364</used>
<free>44</free>
<zone-req>577</zone-req>
<zone-name>224:</zone-name>
<zone-size>224</zone-size>
<count-limit>0</count-limit>
<used>422</used>
<free>20</free>
<zone-req>1950</zone-req>
<zone-name>256:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>204</used>
<free>36</free>
<zone-req>1225</zone-req>
<zone-name>288:</zone-name>
<zone-size>288</zone-size>
<count-limit>0</count-limit>
<used>2</used>
<free>24</free>
<zone-req>10</zone-req>
<zone-name>512:</zone-name>
<zone-size>512</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>7</free>
<zone-req>911</zone-req>
<zone-name>1024:</zone-name>
```

```
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>213</used>
<free>11</free>
<zone-req>1076</zone-req>
<zone-name>2048:</zone-name>
<zone-size>2048</zone-size>
<count-limit>0</count-limit>
<used>199</used>
<free>113</free>
<zone-req>640</zone-req>
<zone-name>4096:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>144</used>
<free>7</free>
<zone-req>2249</zone-req>
<zone-name>Files:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>665</used>
<free>77</free>
<zone-req>16457</zone-req>
<zone-name>MAC labels:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>3998</used>
<free>227</free>
<zone-req>21947</zone-req>
<zone-name>PROC:</zone-name>
<zone-size>544</zone-size>
<count-limit>0</count-limit>
<used>116</used>
<free>10</free>
<zone-req>1394</zone-req>
<zone-name>THREAD:</zone-name>
<zone-size>416</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>17</free>
<zone-req>131</zone-req>
<zone-name>KSEGRP:</zone-name>
<zone-size>88</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>73</free>
<zone-req>131</zone-req>
<zone-name>UPCALL:</zone-name>
<zone-size>44</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>SLEEPQUEUE:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>145</used>
<free>194</free>
<zone-req>145</zone-req>
<zone-name>VMSPACE:</zone-name>
<zone-size>268</zone-size>
```

```

<count-limit>0</count-limit>
<used>57</used>
<free>13</free>
<zone-req>1335</zone-req>
<zone-name>mbuf_packet:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>256</used>
<free>128</free>
<zone-req>49791</zone-req>
<zone-name>mbuf:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>50</used>
<free>466</free>
<zone-req>105183</zone-req>
<zone-name>mbuf_cluster:</zone-name>
<zone-size>2048</zone-size>
<count-limit>25190</count-limit>
<used>387</used>
<free>165</free>
<zone-req>5976</zone-req>
<zone-name>mbuf_jumbo_pagesize:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_9k:</zone-name>
<zone-size>9216</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_16k:</zone-name>
<zone-size>16384</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ACL UMA zone:</zone-name>
<zone-size>388</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>g_bio:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>174</free>
<zone-req>69750</zone-req>
<zone-name>ata_request:</zone-name>
<zone-size>200</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>57</free>
<zone-req>5030</zone-req>
<zone-name>ata_composite:</zone-name>
<zone-size>192</zone-size>
<count-limit>0</count-limit>

```

```
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>GENCFG:</zone-name>
<zone-size>72</zone-size>
<count-limit>1000004</count-limit>
<used>57</used>
<free>102</free>
<zone-req>57</zone-req>
<zone-name>VNODE:</zone-name>
<zone-size>292</zone-size>
<count-limit>0</count-limit>
<used>2718</used>
<free>25</free>
<zone-req>2922</zone-req>
<zone-name>VNODEPOLL:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>S VFS Cache:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2500</used>
<free>76</free>
<zone-req>3824</zone-req>
<zone-name>L VFS Cache:</zone-name>
<zone-size>291</zone-size>
<count-limit>0</count-limit>
<used>51</used>
<free>14</free>
<zone-req>63</zone-req>
<zone-name>NAMEI:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>8</free>
<zone-req>53330</zone-req>
<zone-name>NFSMOUNT:</zone-name>
<zone-size>480</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>NFSNODE:</zone-name>
<zone-size>460</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>PIPE:</zone-name>
<zone-size>404</zone-size>
<count-limit>0</count-limit>
<used>27</used>
<free>9</free>
<zone-req>717</zone-req>
<zone-name>KNOTE:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>42</used>
```

```

<free>64</free>
<zone-req>3311</zone-req>
<zone-name>socket:</zone-name>
<zone-size>412</zone-size>
<count-limit>25191</count-limit>
<used>343</used>
<free>8</free>
<zone-req>2524</zone-req>
<zone-name>unpcb:</zone-name>
<zone-size>140</zone-size>
<count-limit>25200</count-limit>
<used>170</used>
<free>26</free>
<zone-req>2157</zone-req>
<zone-name>ipq:</zone-name>
<zone-size>52</zone-size>
<count-limit>216</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>udpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>19</used>
<free>32</free>
<zone-req>31</zone-req>
<zone-name>inpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>40</used>
<free>28</free>
<zone-req>105</zone-req>
<zone-name>tcpcb:</zone-name>
<zone-size>520</zone-size>
<count-limit>25193</count-limit>
<used>40</used>
<free>16</free>
<zone-req>105</zone-req>
<zone-name>tcptw:</zone-name>
<zone-size>56</zone-size>
<count-limit>5092</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>syncache:</zone-name>
<zone-size>128</zone-size>
<count-limit>15360</count-limit>
<used>0</used>
<free>60</free>
<zone-req>55</zone-req>
<zone-name>tcpreass:</zone-name>
<zone-size>20</zone-size>
<count-limit>1690</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>sackhole:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>

```

```

    <zone-req>0</zone-req>
    <zone-name>ripcb:</zone-name>
    <zone-size>232</zone-size>
    <count-limit>25194</count-limit>
    <used>5</used>
    <free>29</free>
    <zone-req>5</zone-req>
    <zone-name>SWAPMETA:</zone-name>
    <zone-size>276</zone-size>
    <count-limit>94948</count-limit>
    <used>0</used>
    <free>0</free>
    <zone-req>0</zone-req>
    <zone-name>FFS inode:</zone-name>
    <zone-size>132</zone-size>
    <count-limit>0</count-limit>
    <used>1146</used>
    <free>72</free>
    <zone-req>1306</zone-req>
    <zone-name>FFS1 dinode:</zone-name>
    <zone-size>128</zone-size>
    <count-limit>0</count-limit>
    <used>1146</used>
    <free>24</free>
    <zone-req>1306</zone-req>
    <zone-name>FFS2 dinode:</zone-name>
    <zone-size>256</zone-size>
    <count-limit>0</count-limit>
    <used>0</used>
    <free>0</free>
    <zone-req>0</zone-req>
  </vmstat-memstat-zone>
<vmstat-sumstat>
  <cpu-context-switch>934906</cpu-context-switch>
  <dev-intr>1707986</dev-intr>
  <soft-intr>33819</soft-intr>
  <traps>203604</traps>
  <sys-calls>1200636</sys-calls>
  <kernel-thrds>60</kernel-thrds>
  <fork-calls>1313</fork-calls>
  <vfork-calls>21</vfork-calls>
  <rfork-calls>0</rfork-calls>
  <swap-pageins>0</swap-pageins>
  <swap-pagedin>0</swap-pagedin>
  <swap-pageouts>0</swap-pageouts>
  <swap-pagedout>0</swap-pagedout>
  <vnode-pageins>23094</vnode-pageins>
  <vnode-pagedin>23119</vnode-pagedin>
  <vnode-pageouts>226</vnode-pageouts>
  <vnode-pagedout>3143</vnode-pagedout>
  <page-daemon-wakeup>0</page-daemon-wakeup>
  <page-daemon-examined-pages>0</page-daemon-examined-pages>
  <pages-reactivated>8821</pages-reactivated>
  <copy-on-write-faults>48364</copy-on-write-faults>
  <copy-on-write-optimized-faults>31</copy-on-write-optimized-faults>
  <zero-fill-pages-zeroed>74665</zero-fill-pages-zeroed>
  <zero-fill-pages-prezeroed>70061</zero-fill-pages-prezeroed>
  <transit-blocking-page-faults>85</transit-blocking-page-faults>
  <total-vm-faults>191824</total-vm-faults>

<pages-affected-by-kernel-thrd-creat>0</pages-affected-by-kernel-thrd-creat>

```



```

    <pages-affected-by-fork>95343</pages-affected-by-fork>
    <pages-affected-by-vfork>3526</pages-affected-by-vfork>
    <pages-affected-by-rfork>0</pages-affected-by-rfork>
    <pages-freed>221502</pages-freed>
    <pages-freed-by-daemon>0</pages-freed-by-daemon>
    <pages-freed-by-exiting-proc>75630</pages-freed-by-exiting-proc>
    <pages-active>45826</pages-active>
    <pages-inactive>13227</pages-inactive>
    <pages-in-vm-cache>49278</pages-in-vm-cache>
    <pages-wired-down>10640</pages-wired-down>
    <pages-free>70706</pages-free>
    <bytes-per-page>4096</bytes-per-page>
    <swap-pages-used>0</swap-pages-used>
    <peak-swap-pages-used>0</peak-swap-pages-used>
    <total-name-lookups>214496</total-name-lookups>
    <positive-cache-hits>92</positive-cache-hits>
    <negative-cache-hits>5</negative-cache-hits>
    <pass2>0</pass2>
    <cache-deletions>0</cache-deletions>
    <cache-falsehits>0</cache-falsehits>
    <toolong>0</toolong>
</vmstat-sumstat>
<vmstat-intr>
  <intr-name>irq0: clk          </intr-name>
  <intr-cnt>1243455</intr-cnt>
  <intr-rate>999</intr-rate>
  <intr-name>irq4: sio0        </intr-name>
  <intr-cnt>1140</intr-cnt>
  <intr-rate>0</intr-rate>
  <intr-name>irq8: rtc         </intr-name>
  <intr-cnt>159164</intr-cnt>
  <intr-rate>127</intr-rate>
  <intr-name>irq9: cbb1 fxp0   </intr-name>
  <intr-cnt>28490</intr-cnt>
  <intr-rate>22</intr-rate>
  <intr-name>irq10: fxp1       </intr-name>
  <intr-cnt>20593</intr-cnt>
  <intr-rate>16</intr-rate>
  <intr-name>irq14: ata0       </intr-name>
  <intr-cnt>5031</intr-cnt>
  <intr-rate>4</intr-rate>
  <intr-name>Total</intr-name>
  <intr-cnt>1457873</intr-cnt>
  <intr-rate>1171</intr-rate>
</vmstat-intr>
<vm-kernel-state>
  <vm-kmem-map-free>248524800</vm-kmem-map-free>
</vm-kernel-state>
<kernel-direct-mm-size-information>
  <vm-directmm-size-used>4644</vm-directmm-size-used>
  <vm-directmm-size-max>4057334</vm-directmm-size-max>
</kernel-direct-mm-size-information>
</system-virtual-memory-information>
<cli>
  <banner></banner>
</cli>
</rpc-reply>

```

Note: <kernel-direct-mm-size-information> only displays for 64 bit platform.

## show system virtual-memory (QFX Series)

```

user@switch> show system virtual-memory | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/11.1R1/junos">
  <system-virtual-memory-information>
    <vmstat-memstat-malloc>
      <memstat-name>CAM dev queue</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>
      <high-use>-</high-use>
      <memstat-req>1</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>entropy</memstat-name>
      <inuse>1024</inuse>
      <memuse>64</memuse>
      <high-use>-</high-use>
      <memstat-req>1024</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>linker</memstat-name>
      <inuse>481</inuse>
      <memuse>1871</memuse>
      <high-use>-</high-use>
      <memstat-req>1145</memstat-req>
      <memstat-size>16, 32, 64, 4096, 32768, 131072</memstat-size>
      <memstat-name>lockf</memstat-name>
      <inuse>56</inuse>
      <memuse>4</memuse>
      <high-use>-</high-use>
      <memstat-req>5998</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>devbuf</memstat-name>
      <inuse>2094</inuse>
      <memuse>3877</memuse>
      <high-use>-</high-use>
      <memstat-req>2099</memstat-req>

      <memstat-size>16, 32, 64, 128, 512, 1024, 4096, 8192, 16384, 32768, 65536, 131072</memstat-size>

      <memstat-name>temp</memstat-name>
      <inuse>21</inuse>
      <memuse>66</memuse>
      <high-use>-</high-use>
      <memstat-req>3127</memstat-req>

      <memstat-size>16, 32, 64, 128, 256, 512, 2048, 4096, 8192, 16384, 32768, 65536, 131072</memstat-size>

      <memstat-name>ip6ndp</memstat-name>
      <inuse>0</inuse>
      <memuse>0</memuse>
      <high-use>-</high-use>
      <memstat-req>4</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>in6ifmulti</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>
      <high-use>-</high-use>
      <memstat-req>1</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>in6grentry</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>

```

```

<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>iflogical</memstat-name>
<inuse>13</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>64,2048</memstat-size>
<memstat-name>iffamily</memstat-name>
<inuse>28</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>28</memstat-req>
<memstat-size>32,1024,2048</memstat-size>
<memstat-name>rtnexthop</memstat-name>
<inuse>127</inuse>
<memuse>18</memuse>
<high-use>--</high-use>
<memstat-req>129</memstat-req>
<memstat-size>32,256,512,1024,2048,4096</memstat-size>
<memstat-name>metrics</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>inifmulti</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>ingrentry</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>rnode</memstat-name>
<inuse>68</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>76</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rcache</memstat-name>
<inuse>4</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ifdevice</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>ifstat</memstat-name>
<inuse>40</inuse>
<memuse>22</memuse>
<high-use>--</high-use>

```

```
<memstat-req>40</memstat-req>
<memstat-size>512,16384,32768</memstat-size>
<memstat-name>ipfw</memstat-name>
<inuse>42</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>91</memstat-req>

<memstat-size>16,32,64,128,256,512,1024,16384,32768,65536,131072</memstat-size>
<memstat-name>ifmaddr</memstat-name>
<inuse>103</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>103</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rtable</memstat-name>
<inuse>129</inuse>
<memuse>14</memuse>
<high-use>--</high-use>
<memstat-req>139</memstat-req>
<memstat-size>16,32,64,128,1024,16384</memstat-size>
<memstat-name>sysctl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>14847</memstat-req>
<memstat-size>16,32,64,4096,16384,32768</memstat-size>
<memstat-name>ifaddr</memstat-name>
<inuse>29</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>29</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mkey</memstat-name>
<inuse>345</inuse>
<memuse>6</memuse>
<high-use>--</high-use>
<memstat-req>2527</memstat-req>
<memstat-size>16,128</memstat-size>
<memstat-name>pfe_ipc</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1422</memstat-req>

<memstat-size>16,32,64,128,512,1024,2048,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>ifstate</memstat-name>
<inuse>594</inuse>
<memuse>51</memuse>
<high-use>--</high-use>
<memstat-req>655</memstat-req>

<memstat-size>16,32,64,128,256,1024,2048,4096,16384,32768</memstat-size>
<memstat-name>itable16</memstat-name>
<inuse>276</inuse>
<memuse>52</memuse>
<high-use>--</high-use>
<memstat-req>294</memstat-req>
<memstat-size>1024,4096</memstat-size>
<memstat-name>itable32</memstat-name>
```

```

<inuse>160</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>160</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>itable64</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>lr</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pic</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64,512</memstat-size>
<memstat-name>pfestat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>162</memstat-req>
<memstat-size>16,32,128,256,16384</memstat-size>
<memstat-name>gencfg</memstat-name>
<inuse>224</inuse>
<memuse>56</memuse>
<high-use>--</high-use>
<memstat-req>540</memstat-req>
<memstat-size>16,32,64,256,512,32768,65536</memstat-size>
<memstat-name>jsr</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>idl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>16,32,64,128,256,4096,16384,32768,131072</memstat-size>

<memstat-name>rtsmsg</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>module</memstat-name>
<inuse>249</inuse>
<memuse>16</memuse>
<high-use>--</high-use>
<memstat-req>249</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mtx_pool</memstat-name>

```

```
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>DEVFS3</memstat-name>
<inuse>109</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>117</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>DEVFS1</memstat-name>
<inuse>102</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>109</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>pgrp</memstat-name>
<inuse>12</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>session</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>proc</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>subproc</memstat-name>
<inuse>244</inuse>
<memuse>496</memuse>
<high-use>--</high-use>
<memstat-req>1522</memstat-req>
<memstat-size>2048,131072</memstat-size>
<memstat-name>cred</memstat-name>
<inuse>30</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>11409</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>plimit</memstat-name>
<inuse>17</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>133</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>uidinfo</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>32,512</memstat-size>
<memstat-name>sysctluid</memstat-name>
<inuse>1117</inuse>
```

```

<memuse>34</memuse>
<high-use>--</high-use>
<memstat-req>1117</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sysctltmp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>743</memstat-req>
<memstat-size>16,32,64,1024</memstat-size>
<memstat-name>umtx</memstat-name>
<inuse>144</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>144</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>SWAP</memstat-name>
<inuse>2</inuse>
<memuse>209</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>bus</memstat-name>
<inuse>496</inuse>
<memuse>55</memuse>
<high-use>--</high-use>
<memstat-req>1196</memstat-req>
<memstat-size>16,32,64,128,32768</memstat-size>
<memstat-name>bus-sc</memstat-name>
<inuse>23</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>335</memstat-req>

<memstat-size>16,32,64,512,1024,2048,8192,16384,65536,131072</memstat-size>
<memstat-name>devstat</memstat-name>
<inuse>10</inuse>
<memuse>21</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>16,131072</memstat-size>
<memstat-name>eventhandler</memstat-name>
<inuse>35</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>36</memstat-req>
<memstat-size>32,128</memstat-size>
<memstat-name>kobj</memstat-name>
<inuse>93</inuse>
<memuse>186</memuse>
<high-use>--</high-use>
<memstat-req>111</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>DEVFS</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>rman</memstat-name>
<inuse>71</inuse>

```

```
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>433</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sbuf</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>522</memstat-req>
<memstat-size>16,32,32768,131072</memstat-size>
<memstat-name>NULLFS hash</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>taskqueue</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>turnstiles</memstat-name>
<inuse>145</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>145</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>Unitno</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>44</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>iocltops</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>27622</memstat-req>
<memstat-size>16,64,8192,16384,131072</memstat-size>
<memstat-name>iov</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>18578</memstat-req>
<memstat-size>16,64,128,256,512,1024,2048,131072</memstat-size>
<memstat-name>msg</memstat-name>
<inuse>4</inuse>
<memuse>25</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32768,131072</memstat-size>
<memstat-name>sem</memstat-name>
<inuse>4</inuse>
<memuse>7</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16384,32768,131072</memstat-size>
<memstat-name>shm</memstat-name>
<inuse>9</inuse>
<memuse>20</memuse>
```



```

<high-use>--</high-use>
<memstat-req>14</memstat-req>
<memstat-size>32768</memstat-size>
<memstat-name>ttys</memstat-name>
<inuse>321</inuse>
<memuse>61</memuse>
<high-use>--</high-use>
<memstat-req>528</memstat-req>
<memstat-size>512,32768</memstat-size>
<memstat-name>ptys</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>mbuf_tag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>23383</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>soname</memstat-name>
<inuse>115</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>24712</memstat-req>
<memstat-size>16,32,64,256</memstat-size>
<memstat-name>pcb</memstat-name>
<inuse>216</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>484</memstat-req>
<memstat-size>16,32,64,128,1024,2048,4096,16384,32768,65536</memstat-size>
<memstat-name>BIO buffer</memstat-name>
<inuse>43</inuse>
<memuse>86</memuse>
<high-use>--</high-use>
<memstat-req>405</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>vfscache</memstat-name>
<inuse>1</inuse>
<memuse>256</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>cluster_save buffer</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>VFS hash</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>vnodes</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>mount</memstat-name>
<inuse>290</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>535</memstat-req>
<memstat-size>16,32,64,128,256,4096,32768</memstat-size>
<memstat-name>vnodemarker</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>498</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pfs_nodes</memstat-name>
<inuse>25</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>25</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>pfs_vncache</memstat-name>
<inuse>27</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>53</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>STP</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>GEOM</memstat-name>
<inuse>146</inuse>
<memuse>11</memuse>
<high-use>--</high-use>
<memstat-req>1042</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>syncache</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>tlv_stat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>8</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>NFS_daemon</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
```

```

<memstat-name>p1003.1b</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>MD disk</memstat-name>
<inuse>10</inuse>
<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ata_generic</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>16,16384,32768</memstat-size>
<memstat-name>ISofs mount</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>ISofs node</memstat-name>
<inuse>1440</inuse>
<memuse>135</memuse>
<high-use>--</high-use>
<memstat-req>1457</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>CAM SIM</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>CAM XPT</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64,16384</memstat-size>
<memstat-name>CAM periph</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ad_driver</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>pagedep</memstat-name>
<inuse>1</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>106</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>inodedep</memstat-name>

```

```
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>464</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>newblk</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>336</memstat-req>
<memstat-size>64,4096</memstat-size>
<memstat-name>bmsafemap</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>63</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>allocdirect</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>320</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>indirdep</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>17</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>allocindir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>freefrag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>12</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>freeblks</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>40</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>freefile</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>101</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>diradd</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>465</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>mkdir</memstat-name>
<inuse>0</inuse>
```

```

<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>136</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>dirrem</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>168</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>newdirblk</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>savedino</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>157</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>UFS mount</memstat-name>
<inuse>15</inuse>
<memuse>36</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>2048,65536,131072</memstat-size>
<memstat-name>ata_dma</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>UMAHash</memstat-name>
<inuse>1</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>4096,16384,32768,65536</memstat-size>
<memstat-name>cdev</memstat-name>
<inuse>22</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>22</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>file desc</memstat-name>
<inuse>141</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>1583</memstat-req>
<memstat-size>16,1024,2048,16384</memstat-size>
<memstat-name>VM pgdata</memstat-name>
<inuse>2</inuse>
<memuse>65</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>sigio</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>--</high-use>
<memstat-req>20</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kenv</memstat-name>
<inuse>24</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>27</memstat-req>
<memstat-size>16,32,64,131072</memstat-size>
<memstat-name>atkbddev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kqueue</memstat-name>
<inuse>15</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>19</memstat-req>
<memstat-size>1024,4096,32768</memstat-size>
<memstat-name>proc-args</memstat-name>
<inuse>57</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>1001</memstat-req>
<memstat-size>16,32,64,128,256,512,1024</memstat-size>
<memstat-name>isadev</memstat-name>
<inuse>21</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>zombie</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1278</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ithread</memstat-name>
<inuse>69</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>69</memstat-req>
<memstat-size>16,64,256</memstat-size>
<memstat-name>legacydrv</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>memdesc</memstat-name>
<inuse>1</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>nexusdev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
```

```

<memstat-req>2</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>CAM queue</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>$PIR</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>KTRACE</memstat-name>
<inuse>100</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>100</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>kbdmux</memstat-name>
<inuse>5</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>128,2048,65536,131072</memstat-size>
</vmstat-memstat-malloc>
<vmstat-memstat-zone>
  <zone-name>UMA Kegs:</zone-name>
  <zone-size>136</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>1</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Zones:</zone-name>
  <zone-size>120</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>19</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Slabs:</zone-name>
  <zone-size>64</zone-size>
  <count-limit>0</count-limit>
  <used>490</used>
  <free>41</free>
  <zone-req>579</zone-req>
  <zone-name>UMA RCntSlabs:</zone-name>
  <zone-size>104</zone-size>
  <count-limit>0</count-limit>
  <used>276</used>
  <free>20</free>
  <zone-req>276</zone-req>
  <zone-name>UMA Hash:</zone-name>
  <zone-size>128</zone-size>
  <count-limit>0</count-limit>
  <used>4</used>
  <free>26</free>
  <zone-req>5</zone-req>
  <zone-name>16 Bucket:</zone-name>
  <zone-size>76</zone-size>
  <count-limit>0</count-limit>

```

```
<used>30</used>
<free>20</free>
<zone-req>30</zone-req>
<zone-name>32 Bucket:</zone-name>
<zone-size>140</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>23</free>
<zone-req>33</zone-req>
<zone-name>64 Bucket:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>9</free>
<zone-req>33</zone-req>
<zone-name>128 Bucket:</zone-name>
<zone-size>524</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>0</free>
<zone-req>49</zone-req>
<zone-name>VM OBJECT:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>2111</used>
<free>79</free>
<zone-req>25214</zone-req>
<zone-name>MAP:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>7</used>
<free>41</free>
<zone-req>7</zone-req>
<zone-name>KMAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>35336</count-limit>
<used>19</used>
<free>149</free>
<zone-req>2397</zone-req>
<zone-name>MAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2031</used>
<free>153</free>
<zone-req>62417</zone-req>
<zone-name>PV ENTRY:</zone-name>
<zone-size>24</zone-size>
<count-limit>509095</count-limit>
<used>57177</used>
<free>6333</free>
<zone-req>1033683</zone-req>
<zone-name>DP fakepg:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mt_zone:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>238</used>
```



```
<free>57</free>
<zone-req>238</zone-req>
<zone-name>16:</zone-name>
<zone-size>16</zone-size>
<count-limit>0</count-limit>
<used>2114</used>
<free>119</free>
<zone-req>80515</zone-req>
<zone-name>32:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>1335</used>
<free>134</free>
<zone-req>10259</zone-req>
<zone-name>64:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>3529</used>
<free>129</free>
<zone-req>29110</zone-req>
<zone-name>96:</zone-name>
<zone-size>96</zone-size>
<count-limit>0</count-limit>
<used>2062</used>
<free>58</free>
<zone-req>4365</zone-req>
<zone-name>112:</zone-name>
<zone-size>112</zone-size>
<count-limit>0</count-limit>
<used>361</used>
<free>164</free>
<zone-req>24613</zone-req>
<zone-name>128:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>359</used>
<free>61</free>
<zone-req>942</zone-req>
<zone-name>160:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>364</used>
<free>44</free>
<zone-req>577</zone-req>
<zone-name>224:</zone-name>
<zone-size>224</zone-size>
<count-limit>0</count-limit>
<used>422</used>
<free>20</free>
<zone-req>1950</zone-req>
<zone-name>256:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>204</used>
<free>36</free>
<zone-req>1225</zone-req>
<zone-name>288:</zone-name>
<zone-size>288</zone-size>
<count-limit>0</count-limit>
<used>2</used>
<free>24</free>
```

```
<zone-req>10</zone-req>
<zone-name>512:</zone-name>
<zone-size>512</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>7</free>
<zone-req>911</zone-req>
<zone-name>1024:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>213</used>
<free>11</free>
<zone-req>1076</zone-req>
<zone-name>2048:</zone-name>
<zone-size>2048</zone-size>
<count-limit>0</count-limit>
<used>199</used>
<free>113</free>
<zone-req>640</zone-req>
<zone-name>4096:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>144</used>
<free>7</free>
<zone-req>2249</zone-req>
<zone-name>Files:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>665</used>
<free>77</free>
<zone-req>16457</zone-req>
<zone-name>MAC labels:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>3998</used>
<free>227</free>
<zone-req>21947</zone-req>
<zone-name>PROC:</zone-name>
<zone-size>544</zone-size>
<count-limit>0</count-limit>
<used>116</used>
<free>10</free>
<zone-req>1394</zone-req>
<zone-name>THREAD:</zone-name>
<zone-size>416</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>17</free>
<zone-req>131</zone-req>
<zone-name>KSEGRP:</zone-name>
<zone-size>88</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>73</free>
<zone-req>131</zone-req>
<zone-name>UPCALL:</zone-name>
<zone-size>44</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
```

```

<zone-name>SLEEPQUEUE:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>145</used>
<free>194</free>
<zone-req>145</zone-req>
<zone-name>VMSPACE:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>57</used>
<free>13</free>
<zone-req>1335</zone-req>
<zone-name>mbuf_packet:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>256</used>
<free>128</free>
<zone-req>49791</zone-req>
<zone-name>mbuf:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>50</used>
<free>466</free>
<zone-req>105183</zone-req>
<zone-name>mbuf_cluster:</zone-name>
<zone-size>2048</zone-size>
<count-limit>25190</count-limit>
<used>387</used>
<free>165</free>
<zone-req>5976</zone-req>
<zone-name>mbuf_jumbo_pagesize:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_9k:</zone-name>
<zone-size>9216</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_16k:</zone-name>
<zone-size>16384</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ACL_UMA_zone:</zone-name>
<zone-size>388</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>g_bio:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>174</free>
<zone-req>69750</zone-req>
<zone-name>ata_request:</zone-name>

```

```
<zone-size>200</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>57</free>
<zone-req>5030</zone-req>
<zone-name>ata_composite:</zone-name>
<zone-size>192</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>GENCFG:</zone-name>
<zone-size>72</zone-size>
<count-limit>1000004</count-limit>
<used>57</used>
<free>102</free>
<zone-req>57</zone-req>
<zone-name>VNODE:</zone-name>
<zone-size>292</zone-size>
<count-limit>0</count-limit>
<used>2718</used>
<free>25</free>
<zone-req>2922</zone-req>
<zone-name>VNODEPOLL:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>S VFS Cache:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2500</used>
<free>76</free>
<zone-req>3824</zone-req>
<zone-name>L VFS Cache:</zone-name>
<zone-size>291</zone-size>
<count-limit>0</count-limit>
<used>51</used>
<free>14</free>
<zone-req>63</zone-req>
<zone-name>NAMEI:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>8</free>
<zone-req>53330</zone-req>
<zone-name>NFSMOUNT:</zone-name>
<zone-size>480</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>NFSNODE:</zone-name>
<zone-size>460</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>PIPE:</zone-name>
<zone-size>404</zone-size>
```

```
<count-limit>0</count-limit>
<used>27</used>
<free>9</free>
<zone-req>717</zone-req>
<zone-name>KNOTE:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>42</used>
<free>64</free>
<zone-req>3311</zone-req>
<zone-name>socket:</zone-name>
<zone-size>412</zone-size>
<count-limit>25191</count-limit>
<used>343</used>
<free>8</free>
<zone-req>2524</zone-req>
<zone-name>unpcb:</zone-name>
<zone-size>140</zone-size>
<count-limit>25200</count-limit>
<used>170</used>
<free>26</free>
<zone-req>2157</zone-req>
<zone-name>ipq:</zone-name>
<zone-size>52</zone-size>
<count-limit>216</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>udpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>19</used>
<free>32</free>
<zone-req>31</zone-req>
<zone-name>inpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>40</used>
<free>28</free>
<zone-req>105</zone-req>
<zone-name>tcpcb:</zone-name>
<zone-size>520</zone-size>
<count-limit>25193</count-limit>
<used>40</used>
<free>16</free>
<zone-req>105</zone-req>
<zone-name>tcptw:</zone-name>
<zone-size>56</zone-size>
<count-limit>5092</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>syncache:</zone-name>
<zone-size>128</zone-size>
<count-limit>15360</count-limit>
<used>0</used>
<free>60</free>
<zone-req>55</zone-req>
<zone-name>tcpreass:</zone-name>
<zone-size>20</zone-size>
<count-limit>1690</count-limit>
```

```
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>sackhole:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ripcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>5</used>
<free>29</free>
<zone-req>5</zone-req>
<zone-name>SWAPMETA:</zone-name>
<zone-size>276</zone-size>
<count-limit>94948</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>FFS inode:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
<free>72</free>
<zone-req>1306</zone-req>
<zone-name>FFS1 dinode:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
<free>24</free>
<zone-req>1306</zone-req>
<zone-name>FFS2 dinode:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
</vmstat-memstat-zone>
<vmstat-sumstat>
  <cpu-context-switch>934906</cpu-context-switch>
  <dev-intr>1707986</dev-intr>
  <soft-intr>33819</soft-intr>
  <traps>203604</traps>
  <sys-calls>1200636</sys-calls>
  <kernel-thrds>60</kernel-thrds>
  <fork-calls>1313</fork-calls>
  <vfork-calls>21</vfork-calls>
  <rfork-calls>0</rfork-calls>
  <swap-pageins>0</swap-pageins>
  <swap-pagedin>0</swap-pagedin>
  <swap-pageouts>0</swap-pageouts>
  <swap-pagedout>0</swap-pagedout>
  <vnode-pageins>23094</vnode-pageins>
  <vnode-pagedin>23119</vnode-pagedin>
  <vnode-pageouts>226</vnode-pageouts>
  <vnode-pagedout>3143</vnode-pagedout>
  <page-daemon-wakeup>0</page-daemon-wakeup>
  <page-daemon-examined-pages>0</page-daemon-examined-pages>
  <pages-reactivated>8821</pages-reactivated>
```

```

<copy-on-write-faults>48364</copy-on-write-faults>
<copy-on-write-optimized-faults>31</copy-on-write-optimized-faults>
<zero-fill-pages-zeroed>74665</zero-fill-pages-zeroed>
<zero-fill-pages-prezeroed>70061</zero-fill-pages-prezeroed>
<transit-blocking-page-faults>85</transit-blocking-page-faults>
<total-vm-faults>191824</total-vm-faults>

<pages-affected-by-kernel-thrd-creat>0</pages-affected-by-kernel-thrd-creat>
<pages-affected-by-fork>95343</pages-affected-by-fork>
<pages-affected-by-vfork>3526</pages-affected-by-vfork>
<pages-affected-by-rfork>0</pages-affected-by-rfork>
<pages-freed>221502</pages-freed>
<pages-freed-by-daemon>0</pages-freed-by-daemon>
<pages-freed-by-exiting-proc>75630</pages-freed-by-exiting-proc>
<pages-active>45826</pages-active>
<pages-inactive>13227</pages-inactive>
<pages-in-vm-cache>49278</pages-in-vm-cache>
<pages-wired-down>10640</pages-wired-down>
<pages-free>70706</pages-free>
<bytes-per-page>4096</bytes-per-page>
<swap-pages-used>0</swap-pages-used>
<peak-swap-pages-used>0</peak-swap-pages-used>
<total-name-lookups>214496</total-name-lookups>
<positive-cache-hits>92</positive-cache-hits>
<negative-cache-hits>5</negative-cache-hits>
<pass2>0</pass2>
<cache-deletions>0</cache-deletions>
<cache-falsehits>0</cache-falsehits>
<toolong>0</toolong>
</vmstat-sumstat>
<vmstat-intr>
  <intr-name>irq0: clk          </intr-name>
  <intr-cnt>1243455</intr-cnt>
  <intr-rate>999</intr-rate>
  <intr-name>irq4: sio0        </intr-name>
  <intr-cnt>1140</intr-cnt>
  <intr-rate>0</intr-rate>
  <intr-name>irq8: rtc         </intr-name>
  <intr-cnt>159164</intr-cnt>
  <intr-rate>127</intr-rate>
  <intr-name>irq9: cbb1 fxp0    </intr-name>
  <intr-cnt>28490</intr-cnt>
  <intr-rate>22</intr-rate>
  <intr-name>irq10: fxp1       </intr-name>
  <intr-cnt>20593</intr-cnt>
  <intr-rate>16</intr-rate>
  <intr-name>irq14: ata0       </intr-name>
  <intr-cnt>5031</intr-cnt>
  <intr-rate>4</intr-rate>
  <intr-name>Total</intr-name>
  <intr-cnt>1457873</intr-cnt>
  <intr-rate>1171</intr-rate>
</vmstat-intr>
<vm-kernel-state>
  <vm-kmem-map-free>248524800</vm-kmem-map-free>
</vm-kernel-state>
</system-virtual-memory-information>
<cli>
  <banner></banner>
</cli>
</rpc-reply>

```





## show task replication

<b>Syntax</b>	<b>show task replication</b>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.</p> <p>Support for logical systems introduced in Junos OS Release 13.3</p>
<b>Description</b>	Displays nonstop active routing (NSR) status. When you issue this command on the master Routing Engine, the status of nonstop active routing synchronization is also displayed.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show task replication (Issued on the Master Routing Engine) on page 289</a> <a href="#">show task replication (Issued on the Backup Routing Engine) on page 290</a>
<b>Output Fields</b>	Table 22 on page 289 lists the output fields for the <b>show task replication</b> command. Output fields are listed in the approximate order in which they appear.

**Table 22: show task replication Output Fields**

Field Name	Field Description
<b>Stateful replication</b>	Displays whether or not graceful Routing Engine switchover is configured. The status can be <b>Enabled</b> or <b>Disabled</b> .
<b>RE mode</b>	Displays the Routing Engine on which the command is issued: <b>Master</b> , <b>Backup</b> , or <b>Not applicable</b> (when the router has only one Routing Engine).
<b>Protocol</b>	Protocols that are supported by nonstop active routing.
<b>Synchronization Status</b>	Nonstop active routing synchronization status for the supported protocols. States are <b>NotStarted</b> , <b>InProgress</b> , and <b>Complete</b> .

## Sample Output

### show task replication (Issued on the Master Routing Engine)

```

user@host> show task replication
  Stateful Replication: Enabled
    RE mode: Master

  Protocol              Synchronization Status
  OSPF                  NotStarted
  BGP                   Complete
  IS-IS                 NotStarted

```

LDP	Complete
PIM	Complete

**show task replication (Issued on the Backup Routing Engine)**

```
user@host> show task replication
Stateful Replication: Enabled
RE mode: Backup
```

## show version

<b>List of Syntax</b>	<a href="#">Syntax on page 291</a> <a href="#">Syntax (EX Series Switches) on page 291</a> <a href="#">Syntax (TX Matrix Router) on page 291</a> <a href="#">Syntax (TX Matrix Plus Router) on page 291</a> <a href="#">Syntax (MX Series Router) on page 291</a> <a href="#">Syntax (QFX Series) on page 291</a>
<b>Syntax</b>	<pre>show version &lt;brief   detail&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show version &lt;all-members&gt; &lt;brief   detail&gt; &lt;local&gt; &lt;member member-id&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>show version &lt;brief   detail&gt; &lt;all-chassis   all-lcc   lcc number   scc&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>show version &lt;all-chassis   all-lcc   lcc number   sfc number&gt; &lt;brief   detail&gt;</pre>
<b>Syntax (MX Series Router)</b>	<pre>show version &lt;brief   detail&gt; &lt;all-members&gt; &lt;local&gt; &lt;member member-id&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show version &lt;brief   detail&gt; &lt;component component-name   all&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Display the hostname and version information about the software running on the router or switch.</p> <p>Beginning in Junos OS Release 13.3, the <b>show version</b> command output includes the <b>Junos</b> field that displays the Junos OS version running on the device. This field provides a consistent means of identifying the Junos OS version, rather than extracting that information from the list of installed sub-packages.</p>
<b>Options</b>	<p><b>none</b>—Display standard information about the hostname and version of the software running on the router or switch.</p>

**brief | detail**—(Optional) Display the specified level of output.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on all members of the Virtual Chassis configuration.

**component all**—(QFabric systems only) (Optional) Display the host name and version information about the software running on all the components on the QFabric system.

**component *component-name***—(QFabric systems only) (Optional) Display the host name and version information about the software running on a specific QFabric system component. Replace *component-name* with the name of the QFabric system component. The *component-name* can be the name of a diagnostics Routing Engine, Director group, fabric control Routing Engine, fabric manager Routing Engine, Interconnect device, or Node group.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**scc**—(TX Matrix routers only) (Optional) Display the hostname and version information about the software running on the TX Matrix router (or switch-card chassis).

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the host name and version information about the software running on for a specified T640 router (line-card chassis or LCC) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the host name and version information about the software running for a specified T1600 or T4000 router (LCC) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display the hostname and version information about the software running on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

**Additional Information** By default, when you issue the **show version** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) or T1600 or T4000 (in a routing matrix based on a TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) or T1600 or T4000 (in a routing matrix based on a TX Matrix Plus router) backup Routing Engines that are connected to it.

**Required Privilege Level** view

**List of Sample Output** [show version \(Devices Running Junos OS Release 13.3 and Later\) on page 294](#)  
[show version on page 294](#)  
[show version \(TX Matrix Plus Router\) on page 295](#)  
[show version \(TX Matrix Plus Router with 3D SIBs\) on page 297](#)  
[show version \(MX Series Router\) on page 301](#)  
[show version \(QFX3500 Switch\) on page 301](#)  
[show version \(QFabric System\) on page 301](#)  
[show version component all \(QFabric System\) on page 302](#)

## Sample Output

### show version (Devices Running Junos OS Release 13.3 and Later)

The following output is from the MX240 Router and shows the **Junos** field introduced in Junos OS 13.3. Depending on the platform running Junos OS 13.3, you might see different installed sub-packages, but the **Junos** field is common across all platforms that run Junos OS 13.3 and later.

```
user@host > show version
Hostname: lab
Model: mx240
Junos: 13.3R1.4
JUNOS Base OS boot [13.3R1.4]
JUNOS Base OS Software Suite [13.3R1.4]
JUNOS Kernel Software Suite [13.3R1.4]
JUNOS Crypto Software Suite [13.3R1.4]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [13.3R1.4]
JUNOS Packet Forwarding Engine Support (MX Common) [13.3R1.4]
JUNOS Online Documentation [13.3R1.4]
JUNOS Services ACL Container package [13.3R1.4]
JUNOS Services Application Level Gateways [13.3R1.4]
JUNOS AppId Services [13.3R1.4]
JUNOS Border Gateway Function package [13.3R1.4]
JUNOS Services Captive Portal and Content Delivery Container package [13.3R1.4]
JUNOS Services HTTP Content Management package [13.3R1.4]
JUNOS IDP Services [13.3R1.4]
JUNOS Services Jflow Container package [13.3R1.4]
JUNOS Services LL-PDF Container package [13.3R1.4]
JUNOS Services MobileNext Software package [13.3R1.4]
JUNOS Services Mobile Subscriber Service Container package [13.3R1.4]
JUNOS Services NAT [13.3R1.4]
JUNOS Services PTSP Container package [13.3R1.4]
JUNOS Services RPM [13.3R1.4]
JUNOS Services Stateful Firewall [13.3R1.4]
JUNOS Voice Services Container package [13.3R1.4]
JUNOS Services Crypto [13.3R1.4]
JUNOS Services SSL [13.3R1.4]
JUNOS Services IPSec [13.3R1.4]
JUNOS platform Software Suite [13.3R1.4]
JUNOS Runtime Software Suite [13.3R1.4]
JUNOS Routing Software Suite [13.3R1.4]
JUNOS py-base-i386 [13.3R1.4]
```

### show version

```
user@host> show version
Hostname: router1
Model: m20
JUNOS Base OS boot [7.2-20050312.0]
JUNOS Base OS Software Suite [7.2-20050312.0]
JUNOS Kernel Software Suite [7.2R1.7]
JUNOS Packet Forwarding Engine Support (M20/M40) [7.2R1.7]
JUNOS Routing Software Suite [7.2R1.7]
JUNOS Online Documentation [7.2R1.7]
JUNOS Crypto Software Suite [7.2R1.7]

{master}
user@host> show version psd 1
```

```
psd1-re0:
```

```
-----
Hostname: china
Model: t640
JUNOS Base OS boot [9.1I20080311_1959_builder]
JUNOS Base OS Software Suite [9.1-20080321.0]
JUNOS Kernel Software Suite [9.1-20080321.0]
JUNOS Crypto Software Suite [9.1-20080321.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.1-20080321.0]
JUNOS Packet Forwarding Engine Support (T-series) [9.1-20080321.0]
JUNOS Online Documentation [9.1-20080321.0]
JUNOS Routing Software Suite [9.1-20080321.0]
labpkg [7.0]
```

### show version (TX Matrix Plus Router)

```
user@host> show version
```

```
sfc0-re0:
```

```
-----
Hostname: host
Model: txp
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]
```

```
lcc0-re0:
```

```
-----
Hostname: host1
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
```

JUNOS Services Application Level Gateways [12.3-20121019.0]  
JUNOS AppId Services [12.3-20121019.0]  
JUNOS Border Gateway Function package [12.3-20121019.0]  
JUNOS Services Captive Portal and Content Delivery Container package [12.3-20121019.0]  
JUNOS Services HTTP Content Management package [12.3-20121019.0]  
JUNOS IDP Services [12.3-20121019.0]  
JUNOS Services LL-PDF Container package [12.3-20121019.0]  
JUNOS Services NAT [12.3-20121019.0]  
JUNOS Services PTSP Container package [12.3-20121019.0]  
JUNOS Services RPM [12.3-20121019.0]  
JUNOS Services Stateful Firewall [12.3-20121019.0]  
JUNOS Voice Services Container package [12.3-20121019.0]  
JUNOS Services Example Container package [12.3-20121019.0]  
JUNOS Services Crypto [12.3-20121019.0]  
JUNOS Services SSL [12.3-20121019.0]  
JUNOS Services IPSec [12.3-20121019.0]  
JUNOS Runtime Software Suite [12.3-20121019.0]  
JUNOS Routing Software Suite [12.3-20121019.0]

lcc1-re0:

-----  
Hostname: host2  
Model: t1600  
JUNOS Base OS boot [12.3-20121019.0]  
JUNOS Base OS Software Suite [12.3-20121019.0]  
JUNOS Kernel Software Suite [12.3-20121019.0]  
JUNOS Crypto Software Suite [12.3-20121019.0]  
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]  
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]  
JUNOS Online Documentation [12.3-20121019.0]  
JUNOS Services ACL Container package [12.3-20121019.0]  
JUNOS Services Application Level Gateways [12.3-20121019.0]  
JUNOS AppId Services [12.3-20121019.0]  
JUNOS Border Gateway Function package [12.3-20121019.0]  
JUNOS Services Captive Portal and Content Delivery Container package [12.3-20121019.0]  
JUNOS Services HTTP Content Management package [12.3-20121019.0]  
JUNOS IDP Services [12.3-20121019.0]  
JUNOS Services LL-PDF Container package [12.3-20121019.0]  
JUNOS Services NAT [12.3-20121019.0]  
JUNOS Services PTSP Container package [12.3-20121019.0]  
JUNOS Services RPM [12.3-20121019.0]  
JUNOS Services Stateful Firewall [12.3-20121019.0]  
JUNOS Voice Services Container package [12.3-20121019.0]  
JUNOS Services Example Container package [12.3-20121019.0]  
JUNOS Services Crypto [12.3-20121019.0]  
JUNOS Services SSL [12.3-20121019.0]  
JUNOS Services IPSec [12.3-20121019.0]  
JUNOS Runtime Software Suite [12.3-20121019.0]  
JUNOS Routing Software Suite [12.3-20121019.0]

lcc2-re0:

-----  
Hostname: host3  
Model: t1600  
JUNOS Base OS boot [12.3-20121019.0]  
JUNOS Base OS Software Suite [12.3-20121019.0]  
JUNOS Kernel Software Suite [12.3-20121019.0]  
JUNOS Crypto Software Suite [12.3-20121019.0]  
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]



```

JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]

```

```
lcc3-re0:
```

```

-----
Hostname: host4
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]

```

#### show version (TX Matrix Plus Router with 3D SIBs)

```

user@host>show version
sfc0-re0:

```

```

-----
Hostname: sfc0

```

```
Model: txp
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services ACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]
```

```
lcc0-re0:
```

```
-----
Hostname: lcc0
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services ACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
```

```
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]
```

```
lcc2-re0:
```

```
-----
Hostname: lcc2
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]
```

```
lcc4-re0:
```

```
-----
Hostname: lcc4
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
```

JUNOS Services Jflow Container package [13.1-20130306.0]  
JUNOS Services LL-PDF Container package [13.1-20130306.0]  
JUNOS Services MobileNext Software package [13.1-20130306.0]  
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]  
JUNOS Services NAT [13.1-20130306.0]  
JUNOS Services PTSP Container package [13.1-20130306.0]  
JUNOS Services RPM [13.1-20130306.0]  
JUNOS Services Stateful Firewall [13.1-20130306.0]  
JUNOS Voice Services Container package [13.1-20130306.0]  
JUNOS Services Example Container package [13.1-20130306.0]  
JUNOS Services Crypto [13.1-20130306.0]  
JUNOS Services SSL [13.1-20130306.0]  
JUNOS Services IPSec [13.1-20130306.0]  
JUNOS Runtime Software Suite [13.1-20130306.0]  
JUNOS Routing Software Suite [13.1-20130306.0]

lcc6-re0:

-----  
Hostname: lcc6  
Model: t1600  
JUNOS Base OS boot [13.1-20130306.0]  
JUNOS Base OS Software Suite [13.1-20130306.0]  
JUNOS Kernel Software Suite [13.1-20130306.0]  
JUNOS Crypto Software Suite [13.1-20130306.0]  
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]  
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]  
JUNOS Online Documentation [13.1-20130306.0]  
JUNOS Services AACL Container package [13.1-20130306.0]  
JUNOS Services Application Level Gateways [13.1-20130306.0]  
JUNOS AppId Services [13.1-20130306.0]  
JUNOS Border Gateway Function package [13.1-20130306.0]  
JUNOS Services Captive Portal and Content Delivery Container package [13.1-20130306.0]  
JUNOS Services HTTP Content Management package [13.1-20130306.0]  
JUNOS IDP Services [13.1-20130306.0]  
JUNOS Services Jflow Container package [13.1-20130306.0]  
JUNOS Services LL-PDF Container package [13.1-20130306.0]  
JUNOS Services MobileNext Software package [13.1-20130306.0]  
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]  
JUNOS Services NAT [13.1-20130306.0]  
JUNOS Services PTSP Container package [13.1-20130306.0]  
JUNOS Services RPM [13.1-20130306.0]  
JUNOS Services Stateful Firewall [13.1-20130306.0]  
JUNOS Voice Services Container package [13.1-20130306.0]  
JUNOS Services Example Container package [13.1-20130306.0]  
JUNOS Services Crypto [13.1-20130306.0]  
JUNOS Services SSL [13.1-20130306.0]  
JUNOS Services IPSec [13.1-20130306.0]  
JUNOS Runtime Software Suite [13.1-20130306.0]  
JUNOS Routing Software Suite [13.1-20130306.0]

lcc7-re0:

-----  
Hostname: lcc7  
Model: t1600  
JUNOS Base OS boot [13.1-20130306.0]  
JUNOS Base OS Software Suite [13.1-20130306.0]  
JUNOS Kernel Software Suite [13.1-20130306.0]  
JUNOS Crypto Software Suite [13.1-20130306.0]  
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]

```

JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]

```

#### show version (MX Series Router)

```

user@host5> show version
Hostname: host5
Model: mx80
JUNOS Base OS boot [11.3-20110717.0]
JUNOS Base OS Software Suite [11.3-20110717.0]
JUNOS Kernel Software Suite [11.3-20110717.0]
JUNOS Crypto Software Suite [11.3-20110717.0]
JUNOS Packet Forwarding Engine Support (MX80) [11.3-20110717.0]
JUNOS Online Documentation [11.3-20110717.0]
JUNOS Routing Software Suite [11.3-20110717.0]

```

#### show version (QFX3500 Switch)

```

user@switch> show version
Hostname: switch
Model: qfx_s3500
JUNOS Base OS boot [11.1R1]
JUNOS Base OS Software Suite [11.1R1]
JUNOS Kernel Software Suite [11.1R1]
JUNOS Crypto Software Suite [11.1R1]
JUNOS Online Documentation [11.1R1]
JUNOS Enterprise Software Suite [11.1R1]
JUNOS Packet Forwarding Engine Support (QFX) [11.1R1]
JUNOS Routing Software Suite [11.1R1]

```

#### show version (QFabric System)

```

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

```

**show version component all (QFabric System)**

```
user@switch> show version component all
dg1:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3R1.6]

dg0:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3R1.6]

NW-NG-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

FC-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

FC-1:
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

DRE-0:
-
Hostname: dre-0
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
```

```
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
```

```
FM-0:
```

```
-
```

```
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
```

```
nodedevice1:
```

```
-
```

```
Hostname: qfabric
Model: QFX3500
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
```

```
interconnectdevice1:
```

```
-
```

```
Hostname: qfabric
Model: QFX3108
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
warning: from interconnectdevice0: Disconnected
```

## show version fpc

---

<b>Syntax</b>	<code>show version fpc</code> <code>&lt;slot-number&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.4 for EX Series switches.
<b>Description</b>	Display the version of Junos OS for EX Series switches loaded on the line cards in an EX8200 switch.
<b>Options</b>	<b>none</b> —List the version of Junos OS for EX Series switches loaded on the line cards in the EX8200 switch.  <b>&lt;slot-number&gt;</b> —(Optional) Display the version of Junos OS for EX Series switches loaded on the line card slot specified by the slot-number value.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show version on page 291</a></li></ul>
<b>List of Sample Output</b>	<a href="#">show version fpc on page 304</a> <a href="#">show version fpc 5 on page 304</a>

## Sample Output

### show version fpc

```
user@switch> show version fpc
fpc 0 :
JUNOS 11.4I JUNOS 11.4I #0: 2011-12-07 11:33:18 UTC      dmadhuri@bng-junos-pool
40.juniper.net:/b/dmadhuri/nyse_0612/obj-powerpc/bsd/kerne1s/EX8200-LC/kerne1 p
owerpc
fpc 4 :
JUNOS 11.4I JUNOS 11.4I #0: 2011-12-07 11:33:18 UTC      dmadhuri@bng-junos-pool
40.juniper.net:/b/dmadhuri/nyse_0612/obj-powerpc/bsd/kerne1s/EX8200-LC/kerne1 p
owerpc
fpc 5 :
JUNOS 11.4I JUNOS 11.4I #0: 2011-12-07 11:33:18 UTC      dmadhuri@bng-junos-pool
40.juniper.net:/b/dmadhuri/nyse_0612/obj-powerpc/bsd/kerne1s/EX8200-LC/kerne1 p
owerpc
fpc 7 :
JUNOS 11.4I JUNOS 11.4I #0: 2011-12-07 11:33:18 UTC      dmadhuri@bng-junos-pool
40.juniper.net:/b/dmadhuri/nyse_0612/obj-powerpc/bsd/kerne1s/EX8200-LC/kerne1 p
owerpc
```

### show version fpc 5

```
user@switch> show version fpc 5
fpc 5 :
JUNOS 11.4I JUNOS 11.4I #0: 2011-12-07 11:33:18 UTC      dmadhuri@bng-junos-pool
40.juniper.net:/b/dmadhuri/nyse_0612/obj-powerpc/bsd/kerne1s/EX8200-LC/kerne1 p
owerpc
```



# Troubleshooting Procedures

- Troubleshooting Loss of the Root Password on page 305

## Troubleshooting Loss of the Root Password

---

**Problem**    **Description:** If you forget the root password for a switch, use the password recovery procedure to reset the root password.



**NOTE:** You need physical access to the switch to recover the root password.




**TIP:** For a video on recovering the root password for routers, see *Recovering the Root Password*. The procedure is similar for switches.

**Solution**    To recover the root password:

1. Power off your switch by unplugging the power cord or turning off the power at the wall switch.
2. Insert one end of the Ethernet cable into the serial port on the management device and connect the other end to the console port on the back of the switch. See [Figure 1 on page 306](#).

Figure 1: Connecting to the Console Port on the EX Series Switch



3. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate COM port to use (for example, COM1).
4. Configure the port settings as follows:
  - Bits per second: 9600
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: None
5. Power on your switch by plugging in the power cord or turning on the power at the wall switch.
6. When the following prompt appears, press the Spacebar to access the switch's bootstrap loader command prompt:
 

```
Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [kernel] in 1 second...
```



**NOTE:** If the switch is in unattended mode for U-Boot, access to the bootstrap loader command prompt is blocked. If the root password is lost, you must reset the switch to the factory default configuration using the LCD panel. For more information, see [“Reverting to the Default Factory Configuration for the EX Series Switch”](#) on page 620.

7. At the following prompt, type **boot -s** to start up the system in single-user mode:
 

```
loader> boot -s
```
8. At the following prompt, type **recovery** to start the root password recovery procedure:
 

```
Enter full path name of shell or 'recovery' for root password recovery or RETURN for /bin/sh: recovery
```

A series of messages describe consistency checks, mounting of filesystems, and initialization and checkout of management services. Then the CLI prompt appears.

9. Enter configuration mode in the CLI:

```
user@switch> configure
```

10. Set the root password. For example:

```
user@switch# set system root-authentication plain-text-password
```

11. At the following prompt, enter the new root password. For example, juniper1:

```
user@switch# juniper1
```

```
Retype new password:
```

12. At the second prompt, reenter the new root password.

13. If you are finished configuring the network, commit the configuration.

```
root@switch# commit
```

```
commit complete
```

14. Exit configuration mode in the CLI.

```
root@switch# exit
```

15. Exit operational mode in the CLI.

```
root@switch> exit
```

16. At the prompt, enter **y** to reboot the switch.

```
Reboot the system? [y/n] y
```

#### Related Documentation

- *Connecting and Configuring an EX Series Switch (CLI Procedure)*
- *Connecting and Configuring an EX Series Switch (J-Web Procedure)*
- For information about configuring an encrypted root password, configuring SSH keys to authenticate root logins, and configuring special requirements for plain-text passwords, see *Configuring the Root Password*.



## PART 3

# User Interfaces

- [Overview on page 311](#)
- [Configuration on page 317](#)
- [Administration on page 513](#)



## CHAPTER 6

# Overview

- [Software Overview on page 311](#)
- [User Interfaces on page 313](#)

## Software Overview

---

- [Understanding Software Infrastructure and Processes on page 311](#)

### Understanding Software Infrastructure and Processes

Each switch runs the Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches on its general-purpose processors. Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the chassis.

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

With the J-Web interface and the command-line interface (CLI) to the Junos OS, you configure switching features and routing protocols and set the properties of network interfaces on your switch. After activating a software configuration, use either the J-Web or CLI user interface to monitor the switch, manage operations, and diagnose protocol and network connectivity problems.

- [Routing Engine and Packet Forwarding Engine on page 311](#)
- [Junos OS Processes on page 312](#)

### Routing Engine and Packet Forwarding Engine

---

A switch has two primary software processing components:

- **Packet Forwarding Engine**—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.
- **Routing Engine**—Provides three main functions:
  - Creates the packet forwarding switch fabric for the switch, providing route lookup, filtering, and switching on incoming data packets, then directing outbound packets to the appropriate interface for transmission to the network

- Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
- Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

### Junos OS Processes

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

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

Table 6 on page 30 describes the primary Junos OS processes.

**Table 23: Junos OS Processes**

Process	Name	Description
Chassis process	chassisd	<p>Detects hardware on the system that is used to configure network interfaces.</p> <p>Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered.</p> <p>Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.</p>
Ethernet switching process	eswd	<p>Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree protocol and access port security. The process is also responsible for managing Ethernet switching interfaces, VLANs, and VLAN interfaces.</p> <p>Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.</p>
Forwarding process	pfem	<p>Defines how routing protocols operate on the switch. The overall performance of the switch is largely determined by the effectiveness of the forwarding process.</p>
Interface process	dcd	<p>Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.</p>
Management process	mgd	<p>Provides communication between the other processes and an interface to the configuration database.</p> <p>Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured.</p> <p>Interacts with the other processes when commands are issued through one of the user interfaces on the switch.</p> <p>If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.</p>
Routing protocol process	rpd	<p>Defines how routing protocols such as RIP, OSPF, and BGP operate on the device, including selecting routes and maintaining forwarding tables.</p>



- Related Documentation**
- [For more information about processes, see \*Junos OS Network Operations Guide\*](#)
  - [For more information about basic system parameters, supported protocols, and software processes, see \*Junos OS System Basics Configuration Guide\*](#)

---

## User Interfaces

---

- [CLI User Interface Overview on page 313](#)
- [EX Series Switches Hardware and CLI Terminology Mapping on page 315](#)

### CLI User Interface Overview

You can use two interfaces to monitor, configure, troubleshoot, and manage a Juniper Networks EX Series Ethernet Switch: the J-Web graphical user interface and the Junos operating system (Junos OS) command-line interface (CLI). Both of these user interfaces are shipped with the switch. This topic describes the CLI. For information about the J-Web user interface, see *J-Web User Interface for EX Series Switches Overview*.

- [CLI Overview on page 313](#)
- [CLI Help and Command Completion on page 313](#)
- [CLI Command Modes on page 313](#)

---

#### CLI Overview

---

Junos operating system (Junos OS) CLI is a Juniper Networks specific command shell that runs on top of a UNIX-based operating system kernel. The CLI provides command help and command completion.

The CLI also provides a variety of UNIX utilities, such as Emacs-style keyboard sequences that allow you to move around on a command line and scroll through recently executed commands, regular expression matching to locate and replace values and identifiers in a configuration, filter command output, or log file entries, store and archive router files on a UNIX-based file system, and exit from the CLI environment and create a UNIX C shell or Bourne shell to navigate the file system, manage switch processes, and so on.

---

#### CLI Help and Command Completion

---

To access CLI Help, type a question mark (?) at any level of the hierarchy. The system displays a list of the available commands or statements and a short description of each.

To complete a command, statement, or option that you have partially typed, press the Tab key or the Spacebar. If the partially typed letters uniquely identify a command, the complete command name appears. Otherwise, a beep indicates that you have entered an ambiguous command and the possible completions are displayed. This completion feature also applies to other strings, such as filenames, interface names, usernames, and configuration statements.

---

#### CLI Command Modes

---

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

In operational mode, you enter commands to monitor and troubleshoot switch hardware and software and network connectivity. Operational mode is indicated by the `>` prompt—for example, `user@switch>`.

In configuration mode, you can define all properties of the Juniper Networks Junos operating system (Junos OS), including interfaces, VLANs, Virtual Chassis information, routing protocols, user access, and several system hardware properties.

To enter configuration mode, enter the **configure** command: .

```
user@switch> configure
```

Configuration mode is indicated by the `#` prompt, and includes the current location in the configuration hierarchy—for example:

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

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

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

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

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

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



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

#### Related Documentation

- [EX Series Switch Software Features Overview](#)
- [Junos OS CLI User Guide](#)

## EX Series Switches Hardware and CLI Terminology Mapping

The terms used to describe hardware components in EX Series switches documentation are sometimes different from the terms used in the Junos OS command line interface (CLI).

See the following topics to map the hardware terms used in EX Series switches documentation to the corresponding terms used in the CLI:

- *EX2200 Switch Hardware and CLI Terminology Mapping*
- *EX3200 Switch Hardware and CLI Terminology Mapping*
- *EX4200 Switch Hardware and CLI Terminology Mapping*
- *EX4500 Switch Hardware and CLI Terminology Mapping*
- *EX6210 Switch Hardware and CLI Terminology Mapping*
- *EX8208 Switch Hardware and CLI Terminology Mapping*
- *EX8216 Switch Hardware and CLI Terminology Mapping*

### **Related Documentation**

- *EX2200 Switches Hardware Overview*
- *EX3200 Switches Hardware Overview*
- *EX4200 Switches Hardware Overview*
- *EX4500 Switches Hardware Overview*
- *EX6210 Switch Hardware Overview*
- *EX8208 Switch Hardware Overview*
- *EX8216 Switch Hardware Overview*



## CHAPTER 7

# Configuration

- Configuration Tasks on page 317
- Statement Hierarchies on page 319

## Configuration Tasks

---

- Using the J-Web CLI Terminal on page 317

### Using the J-Web CLI Terminal



**NOTE:** This topic applies only to the J-Web Application package.

The J-Web CLI terminal provides access to the Junos OS command-line interface (CLI) through the J-Web interface. The functionality and behavior of the CLI available through the CLI Terminal page is the same as that of the Junos OS CLI available through the switch console. The CLI terminal supports all CLI commands and other features such as CLI help and autocompletion. Using the CLI terminal page, you can fully configure, monitor, and manage the switch.

This topic covers:

- Configuring the Web Browser on page 317
- Setting Domain Name, Hostname, and Name Server on page 318
- Enabling SSH on your system on page 318
- Sample Configuration on an EX Series Switch on page 318

### Configuring the Web Browser

---

Configure your Web browser as follows:

- Install Java Runtime Environment (JRE) version 1.4 or later on your system. JRE is a software package that must be installed on the client system to run Java applications. You can download the latest version of JRE from the Java software website <http://www.java.com/>. Installing JRE installs Java plug-ins, which once installed, load automatically and transparently to render Java applets.



**NOTE:** By default Mozilla Firefox has blocked JRE versions earlier than 1.6.0\_31 and 1.7.0 through 1.7.0\_2. However, Mozilla Firefox users can still click **Add-ons > Plugin** to enable Java.

- Set your browser to support and enable Java applets. To know more about checking the status of java applets in your browser see [http://java.com/en/download/help/enable\\_browser.xml](http://java.com/en/download/help/enable_browser.xml).

---

### Setting Domain Name, Hostname, and Name Server

Configure the domain name and hostname of the switch on your system. Ensure that the DNS server setting is correct. DNS name resolution must happen properly. Ensure that there is connectivity between the client and the management device.

You can set the domain name, hostname, and the DNS name server either through the J-Web interface or the CLI:

- To set through the J-Web interface:

See “Configuring System Identity for an EX Series Switch (J-Web Procedure)” on [page 36](#) for more information.

- To set through the CLI:

```
set system domain-name domain-name
```

```
set system host-name host-name
```

```
set system name-server dns-ip-address
```

---

### Enabling SSH on your system

SSH provides a secure method of logging in to the switch, and encrypting traffic so that it is not intercepted. If SSH is not enabled on the system, the CLI terminal page displays the error message:

To enable SSH on your system, do the following:

```
set system services ssh
```

---

### Sample Configuration on an EX Series Switch

1. Type the **configure** command to enter the configuration mode:

```
user@switch> configure
```

2. Log in as host:

```
user@switch# set system host-name host
```

3. Configure the encrypted password; for example:

```
user@switch# set system root-authentication encrypted-password  
"$1$mr3D4eVf$mc7y54e6hk4JulpwWPao6."
```

4. Map the hostname to the IP address:

```
user@switch# set system static-host-mapping host inet 10.9.221.31
```

5. Configure the IP address for the DNS server:

```
user@switch# set system name-server 10.0.220.1
```

6. Enable the system services by using:

```
set system services:user@switch# set system services ssh
```

7. Select **Troubleshoot > CLI Terminal**. The password window is displayed.

8. Enter the password, and click **OK**. The CLI Terminal window appears on the J-Web page.



**NOTE:** If you exit from the CLI terminal, the connection is lost. Click **CLI Terminal** if you want to connect again.

#### Related Documentation

- [CLI User Interface Overview on page 313](#)
- [Understanding J-Web Configuration Tools](#)

## Statement Hierarchies

- [\[edit access\] Configuration Statement Hierarchy on EX Series Switches on page 322](#)
- [\[edit accounting-options\] Configuration Statement Hierarchy on EX Series Switches on page 325](#)
- [\[edit chassis\] Configuration Statement Hierarchy on EX Series Switches on page 327](#)
- [\[edit class-of-service\] Configuration Statement Hierarchy on EX Series Switches on page 328](#)
- [\[edit ethernet-switching-options\] Configuration Statement Hierarchy on EX Series Switches on page 331](#)
- [\[edit event-options\] Configuration Statement Hierarchy on EX Series Switches on page 334](#)
- [\[edit firewall\] Configuration Statement Hierarchy on EX Series Switches on page 336](#)
- [\[edit forwarding-options\] Configuration Statement Hierarchy on EX Series Switches on page 338](#)
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)
- [\[edit interfaces ae\] Configuration Statement Hierarchy on EX Series Switches on page 343](#)
- [\[edit interfaces ge\] Configuration Statement Hierarchy on EX Series Switches on page 346](#)
- [\[edit interfaces gr\] Configuration Statement Hierarchy on EX Series Switches on page 350](#)

- [\[edit interfaces interface-range\] Configuration Statement Hierarchy on EX Series Switches on page 353](#)
- [\[edit interfaces lo\] Configuration Statement Hierarchy on EX Series Switches on page 360](#)
- [\[edit interfaces me\] Configuration Statement Hierarchy on EX Series Switches on page 363](#)
- [\[edit interfaces vlan\] Configuration Statement Hierarchy on EX Series Switches on page 367](#)
- [\[edit interfaces vme\] Configuration Statement Hierarchy on EX Series Switches on page 370](#)
- [\[edit interfaces xe\] Configuration Statement Hierarchy on EX Series Switches on page 373](#)
- [\[edit poe\] Configuration Statement Hierarchy on EX Series Switches on page 377](#)
- [\[edit policy-options\] Configuration Statement Hierarchy on EX Series Switches on page 378](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)
- [\[edit protocols bfd\] Configuration Statement Hierarchy on EX Series Switches on page 391](#)
- [\[edit protocols bgp\] Configuration Statement Hierarchy on EX Series Switches on page 392](#)
- [\[edit protocols connections\] Configuration Statement Hierarchy on EX Series Switches on page 401](#)
- [\[edit protocols dcbx\] Configuration Statement Hierarchy on EX Series Switches on page 403](#)
- [\[edit protocols dot1x\] Configuration Statement Hierarchy on EX Series Switches on page 404](#)
- [\[edit protocols igmp\] Configuration Statement Hierarchy on EX Series Switches on page 406](#)
- [\[edit protocols igmp-snooping\] Configuration Statement Hierarchy on EX Series Switches on page 407](#)
- [\[edit protocols isis\] Configuration Statement Hierarchy on EX Series Switches on page 408](#)
- [\[edit protocols lacp\] Configuration Statement Hierarchy on EX Series Switches on page 411](#)
- [\[edit protocols link-management\] Configuration Statement Hierarchy on EX Series Switches on page 412](#)
- [\[edit protocols lldp\] Configuration Statement Hierarchy on EX Series Switches on page 413](#)
- [\[edit protocols lldp-med\] Configuration Statement Hierarchy on EX Series Switches on page 415](#)
- [\[edit protocols mld\] Configuration Statement Hierarchy on EX Series Switches on page 416](#)



- [\[edit protocols mld-snooping\] Configuration Statement Hierarchy on EX Series Switches on page 417](#)
- [\[edit protocols mpls\] Configuration Statement Hierarchy on EX Series Switches on page 418](#)
- [\[edit protocols msdp\] Configuration Statement Hierarchy on EX Series Switches on page 429](#)
- [\[edit protocols mstp\] Configuration Statement Hierarchy on EX Series Switches on page 431](#)
- [\[edit protocols mvrp\] Configuration Statement Hierarchy on EX Series Switches on page 433](#)
- [\[edit protocols neighbor-discovery\] Configuration Statement Hierarchy on EX Series Switches on page 434](#)
- [\[edit protocols oam\] Configuration Statement Hierarchy on EX Series Switches on page 435](#)
- [\[edit protocols ospf\] Configuration Statement Hierarchy on EX Series Switches on page 438](#)
- [\[edit protocols ospf3\] Configuration Statement Hierarchy on EX Series Switches on page 441](#)
- [\[edit protocols pim\] Configuration Statement Hierarchy on EX Series Switches on page 444](#)
- [\[edit protocols rip\] Configuration Statement Hierarchy on EX Series Switches on page 447](#)
- [\[edit protocols ripng\] Configuration Statement Hierarchy on EX Series Switches on page 450](#)
- [\[edit protocols router-advertisement\] Configuration Statement Hierarchy on EX Series Switches on page 451](#)
- [\[edit protocols router-discovery\] Configuration Statement Hierarchy on EX Series Switches on page 452](#)
- [\[edit protocols rstp\] Configuration Statement Hierarchy on EX Series Switches on page 453](#)
- [\[edit protocols rsvp\] Configuration Statement Hierarchy on EX Series Switches on page 455](#)
- [\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches on page 459](#)
- [\[edit protocols stp\] Configuration Statement Hierarchy on EX Series Switches on page 460](#)
- [\[edit protocols uplink-failure-detection\] Configuration Statement Hierarchy on EX Series Switches on page 461](#)
- [\[edit protocols vrrp\] Configuration Statement Hierarchy on EX Series Switches on page 462](#)
- [\[edit protocols vstp\] Configuration Statement Hierarchy on EX Series Switches on page 463](#)

- [\[edit redundant-power-system\] Configuration Statement Hierarchy on EX Series Switches on page 465](#)
- [\[edit routing-options\] Configuration Statement Hierarchy on EX Series Switches on page 465](#)
- [\[edit security\] Configuration Statement Hierarchy on EX Series Switches on page 482](#)
- [\[edit services\] Configuration Statement Hierarchy on EX Series Switches on page 485](#)
- [\[edit snmp\] Configuration Statement Hierarchy on EX Series Switches on page 488](#)
- [\[edit system\] Configuration Statement Hierarchy on EX Series Switches on page 495](#)
- [\[edit virtual-chassis\] Configuration Statement Hierarchy on page 508](#)
- [\[edit vlans\] Configuration Statement Hierarchy on EX Series Switches on page 509](#)

## **[edit access] Configuration Statement Hierarchy on EX Series Switches**

This topic lists supported and unsupported configuration statements in the **[edit access]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit access\] Hierarchy Level on page 322](#)
- [Unsupported Statements in the \[edit access\] Hierarchy Level on page 324](#)

### **Supported Statements in the [edit access] Hierarchy Level**

---

The following hierarchy shows the **[edit access]** configuration statements supported on EX Series switches:

```
access {  
  address-assignment {  
    abated-utilization percentage;  
    abated-utilization-v6 percentage;  
    high-utilization percentage;  
    high-utilization-v6 percentage;  
    pool pool-name {  
      family inet {  
        dhcp-attributes {  
          boot-file filename;  
          boot-server hostname;  
          domain-name domain-name;  
          grace-period seconds;  
          maximum-lease-time (seconds | infinite);  
          name-server {  
            address;  
          }  
        }  
      }  
    }  
  }  
}
```

```

netbios-node-type (b-node | h-node | m-node | p-node);
option option-index (array (byte | flag | integer | ip-address | short | string |
    unsigned-integer | unsigned-short) [ type-values ] | byte 8-bit-value |
    flag (false | off | on | true) | integer signed-32-bit-value | ip-address address |
    short signed-16-bit-value | string text-string | unsigned-integer 32-bit-value |
    unsigned-short 16-bit-value);
router {
    address;
}
server-identifier ipv4-address;
tftp-server hostname;
wins-server {
    address;
}
}
host hostname {
    hardware-address mac-address;
    ip-address ip-address;
}
network ip-prefix </prefix-length>;
range name {
    high upper-limit;
    low lower-limit;
}
}
link pool-name;
}
}
address-pool pool-name {
    address address-or-prefix;
    address-range <low lower-limit> <high upper-limit>;
}
profile profile-name {
    accounting (Access Profile) {
        accounting-stop-on-access-deny;
        accounting-stop-on-failure;
        coa-immediate-update;
        immediate-update;
        order (radius | none);
        statistics (time | volume-time);
    }
}
authentication-order (ldap | password | radius);
client client-name {
    chap-secret chap-secret;
    firewall-user {
        password password;
    }
    no-rfc2486;
    pap-password password;
}
radius {
    accounting-server server-address;
    attributes {
        exclude [exclude-options];
        ignore [ignore-options];
    }
}

```

```

    }
    authentication-server server-address;
  }
  radius-options {
    revert-interval interval;
  }
  session-options {
    client-group [group-names];
    client-idle-timeout minutes;
    client-session-timeout minutes;
  }
  radius-options {
    revert-interval interval;
  }
  radius-server server-address {
    port port-number;
    retry attempts;
    routing-instance instance-name;
    secret password;
    source-address address;
    timeout minutes;
  }
}

```

### Unsupported Statements in the [edit access] Hierarchy Level

All statements in the [edit access] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 24: Unsupported [edit access] Configuration Statements on EX Series Switches**

Statement	Hierarchy Level
<b>NOTE:</b> Variables, such as <i>filename</i> , are not shown in the statements or hierarchies.	
aaa	[edit access terminate-code]
administrative-reset	[edit access terminate-code aaa shutdown]
authentication-denied	[edit access terminate-code aaa deny]
client-request	[edit access terminate-code aaa dhcp]
compliance	[edit access ppp-options]
deny	[edit access terminate-code aaa]
dhcp	[edit access terminate-code]
group-profile	[edit access]
ike	[edit access profile client]

Table 24: Unsupported [edit access] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy Level
initiate-dead-peer-detection	[edit access profile client ike]
lost-carrier	[edit access terminate-code dhcp]
nak	[edit access terminate-code dhcp]
nas-logout	[edit access terminate-code dhcp]
no-offers	[edit access terminate-code dhcp]
no-resources	[edit access terminate-code aaa deny]
ppp-options	[edit access]
preference	[edit access profile client ike reverse-route]
remote-reset	[edit access terminate-code aaa shutdown]
rfc	[edit access ppp-options compliance]
reverse-route	[edit access profile client ike]
server-request-timeout	[edit access terminate-code aaa deny]
shutdown	[edit access terminate-code aaa]
terminate-code	[edit access]

- Related Documentation**
- [Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch on page 1843](#)
  - [Configuring 802.1X RADIUS Accounting \(CLI Procedure\) on page 1909](#)
  - [Security Features for EX Series Switches Overview on page 4693](#)

## [edit accounting-options] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit accounting-options]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit accounting-options\] Hierarchy Level on page 326](#)
- [Unsupported Statements in the \[edit accounting-options\] Hierarchy Level on page 326](#)

### Supported Statements in the [edit accounting-options] Hierarchy Level

The following hierarchy shows the **[edit accounting-options]** configuration statements supported on EX Series switches:

```
accounting-options {  
  class-usage-profile profile-name {  
    destination-classes destination-class-name;  
    file filename;  
    interval minutes;  
    source-classes source-class-name;  
  }  
  file {  
    archive-sites site-url {  
      password password;  
    }  
    files number;  
    size bytes;  
    start-time time;  
    transfer-interval minutes;  
  }  
  filter-profile profile-name {  
    counters counter-name;  
    file filename;  
    interval minutes;  
  }  
  interface-profile profile-name {  
    fields [interface-profile-fields];  
    file file-name;  
    interval minutes;  
  }  
  mib-profile profile-name {  
    file filename;  
    interval minutes;  
    object-names [mib-object-names];  
    operation [snmp-operations];  
  }  
  routing-engine-profile profile-name {  
    fields [field-names];  
    file filename;  
    interval minutes;  
  }  
}
```

### Unsupported Statements in the [edit accounting-options] Hierarchy Level

All statements in the **[edit accounting-options]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

- Related Documentation**
- *Network Management Administration Guide for Routing Devices*
  - *SNMP MIBs and Traps Reference*

## [edit chassis] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit chassis]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit chassis\] Hierarchy Level on page 327](#)

### Supported Statements in the [edit chassis] Hierarchy Level

The following hierarchy shows the **[edit chassis]** configuration statements supported on EX Series switches:

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
      lacp {
        link-protection non-revertive;
        system-priority system-priority-number
      }
    }
  }
  alarm {
    ethernet {
      link-down (ignore | red | yellow);
    }
    management-ethernet {
      link-down (ignore | red | yellow);
    }
  }
  container-devices {
    device-count device-count-number;
  }
  disk-partition {
    /config {
      level (full | high) {
        free-space (free-space-threshold-value | mb | percent);
      }
    }
  }
  /var {
```

```
        level (full | high) {
            free-space (free-space-threshold-value | mb | percent);
        }
    }
}
fpc slot-number {
    pic pic-number {
        no-multi-rate;
        q-pic-large-buffer (large-scale | small-scale);
    }
}
maximum-ecmp maximum-ecmp-routes;
lcd-menu {
    fpc slot-number {
        menu-item menu-name;
        disable;
    }
    pseudowire-service {
        device-count device-count-number;
    }
    psu {
        redundancy {
            n-plus-n;
        }
        redundancy {
            graceful-switchover;
        }
    }
    slow-pfe-alarm;
}
```

**Related Documentation**

- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\) on page 33](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 5116](#)
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 2512](#)
- [Configuring the Power Priority of Line Cards \(CLI Procedure\)](#)
- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade \(CLI Procedure\)](#)

## [edit class-of-service] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit class-of-service]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.



- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit class-of-service\] Hierarchy Level on page 329](#)
- [Unsupported Statements in the \[edit class-of-service\] Hierarchy Level on page 330](#)

### Supported Statements in the [edit class-of-service] Hierarchy Level

The following hierarchy shows the **[edit class-of-service]** configuration statements supported on EX Series switches:

```
class-of-service {
  classifiers {
    (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) {
          code-points [ aliases ] [ 6 bit-patterns ];
        }
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) {
      alias-name bits;
    }
  }
  drop-profiles {
    profile-name {
      interpolate {
        drop-probability [values];
        fill-level [values]
      }
    }
  }
  forwarding-classes {
    class class-name
    queue queue-number;
  }
  interfaces interface-name {
    scheduler-map map-name;
    shaping-rate rate;
    unit (logical-unit-number | *) {
      classifiers {
        (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) (classifier-name | default);
      }
      forwarding-class class-name ;
    }
  }
  rewrite-rules {
    (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) (rewrite-rule-name | default);
  }
}
```

```

rewrite-rules {
  (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence ) rewrite-name {
    import (default | rewrite-name);
    forwarding-class class-name {
      loss-priority (high | low | medium-high | medium-low) code-point (alias | bits);
    }
  }
}
scheduler-maps {
  map-name {
    forwarding-class class-name {
      scheduler scheduler-name;
    }
  }
}
schedulers {
  scheduler-name {
    buffer-size (exact | percent percentage | remainder);
    drop-profile-map {
      loss-priority (any | high | medium-high | medium-low);
      protocol any;
      {
        drop-profile profile-name
      }
    }
    excess-rate {
      percent percentage;
    }
    priority (low | strict-high);
    shaping-rate (rate | percent percentage);
    transmit-rate (EX Series Switches) (rate | percent percentage | remainder) ;
  }
}
shared-buffer {
  percent;
}
traceoptions {
  file (file-name | files files | match match | no-world-readable | size size | world-readable);
  flag ( all | asynch | chassis-scheduler | cos-adjustment | dynamic | hardware-database
    | init | parse | performance-monitor | process | restart | route-socket | show | snmp |
    util);
  no-remote-trace;
}
tri-color;
}

```

### Unsupported Statements in the [edit class-of-service] Hierarchy Level

All statements in the [edit class-of-service] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Defining CoS Code-Point Aliases \(CLI Procedure\) on page 2101](#) or [Defining CoS Code-Point Aliases \(J-Web Procedure\) on page 2101](#)

- [Defining CoS Classifiers \(CLI Procedure\) on page 2103](#) or [Defining CoS Classifiers \(J-Web Procedure\) on page 2105](#)
- [Defining CoS Forwarding Classes \(CLI Procedure\) on page 2107](#) or [Defining CoS Forwarding Classes \(J-Web Procedure\) on page 2107](#)
- [Configuring CoS Tail Drop Profiles \(CLI Procedure\)](#)
- [Defining CoS Schedulers and Scheduler Maps \(CLI Procedure\) on page 2109](#) or [Defining CoS Schedulers \(J-Web Procedure\) on page 2111](#)
- [Defining CoS Rewrite Rules \(CLI Procedure\) on page 2120](#) or [Defining CoS Rewrite Rules \(J-Web Procedure\) on page 2121](#)

## [\[edit ethernet-switching-options\]](#) Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [\[edit ethernet-switching-options\]](#) hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit ethernet-switching-options\] Hierarchy Level on page 331](#)
- [Unsupported Statements in the \[edit ethernet-switching-options\] Hierarchy Level on page 334](#)

### [Supported Statements in the \[edit ethernet-switching-options\] Hierarchy Level](#)

The following hierarchy shows the [\[edit ethernet-switching-options\]](#) configuration statements supported on EX Series switches:

```
ethernet-switching-options {
  analyzer {
    name {
      input {
        egress {
          interface (all | interface-name);
        }
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
      }
    }
    loss-priority priority;
    output {
      interface interface-name;
```

```
        vlan (vlan-id | vlan-name);
    }
    ratio number;
}
}
authentication-whitelist {
    interface;
    vlan-assignment;
}
bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]) {
        (disable | drop | shutdown);
    }
}
dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
}
interfaces interface-name {
    no-mac-learning;
}
mac-lookup-length number-of-entries;
}
mac-notification {
    notification-interval seconds;
}
mac-table-aging-time seconds;
port-error-disable {
    disable-timeout timeout;
}
redundant-trunk-group {
    group name {
        description;
        interface interface-name {
            primary;
        }
        preempt-cutover-timer seconds;
    }
}
secure-access-port {
    dhcp-snooping-file {
        location local_pathname | remote_URL;
        timeout seconds;
        write-interval seconds;
    }
    interface (all | interface-name) {
        allowed-mac {
            mac-address-list;
        }
        (dhcp-trusted | no-dhcp-trusted );
        fcoe-trusted;
        mac-limit limit action action;
        no-allowed-mac-log;
        static-ip ip-address {
            mac mac-address;
            vlan vlan-name;
        }
    }
}
```

```

    }
  }
  uac-policy;
}
vlan (all | vlan-name) {
  (arp-inspection | no-arp-inspection );
  dhcp-option82 {
    disable;
    circuit-id {
      prefix hostname;
      use-interface-description;
      use-vlan-id;
    }
    remote-id {
      prefix (hostname | mac | none);
      use-interface-description;
      use-string string;
    }
    vendor-id [string];
  }
  (examine-dhcp | no-examine-dhcp);
  examine-fip {
    fc-map fc-map-value;
  }
  (ip-source-guard | no-ip-source-guard);
  mac-move-limit limit action action;
}
}
static {
  vlan vlan-id {
    mac mac-address next-hop interface-name;
  }
}
storm-control {
  action-shutdown;
  interface (all | interface-name) {
    bandwidth bandwidth;
    multicast;
    no-broadcast;
    no-multicast;
    no-registered-multicast;
    no-unknown-unicast;
    no-unregistered-multicast;
  }
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
}
unknown-unicast-forwarding {
  vlan (all | vlan-name) {
    interface interface-name;
  }
}
}
voip {

```

```
interface (all | [interface-name | access-ports]) {
    forwarding-class (assured-forwarding | best-effort | expedited-forwarding |
        network-control);
    vlan vlan-name;
}
}
```

### Unsupported Statements in the [edit ethernet-switching-options] Hierarchy Level

All statements in the **[edit ethernet-switching-options]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- *Example: Setting Up Q-in-Q Tunneling on EX Series Switches*
- *Example: Configuring Redundant Trunk Links for Faster Recovery*
- *Configuring MAC Table Aging (CLI Procedure)*
- *Configuring MAC Notification (CLI Procedure)*
- *Configuring Q-in-Q Tunneling (CLI Procedure)*
- *Configuring Redundant Trunk Links for Faster Recovery (CLI Procedure)*
- *Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)*

### **[edit event-options] Configuration Statement Hierarchy on EX Series Switches**

This topic lists supported and unsupported configuration statements in the **[edit event-options]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit event-options\] Hierarchy Level on page 334](#)
- [Unsupported Statements in the \[edit event-options\] Hierarchy Level on page 336](#)

### Supported Statements in the [edit event-options] Hierarchy Level

The following hierarchy shows the **[edit event-options]** configuration statements supported on EX Series switches:

```
event-options {
    destinations {
        destination-name {
            archive-sites {
                url <password password>;
            }
        }
    }
}
```

```

    }
    transfer-delay seconds;
  }
}
event-script {
  file filename {
    checksum (md5 | sha-256 | sha1) hash;
    refresh;
    refresh-from url;
    remote-execution {
      remote-hostname {
        passphrase user-password;
        username user-login;
      }
    }
    source url;
  }
  refresh;
  refresh-from url;
  traceoptions {
    file <filename> <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
generate-event event-name {
  time-interval seconds;
  time-of-day hh:mm:ss;
}
policy policy-name {
  attributes-match {
    event1.attribute-name equals event2.attribute-name;
    event.attribute-name matches regular-expression;
    event1.attribute-name starts-with event2.attribute-name;
  }
  events [ events ];
  then {
    event-script filename {
      arguments {
        argument-name argument-value;
      }
      destination destination-name {
        retry-count number retry-interval seconds;
        transfer-delay seconds;
      }
      output-filename filename;
      output-format (text | xml);
      user-name username;
    }
    execute-commands {
      commands {
        "command";
      }
      destination destination-name {
        retry-count number retry-interval seconds;

```

```
        transfer-delay seconds;
    }
    output-filename filename;
    output-format (text | xml);
    user-name username;
}
ignore;
raise-trap;
upload filename (filename | committed) destination destination-name {
    retry-count number retry-interval seconds;
    transfer-delay seconds;
    user-name username;
}
}
within seconds {
    events [ events ];
    not events [ events ];
    trigger (after number | on number | until number);
}
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
}
```

#### Unsupported Statements in the [edit event-options] Hierarchy Level

---

All statements in the [edit event-options] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

**Related Documentation**

- [Event Scripts Overview](#)

### [edit firewall] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit firewall] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit firewall\] Hierarchy Level on page 337](#)
- [Unsupported Statements in the \[edit firewall\] Hierarchy Level on page 337](#)



### Supported Statements in the [edit firewall] Hierarchy Level

The following hierarchy shows the [edit firewall] configuration statements supported on EX Series switches:

```

firewall {
  family family-name {
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
  policer policer-name {
    filter-specific;
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      policer-action;
    }
  }
  three-color-policer policer-name {
    action {
      loss-priority high then discard;
    }
    filter-specific;
    single-rate {
      (color-aware | color-blind);
      committed-burst-size bytes;
      committed-information-rate bps;
      excess-burst-size bytes;
    }
    two-rate {
      (color-aware | color-blind);
      committed-burst-size bytes;
      committed-information-rate bps;
      peak-information-rate bps;
      peak-burst-size bytes;
    }
  }
}

```

### Unsupported Statements in the [edit firewall] Hierarchy Level

All statements in the [edit firewall] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

- Related Documentation**
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
  - [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
  - [Configuring Policers to Control Traffic Rates \(CLI Procedure\) on page 4818](#)
  - [Firewall Filter Configuration Statements Supported by Junos OS for EX Series Switches on page 4832](#)
  - [Firewall Filters for EX Series Switches Overview on page 4696](#)

## [\[edit forwarding-options\]](#) Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [\[edit forwarding-options\]](#) hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit forwarding-options\] Hierarchy Level on page 338](#)
- [Unsupported Statements in the \[edit forwarding-options\] Hierarchy Level on page 340](#)

### Supported Statements in the [\[edit forwarding-options\]](#) Hierarchy Level

The following hierarchy shows the [\[edit forwarding-options\]](#) configuration statements supported on EX Series switches:

```
forwarding-options {
  dhcp-relay {
    group group-name {
      interface interface-name {
        overrides {
          always-write-giaddr;
          always-write-option-82;
          client-discover-match <option60-and-option82>;
          interface-client-limit number;
          layer2-unicast-replies;
          no-arp;
          trust-option-82;
        }
      }
    }
    exclude {
      overrides {
        ...
      }
    }
  }
}
```

```

        trace;
        upto upto-interface-name;
    }
    overrides {
        ...
    }
    relay-option {
        ...
    }
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
}
helpers{
    bootp {
        client-response-ttl number;
        description text-description;
        dhcp-option82 {
            circuit-id {
                prefix (Circuit ID for Option 82) hostname;
                use-interface-description;
                use-vlan-id;
            }
            disable;
            remote-id {
                prefix hostname | mac | none;
                use-interface-description;
                use-string string;
            }
            vendor-id <string>;
        }
    }
    interface (interface-name | interface-group) {
        broadcast;
        client-response-ttl number;
        description text-description;
        dhcp-option82 {
            circuit-id {
                prefix (Circuit ID for Option 82) hostname;
                use-interface-description;
                use-vlan-id;
            }
            disable;
            remote-id {
                prefix hostname | mac | none;
                use-interface-description;
                use-string string;
            }
        }
    }
}

```

```

        vendor-id <string>;
    }
    maximum-hop-count number;
    minimum-wait-time seconds;
    no-listen;
    server address {
        routing-instance [ routing-instance-names ];
    }
}
maximum-hop-count number;
minimum-wait-time seconds;
no-listen;
relay-agent-option;
server address {
    routing-instance [ routing-instance-names ];
}
source-address-giaddr;
}
}

```

### Unsupported Statements in the [edit forwarding-options] Hierarchy Level

All statements in the [edit forwarding-options] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 25: Unsupported [edit forwarding-options] Configuration Statements on EX Series Switches**

Statement	Hierarchy Level
<i>NOTE:</i> Variables, such as <i>filename</i> , are not shown in the statements or hierarchies.	
accounting	[edit forwarding-options]
aggregate-export-interval	[edit forwarding-options accounting output]
broadcast	[edit forwarding-options helpers domain interface] [edit forwarding-options helpers port interface] [edit forwarding-options helpers tftp interface]
description	[edit forwarding-options helpers domain] [edit forwarding-options helpers domain interface] [edit forwarding-options helpers port interface] [edit forwarding-options helpers tftp] [edit forwarding-options helpers tftp interface]
domain	[edit forwarding-options helpers]
engine-id	[edit forwarding-options accounting output interface]
file	[edit forwarding-options helpers traceoption]
flag	[edit forwarding-options helpers traceoption]

Table 25: Unsupported [edit forwarding-options] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy Level
flow-active-timeout	[edit forwarding-options accounting output]
flow-inactive-timeout	[edit forwarding-options accounting output]
hash-seed	[edit forwarding-options load-balance per-prefix]
indexed-next-hop	[edit forwarding-options load-balance]
interface	[edit forwarding-options accounting output] [edit forwarding-options helpers domain] [edit forwarding-options helpers port] [edit forwarding-options helpers tftp]
level	[edit forwarding-options helpers traceoption]
load-balance	[edit forwarding-options]
no-listen	[edit forwarding-options helpers domain interface] [edit forwarding-options helpers port interface] [edit forwarding-options helpers tftp interface]
no-remote-trace	[edit forwarding-options helpers traceoption]
output	[edit forwarding-options accounting]
per-prefix	[edit forwarding-options load-balance]
port	[edit forwarding-options helpers]
routing-instance	[edit forwarding-options helpers domain interface server] [edit forwarding-options helpers domain server] [edit forwarding-options helpers port interface server] [edit forwarding-options helpers rtsdb-client-traceoptions] [edit forwarding-options helpers tftp interface server] [edit forwarding-options helpers tftp server]
rtsdb-client-traceoptions	[edit forwarding-options helpers]
server	[edit forwarding-options helpers domain] [edit forwarding-options helpers domain interface] [edit forwarding-options helpers port] [edit forwarding-options helpers port interface] [edit forwarding-options helpers tftp] [edit forwarding-options helpers tftp interface]
source-address	[edit forwarding-options accounting output interface]
tftp	[edit forwarding-options helpers]

Table 25: Unsupported [edit forwarding-options] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy Level
traceoptions	[edit forwarding-options helpers]
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i></li> <li>• <i>Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li> <li>• <i>DHCP/BOOTP Relay for Switches Overview</i></li> <li>• <i>For more information about the [edit forwarding-options] hierarchy and its options, see Junos OS Policy Framework Configuration Guide</i></li> </ul>

## [edit interfaces] Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the [edit interfaces] hierarchy:

- [edit interfaces ae] Configuration Statement Hierarchy on EX Series Switches on page 343
- [edit interfaces ge] Configuration Statement Hierarchy on EX Series Switches on page 346
- [edit interfaces gr] Configuration Statement Hierarchy on EX Series Switches on page 350
- [edit interfaces interface-range] Configuration Statement Hierarchy on EX Series Switches on page 353
- [edit interfaces lo] Configuration Statement Hierarchy on EX Series Switches on page 360
- [edit interfaces me] Configuration Statement Hierarchy on EX Series Switches on page 363
- [edit interfaces vlan] Configuration Statement Hierarchy on EX Series Switches on page 367
- [edit interfaces vme] Configuration Statement Hierarchy on EX Series Switches on page 370
- [edit interfaces xe] Configuration Statement Hierarchy on EX Series Switches on page 373

### Related Documentation

- *EX Series Switches Interfaces Overview on page 2577*
- *Configuring Aggregated Ethernet Links (CLI Procedure) on page 2667*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring a Layer 3 Subinterface (CLI Procedure) on page 2689*

- *Configuring Routed VLAN Interfaces (CLI Procedure)*
- *Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)*
- [Junos OS Interfaces Fundamentals Configuration Guide](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

## [edit interfaces ae] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces ae]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces ae\] Hierarchy Level on page 343](#)
- [Unsupported Statements in the \[edit interfaces ae\] Hierarchy Level on page 346](#)

### Supported Statements in the [edit interfaces ae] Hierarchy Level

The following hierarchy shows the **[edit interfaces ae]** configuration statements supported on EX Series switches.

```

interfaces {
  ae-fpc/pic/port {
    accounting-profile name;
    aggregated-ether-options {
      ethernet-switch-profile {
        tag-protocol-id identifier;
      }
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        admin-key key;
        periodic interval;
        system-id mac-address;
      }
      (link-protection | no-link-protection);
      link-speed speed;
      (loopback | no-loopback);
      minimum-links number;
    }
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
  }
}

```

```
mtu bytes;  
no-gratuitous-arp-request;  
traceoptions {  
    flag flag;  
}  
(traps | no-traps);  
unit logical-unit-number {  
    accounting-profile name;  
    arp-resp;  
    bandwidth rate;  
    description text;  
    disable;  
    family ccc;  
    family ethernet-switching {  
        filter {  
            input filter-name;  
            output filter-name;  
        }  
        native-vlan-id vlan-id-number;  
        port-mode (access | trunk);  
        vlan {  
            members [ members];  
        }  
    }  
}  
family inet {  
    address ipv4-address {  
        arp ip-address (mac | multicast-mac) mac-address <publish>;  
        broadcast address;  
        preferred;  
        primary;  
        vrrp-group group-number {  
            (accept-data | no-accept-data);  
            advertise-interval seconds;  
            authentication-key key;  
            authentication-type authentication;  
            fast-interval milliseconds;  
            (preempt | no-preempt) {  
                hold-time seconds;  
            }  
            priority number;  
            track {  
                interface interface-name {  
                    bandwidth-threshold bandwidth;  
                    priority-cost number;  
                }  
                priority-hold-time seconds;  
                route ip-address/mask routing-instance instance-name priority-cost cost;  
            }  
            virtual-address [ addresses ];  
            virtual-link-local-address address;  
            vrrp-inherit-from {  
                active-group group-number;  
                active-interface interface-name;  
            }  
        }  
    }  
}
```



```

dhcp {
  client-identifier (ascii client-id | hexadecimal client-id);
  lease-time (seconds | infinte);
  retransmission-attempt number;
  retransmission-interval sections;
  server-address ip-address;
  update-server server;
  vendor-id id;
}
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check;
targeted-broadcast;
}
family inet6 {
  address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
      accept-data | no-accept-data;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds;
      preempt | no-preempt {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          bandwidth-threshold bandwidth priority-cost number;
          priority-cost number;
        }
        priority-hold-time seconds;
        route ( address | routing-instance routing-instance-name );
      }
      virtual-inet6-address [addresses];
      virtual-link-local-address ipv6-address;
    }
    vrrp-inherit-from {
      active-group group-name;
      active-interface interface-name;
    }
  }
}
(dad-disable | no-dad-disable);
filter {
  input filter-name;
  output filter-name;
}

```

```

    }
    mtu bytes;
    no-neighbor-learn;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
}
vlan-tagging;
}
}

```

### Unsupported Statements in the [edit interfaces ae] Hierarchy Level

All statements in the [edit interfaces ae] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 26: Unsupported [edit interfaces ae] Configuration Statements on EX Series Switches**

Statement	Hierarchy
NOTE: Variables, such as <i>interface-range</i> , are not shown in the statements or hierarchies.	
family fibre-channel	[edit interfaces ae unit]
source-address-filter	[edit interfaces ae aggregated-ether-options]
source-address-filtering   no-source-address-filtering	[edit interfaces ae aggregated-ether-options]

**Related Documentation**

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)

### [edit interfaces ge] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit interfaces ge] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.

- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces ge\] Hierarchy Level on page 347](#)
- [Unsupported Statements in the \[edit interfaces ge\] Hierarchy Level on page 350](#)

### [Supported Statements in the \[edit interfaces ge\] Hierarchy Level](#)

The following hierarchy shows the **[edit interfaces ge]** configuration statements supported on EX Series switches.

```

interfaces {
  ge-fpc/pic/port {
    accounting-profile name;
    description text;
    disable;
    ether-options {
      802.3ad {
        aex;
        (backup | primary);
        lacp {
          force-up;
        }
      }
      (auto-negotiation | no-auto-negotiation);
      (flow-control | no-flow-control);
      ieee-802-3az-eee;
      link-mode mode;
      (loopback | no-loopback);
      speed (auto-negotiation | speed);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    optics-options {
      alarm alarm-type;
      warning alarm-type;
      wavelength nanometers;
    }
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
    }
  }
}

```

```
family ccc;
family ethernet-switching {
  filter {
    input filter-name;
    output filter-name;
  }
  native-vlan-id vlan-id-number;
  port-mode (access | trunk);
  vlan {
    members [ members ];
  }
}
family inet {
  address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    preferred;
    primary;
    vrrp-group group-number {
      (accept-data | no-accept-data);
      advertise-interval seconds;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      (preempt | no-preempt) {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          bandwidth-threshold bandwidth;
          priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
      }
      virtual-address [ addresses ];
      virtual-link-local-address address;
      vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
      }
    }
  }
}
dhcp {
  client-identifier (ascii client-id | hexadecimal client-id);
  lease-time (seconds | infinite);
  retransmission-attempt number;
  retransmission-interval sections;
  server-address ip-address;
  update-server
  vendor-id
}
filter {
  input filter-name;
  output filter-name;
```

```

}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check;
targeted-broadcast;
}
family inet6 {
  address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
      accept-data | no-accept-data;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds;
      preempt | no-preempt {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          bandwidth-threshold bandwidth priority-cost number;
          priority-cost number;
        }
        priority-hold-time seconds;
        route ( address | routing-instance routing-instance-name );
      }
      virtual-inet6-address [addresses];
      virtual-link-local-address ipv6-address;
      vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
      }
    }
  }
}
(dad-disable | no-dad-disable);
filter {
  group group-name;
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
  input policer-name;
  output policer-name;
}
rpf-check;
}
family iso {
  address interface-address;

```

```
        mtu bytes;
    }
    family mpls {
        mtu bytes;
    }
    proxy-arp (restricted | unrestricted);
    swap-by-poppush;
    (traps | no-traps);
    vlan-id vlan-id-number;
}
vlan-tagging;
}
```

---

### Unsupported Statements in the [edit interfaces ge] Hierarchy Level

All statements in the **[edit interfaces ge]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)

## [edit interfaces gr] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces gr]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces gr\] Hierarchy Level on page 350](#)
- [Unsupported Statements in the \[edit interfaces gr\] Hierarchy Level on page 353](#)

---

### Supported Statements in the [edit interfaces gr] Hierarchy Level

The following hierarchy shows the **[edit interfaces gr]** configuration statements supported on EX Series switches.

```
interfaces {
  gr-fpc/pic/port {
    accounting-profile name;
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    traceoptions {
      flag flag;
    }
  }
}
```

```

}
(traps | no-traps);
unit logical-unit-number {
  accounting-profile name;
  bandwidth rate;
  description text;
  disable;
  family ccc;
  family inet {
    accounting {
      destination-class-usage;
      source-class-usage {
        direction;
      }
    }
  }
  address ipv4-address {
    destination address;
    preferred;
    primary;
    vrrp-group group-number {
      (accept-data | no-accept-data);
      advertise-interval seconds;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      (preempt | no-preempt) {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          bandwidth-threshold bandwidth;
          priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
      }
      virtual-address [ addresses ];
      virtual-link-local-address address;
      vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
      }
    }
  }
}
dhcp {
  client-identifier (ascii client-id | hexadecimal client-id);
  lease-time (seconds | infinite);
  retransmission-attempt number;
  retransmission-interval sections;
  server-address ip-address;
  update-server
  vendor-id
}
filter {
  input filter-name;

```

```
        output filter-name;
    }
    mtu bytes;
    no-neighbor-learn;
    no-redirects;
    primary;
    rpf-check;
    targeted-broadcast;
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            direction;
        }
    }
}
address address {
    destination address;
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ( address | routing-instance routing-instance-name );
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-name;
            active-interface interface-name;
        }
    }
}
(dad-disable | no-dad-disable);
filter {
    group group-name;
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
```



```

    no-neighbor-learn;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
(traps | no-traps);
tunnel {
    destination destination-address;
    flow-label label;
    source source-address;
    ttl number;
}
}
}
}

```

#### Unsupported Statements in the [edit interfaces gr] Hierarchy Level

All statements in the [edit interfaces gr] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)

### [edit interfaces interface-range] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit interfaces interface-range] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces interface-range\] Hierarchy Level on page 354](#)
- [Unsupported Statements in the \[edit interfaces interface-range\] Hierarchy Level on page 357](#)

### Supported Statements in the [edit interfaces interface-range] Hierarchy Level

The following hierarchy shows the [edit interfaces interface-range] configuration statements supported on EX Series switches.

```
interfaces {
  interface-range name {
    accounting-profile name;
    aggregated-ether-options {
      ethernet-switch-profile {
        tag-protocol-id identifier;
      }
      (flow-control | no-flow-control);
      ieee-802-3az-eee;
      lacp {
        (active | passive);
        admin-key key;
        periodic interval;
        system-id mac-address;
      }
      (link-protection | no-link-protection);
      link-speed speed;
      (loopback | no-loopback);
      minimum-links number;
      rebalance-periodic;
      source-address-filter filter;
      source-filtering | no-source-filtering;
    }
    description text;
    disable;
    ether-options {
      802.3ad {
        aex;
        (backup | primary);
        lacp {
          force-up;
        }
      }
      (auto-negotiation | no-auto-negotiation);
      (flow-control | no-flow-control);
      link-mode mode;
      (loopback | no-loopback);
      speed (auto-negotiation | speed);
    }
    framing;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    member interface-name;
    member-range starting-interface name to ending-interface name;
    mtu bytes;
    no-gratuitous-arp-request;
    optics-options {
      alarm alarm-type;
      warning alarm-type;
      wavelength nanometers;
    }
  }
}
```

```

}
services-options;
speed speed;
traceoptions {
    flag flag;
}
(traps | no-traps);
unit logical-unit-number {
    accept-source-mac;
    accounting-profile name;
    arp-resp;
    bandwidth rate;
    description text;
    disable;
    family ccc;
    family ethernet-switching {
        filter {
            input filter-name;
            output filter-name;
        }
        native-vlan-id vlan-id-number;
        port-mode (access | trunk);
        vlan {
            members [ members ];
        }
    }
}
family inet {
    accounting {
        destination-class-usage;
        source-class-usage;
    }
    address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        destination-class-usage;
        destination-profile;
        master-only;
        preferred;
        primary;
        vrrp-group group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bandwidth;
                    priority-cost number;
                }
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;

```

```

    }
    virtual-address [ addresses ];
    virtual-link-local-address address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
}
dhcp {
    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinite);
    retransmission-attempt number;
    retransmission-interval sections;
    server-address ip-address;
    update-server
    vendor-id
}
filter {
    input filter-name;
    output filter-name;
}
ipsec-sa;
mtu bytes;
multicast-only;
negotiate-address;
next-hop-tunnel;
no-neighbor-learn;
no-redirects;
primary;
receive-option-packets;
rpf-check;
targeted-broadcast;
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage;
    }
    address address {
        eui-64;
        ndp ip-address (mac | multicast-mac) mac-address <publish>;
        preferred;
        primary;
        vrrp-inet6-group group-id {
            accept-data | no-accept-data;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            inet6-advertise-interval milliseconds;
            preempt | no-preempt {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {

```

```

        bandwidth-threshold bandwidth priority-cost number;
        priority-cost number;
    }
    priority-hold-time seconds;
    route ( address | routing-instance routing-instance-name );
}
virtual-inet6-address [addresses];
virtual-link-local-address ipv6-address;
}
vrrp-inherit-from {
    active-group group-name;
    active-interface interface-name;
}
}
(dad-disable | no-dad-disable);
filter {
    group group-name;
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
interface-shared-with;
interleave-fragments;
inverse-arp;
link-layer-overhead;
minimum-links;
mtu;
proxy-arp (restricted | unrestricted);
swap-by-poppush;
(traps | no-traps);
vlan-id vlan-id-number;
}
vlan-tagging;
}

```

### Unsupported Statements in the [edit interfaces interface-range] Hierarchy Level

All statements in the [edit interfaces interface-range] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 27: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches**

Statement	Hierarchy
<b>NOTE:</b> Variables, such as <i>interface-range</i> , are not shown in the statements or hierarchies.	
aggregated-sonet-options and all substatements	[edit interfaces interface-range]
allow-any-vci	[edit interfaces interface-range unit]
atm-l2circuit-mode	[edit interfaces interface-range unit]
atm-options and all substatements	[edit interfaces interface-range]
cell-bundle-size	[edit interfaces interface-range unit]
clear-don't-fragment-bit	[edit interfaces interface-range unit]
clocking	[edit interfaces interface-range]
compression-device	[edit interfaces interface-range unit]
container-options and all substatements	[edit interfaces interface-range]
copy-tos-to-outer-ip-header	[edit interfaces interface-range unit]
dce	[edit interfaces interface-range]
disable-mlppp-inter-ppp-pfc	[edit interfaces interface-range unit]
dlci	[edit interfaces interface-range unit]
drop-timeout	[edit interfaces interface-range unit]
ds0-options and all substatements	[edit interfaces interface-range]
e1-options and all substatements	[edit interfaces interface-range]
e3-options and all substatements	[edit interfaces interface-range]
epd-threshold	[edit interfaces interface-range unit]
family mlfr-end-to-end and all substatements	[edit interfaces interface-range unit]
family mlfr-uni-uni and all substatements	[edit interfaces interface-range unit]
family mlppp and all substatements	[edit interfaces interface-range unit]
fragment-threshold	[edit interfaces interface-range unit]

Table 27: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches (*continued*)

Statement	Hierarchy
ggsn-options and all substatements	[edit interfaces interface-range]
keepalives   no-keepalives	[edit interfaces interface-range] [edit interfaces interface-range unit]
lmi	[edit interfaces interface-range]
lsq-failure-options	[edit interfaces interface-range]
mlfr-uni-nni-bundle-options and all substatements	[edit interfaces interface-range]
mrru	[edit interfaces interface-range unit]
multicast-dlci	[edit interfaces interface-range unit]
multilink-max-classes	[edit interfaces interface-range unit]
multipoint	[edit interfaces interface-range unit]
multipoint-destination	[edit interfaces interface-range unit family inet address]
multiservice-options and all substatements	[edit interfaces interface-range]
oam-liveness	[edit interfaces interface-range unit]
oam-period	[edit interfaces interface-range unit]
passive-monitor-mode	[edit interfaces interface-range unit]
peer-unit	[edit interfaces interface-range unit]
plp-to-clp	[edit interfaces interface-range unit]
point-to-point	[edit interfaces interface-range unit]
ppp-options and all substatements	[edit interfaces interface-range] [edit interfaces interface-range unit]
receive-lsp	[edit interfaces interface-range unit]
satop-options and all substatements	[edit interfaces interface-range]
serial-options and all substatements	[edit interfaces interface-range]
service-domain	[edit interfaces interface-range unit]

Table 27: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches (*continued*)

Statement	Hierarchy
shaping	[edit interfaces interface-range unit]
short-sequence	[edit interfaces interface-range unit]
shdsl-options and all substatements	[edit interfaces interface-range]
t1-options and all substatements	[edit interfaces interface-range]
t3-options and all substatements	[edit interfaces interface-range]
transmit-lsp	[edit interfaces interface-range unit]
transmit-weight	[edit interfaces interface-range unit]
trunk-id	[edit interfaces interface-range unit]
tunnel	[edit interfaces interface-range unit]
vci	[edit interfaces interface-range unit]
vci-range	[edit interfaces interface-range unit]
vpi	[edit interfaces interface-range unit]

**Related Documentation** • [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)

## [edit interfaces lo] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces lo]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces lo\] Hierarchy Level on page 361](#)
- [Unsupported Statements in the \[edit interfaces lo\] Hierarchy Level on page 363](#)



### Supported Statements in the [edit interfaces lo] Hierarchy Level

The following hierarchy shows the [edit interfaces lo] configuration statements supported on EX Series switches.

```

interfaces {
  lo0 {
    accounting-profile name;
    description text;
    disable;
    hold-time down milliseconds up milliseconds ;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family ccc;
      family inet {
        address ipv4-address {
          preferred;
          primary;
        }
        vrrp-group group-number {
          (accept-data | no-accept-data);
          advertise-interval seconds;
          authentication-key key;
          authentication-type authentication;
          fast-interval milliseconds;
          (preempt | no-preempt) {
            hold-time seconds;
          }
        }
        priority number;
        track {
          interface interface-name {
            bandwidth-threshold bandwidth;
            priority-cost number;
          }
          priority-hold-time seconds;
          route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
        virtual-link-local-address address;
        vrrp-inherit-from {
          active-group group-number;
          active-interface interface-name;
        }
      }
    }
  }
  dhcp {
    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinite);
  }
}

```

```
    retransmission-attempt number;  
    retransmission-interval sections;  
    server-address ip-address;  
    update-server  
    vendor-id  
  }  
  filter {  
    input filter-name;  
    output filter-name;  
  }  
  no-neighbor-learn;  
  no-redirects;  
  primary;  
}  
family inet6 {  
  address address {  
    preferred;  
    primary;  
    vrrp-inet6-group group-id {  
      accept-data | no-accept-data;  
      authentication-key key;  
      authentication-type authentication;  
      fast-interval milliseconds;  
      inet6-advertise-interval milliseconds;  
      preempt | no-preempt {  
        hold-time seconds;  
      }  
    }  
    priority number;  
    track {  
      interface interface-name {  
        bandwidth-threshold bandwidth priority-cost number;  
        priority-cost number;  
      }  
      priority-hold-time seconds;  
      route ( address | routing-instance routing-instance-name );  
    }  
    virtual-inet6-address [addresses];  
    virtual-link-local-address ipv6-address;  
    vrrp-inherit-from {  
      active-group group-name;  
      active-interface interface-name;  
    }  
  }  
  (dad-disable | no-dad-disable);  
  filter {  
    group group-name;  
    input filter-name;  
    output filter-name;  
  }  
  no-neighbor-learn;  
  policer {  
    input policer-name;  
    output policer-name;  
  }  
}  
family iso {
```

```
        address interface-address;
    }
    family mpls;
    (traps | no-traps);
}
}
```

Unsupported Statements in the [edit interfaces lo] Hierarchy Level

All statements in the [edit interfaces lo] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 28: Unsupported [edit interfaces lo] Configuration Statements for EX Series Switches

Statement	Hierarchy
layer2-policer	[edit interfaces lo unit]
any	[edit interfaces lo unit family]
tcc	[edit interfaces lo unit family]
policer	[edit interfaces lo unit family inet]
unnumbered-address	[edit interfaces lo unit family inet]

- Related Documentation
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)
  - [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

[edit interfaces me] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit interfaces me] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces me\] Hierarchy Level on page 364](#)
- [Unsupported Statements in the \[edit interfaces me\] Hierarchy Level on page 366](#)

### Supported Statements in the [edit interfaces me] Hierarchy Level

---

The following hierarchy shows the [edit interfaces me] configuration statements supported on EX Series switches.

```
interfaces {
  me0 {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    no-gratuitous-arp-request;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family ethernet-switching {
        filter {
          input filter-name;
          output filter-name;
        }
        native-vlan-id vlan-id-number;
        port-mode (access | trunk);
        vlan {
          members [ members ];
        }
      }
    }
    family inet {
      accounting {
        destination-class-usage;
        source-class-usage {
          input;
          output;
        }
      }
      address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        master-only;
        preferred;
        primary;
      }
      dhcp {
        client-identifier (ascii client-id | hexadecimal client-id);
        lease-time (seconds | infinte);
        retransmission-attempt number;
        retransmission-interval sections;
        server-address ip-address;
      }
    }
  }
}
```

```

        update-server
        vendor-id
    }
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    no-neighbor-learn;
    primary;
    rpf-check;
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
}
(dad-disable | no-dad-disable);
filter {
    group group-name;
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
swap-by-poppush;
(traps | no-traps);
vlan-id vlan-id-number;
}
vlan-tagging;
}
}

```

### Unsupported Statements in the [edit interfaces me] Hierarchy Level

All statements in the [edit interfaces me] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 29: Unsupported [edit interfaces me] Configuration Statements for EX Series Switches**

Statement	Hierarchy
encapsulation	[edit interfaces me]
link-mode	[edit interfaces me]
encapsulation	[edit interfaces me unit]
layer2-policer	[edit interfaces me unit]
native-inner-vlan-id	[edit interfaces me unit]
vlan-id-list	[edit interfaces me unit]
vlan-id-range	[edit interfaces me unit]
ccc	[edit interfaces me unit family]
tcc	[edit interfaces me unit family]
vpls	[edit interfaces me unit family]
no-redirects	[edit interfaces me unit family inet]
policer	[edit interfaces me unit family inet]
sampling	[edit interfaces me unit family inet]
service	[edit interfaces me unit family inet]
unnumbered-address	[edit interfaces me unit family inet]
vrrp-group	[edit interfaces me unit family inet address]
service	[edit interfaces me unit family inet6]
vrrp-inet6-group	[edit interfaces me unit family inet6 address]

- Related Documentation**
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)
  - [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

## [edit interfaces vlan] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces vlan]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces vlan\] Hierarchy Level on page 367](#)
- [Unsupported Statements in the \[edit interfaces vlan\] Hierarchy Level on page 370](#)

### Supported Statements in the [edit interfaces vlan] Hierarchy Level

The following hierarchy shows the **[edit interfaces vlan]** configuration statements supported on EX Series switches.

```

interfaces {
  vlan {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family inet {
        accounting {
          destination-class-usage;
          source-class-usage {
            input;
            output;
          }
        }
      }
      address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        master-only;
      }
    }
  }
}

```

```
preferred;
primary;
vrrp-group group-number {
  (accept-data | no-accept-data);
  advertise-interval seconds;
  authentication-key key;
  authentication-type authentication;
  fast-interval milliseconds;
  (preempt | no-preempt) {
    hold-time seconds;
  }
  priority number;
  track {
    interface interface-name {
      bandwidth-threshold bandwidth;
      priority-cost number;
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
  }
  virtual-address [ addresses ];
  virtual-link-local-address address;
  vrrp-inherit-from {
    active-group group-number;
    active-interface interface-name;
  }
}
}
dhcp {
  client-identifier (ascii client-id | hexadecimal client-id);
  lease-time (seconds | infinte);
  retransmission-attempt number;
  retransmission-interval sections;
  server-address ip-address;
  update-server
  vendor-id
}
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
primary;
rpf-check;
}
family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
}
address address {
  eui-64;
```



```

ndp ip-address (mac | multicast-mac) mac-address <publish>;
preferred;
primary;
vrrp-inet6-group group-id {
    accept-data | no-accept-data;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    inet6-advertise-interval milliseconds;
    preempt | no-preempt {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bandwidth priority-cost number;
            priority-cost number;
        }
        priority-hold-time seconds;
        route ( address | routing-instance routing-instance-name );
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
    }
}
}
(dad-disable | no-dad-disable);
filter {
    group group-name;
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
}
proxy-arp (restricted | unrestricted);
(traps | no-traps);
}
}
}

```

### Unsupported Statements in the [edit interfaces vlan] Hierarchy Level

---

All statements in the [edit interfaces vlan] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)

### [edit interfaces vme] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit interfaces vme] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces vme\] Hierarchy Level on page 370](#)
- [Unsupported Statements in the \[edit interfaces vme\] Hierarchy Level on page 373](#)

### Supported Statements in the [edit interfaces vme] Hierarchy Level

---

The following hierarchy shows the [edit interfaces vme] configuration statements supported on EX Series switches.

```
interfaces {
  vme {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family inet {
        accounting {
```

```

destination-class-usage;
source-class-usage {
    input;
    output;
}
}
address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    master-only;
    preferred;
    primary;
    vrrp-group group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
        virtual-link-local-address address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
dhcp {
    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinte);
    retransmission-attempt number;
    retransmission-interval sections;
    server-address ip-address;
    update-server
    vendor-id
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
primary;
rpf-check;
}

```

```
family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
  address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
      accept-data | no-accept-data;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds;
      preempt | no-preempt {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          bandwidth-threshold bandwidth priority-cost number;
          priority-cost number;
        }
        priority-hold-time seconds;
        route ( address | routing-instance routing-instance-name );
      }
      virtual-inet6-address [addresses];
      virtual-link-local-address ipv6-address;
      vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
      }
    }
  }
  (dad-disable | no-dad-disable);
  filter {
    group group-name;
    input filter-name;
    output filter-name;
  }
  mtu bytes;
  no-neighbor-learn;
  policer {
    input policer-name;
    output policer-name;
  }
  rpf-check;
}
family iso {
  address interface-address;
  mtu bytes;
```

```

    }
    family mpls {
        mtu bytes;
    }
    (traps | no-traps);
    vlan-id vlan-id-number;
}
vlan-tagging;
}
}

```

### Unsupported Statements in the [edit interfaces vme] Hierarchy Level

All statements in the [edit interfaces vme] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

- Related Documentation**
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)
  - [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

## [edit interfaces xe] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit interfaces xe] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces xe\] Hierarchy Level on page 373](#)
- [Unsupported Statements in the \[edit interfaces xe\] Hierarchy Level on page 377](#)

### Supported Statements in the [edit interfaces xe] Hierarchy Level

The following hierarchy shows the [edit interfaces xe] configuration statements supported on EX Series switches.

```

interfaces {
  xe-fpc/pic/port {
    accounting-profile name;
    clocking (external | internal);
    description text;
    disable;
    ether-options {
      802.3ad {
        aex;
        (backup | primary);
      }
    }
  }
}

```

```
    lacp {
        force-up;
    }
}
(flow-control | no-flow-control);
(loopback | no-loopback);
}
framing (lan-phy | wan-phy);
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
optics-options {
    alarm alarm-type;
    warning alarm-type;
    wavelength nanometers;
}
traceoptions {
    flag flag;
}
(traps | no-traps);
unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    family ccc;
    family ethernet-switching {
        filter {
            input filter-name;
            output filter-name;
        }
        native-vlan-id vlan-id-number;
        port-mode (access | trunk);
        vlan {
            members [ members];
        }
    }
}
family inet {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    preferred;
    primary;
    vrrp-group group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-key key;
        authentication-type authentication;
```

```

    fast-interval milliseconds;
    (preempt | no-preempt) {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bandwidth;
            priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [ addresses ];
    virtual-link-local-address address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
}
dhcp {
    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinite);
    retransmission-attempt number;
    retransmission-interval sections;
    server-address ip-address;
    update-server
    vendor-id
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check;
targeted-broadcast;
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
    }
}

```

```

    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    inet6-advertise-interval milliseconds;
    preempt | no-preempt {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bandwidth priority-cost number;
            priority-cost number;
        }
        priority-hold-time seconds;
        route ( address | routing-instance routing-instance-name );
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
    }
    }
    }
    (dad-disable | no-dad-disable);
    filter {
        group group-name;
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    no-neighbor-learn;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check;
    }
    family iso {
        address interface-address;
        mtu bytes;
    }
    family mpls {
        mtu bytes;
    }
    proxy-arp (restricted | unrestricted);
    swap-by-poppush;
    (traps | no-traps);
    vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
    }
    vlan-tagging;
    }
    }

```



### Unsupported Statements in the `[edit interfaces xe]` Hierarchy Level

All statements in the `[edit interfaces xe]` hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

**Related Documentation**

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)

### `[edit poe]` Configuration Statement Hierarchy on EX Series Switches

This topic lists supported configuration statements in the `[edit poe]` hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the `\[edit poe\]` Hierarchy Level on page 377](#)

### Supported Statements in the `[edit poe]` Hierarchy Level

The following hierarchy shows the `[edit poe]` configuration statements supported on EX Series switches except for EX6200 and EX8200 switches:

```
poe {
  guard-band watts;
  interface (all | interface-name) {
    disable;
    maximum-power watts;
    priority (high | low);
    telemetries {
      disable;
      duration hours;
      interval minutes;
    }
  }
  lldp-priority;
  management (class | static);
  notification-control {
    fpc slot-number {
      disable;
    }
  }
}
```

The following hierarchy shows the `[edit poe]` configuration statements supported on EX Series switches for EX6200 and EX8200 switches:

```
poe {
  fpc (all | slot-number) {
```

```
guard-band watts;
lldp-priority;
management (class | static);
maximum-power watts;
}
interface (all | interface-name) {
  af-mode;
  disable;
  maximum-power watts;
  priority (high | low);
  telemetries {
    disable;
    duration hours;
    interval minutes;
  }
}
notification-control {
  fpc slot-number {
    disable;
  }
}
}
```

**Related Documentation**

- [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)
- [Example: Configuring PoE on an EX6200 or EX8200 Switch](#)
- [Configuring PoE on EX Series Switches \(CLI Procedure\) on page 4440](#)
- [Understanding PoE on EX Series Switches on page 4423](#)

## [\[edit policy-options\]](#) Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [\[edit policy-options\]](#) hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).
- [Supported Statements in the \[edit policy-options\] Hierarchy Level on page 378](#)
- [Unsupported Statements in the \[edit policy-options\] Hierarchy Level on page 389](#)

### [Supported Statements in the \[edit policy-options\] Hierarchy Level](#)

---

The following hierarchy shows the [\[edit policy-options\]](#) configuration statements supported on EX Series switches:

```
policy-options {
```

```

as-path name regular-expression {
    dynamic-db;
}
as-path-group group-name {
    as-path name regular-expression;
    dynamic-db;
}
community name {
    dynamic-db;
    invert-match;
    members [ community-ids ];
}
condition condition-name {
    dynamic-db;
    if-route-exists address table table-name;
}
damping name {
    disable;
    half-life minutes;
    max-suppress minutes;
    reuse number;
    suppress number;
}
policy-statement policy-name {
    dynamic-db;
    from {
        aggregate-contributor;
        area area-id;
        as-path [ regular-expression-names ];
        as-path-group [ as-path-group-names ];
        color preference;
        color2 preference;
        community [ community-names ];
        condition [ conditions ];
        external {
            type (1 | 2);
        }
        family family-name;
        instance instance-name;
        interface [ interface-names ];
        level isis-level;
        local-preference value;
        metric metric-value;
        metric2 metric-value;
        metric3 metric-value;
        metric4 metric-value;
        multicast-scope (scope-value | global | link-local | node-local | organization-local |
            site-local) (orhigher | orlower);
        neighbor [ ip-addresses ];
        next-hop [ ip-addresses ];
        origin (egp | igp | incomplete);
        policy [ policy-names ];
        preference preference;
        preference2 preference;
        prefix-list prefix-list-name;
        prefix-list-filter prefix-list-name (exact | longer | orlonger) {

```

```

(accept | reject);
as-path-expand (as-number | last-as) <count number>;
as-path-prepend as-number;
class class-name;
color (preference | add number | subtract number);
color2 (preference | add number | subtract number);
community (add | delete | set | + | - | =) community-name;
cos-next-hop-map map-name;
damping list-name;
default-action (accept | reject);
destination-class class-name;
dynamic-db;
external {
    type (1 | 2);
}
forwarding-class class-name;
install-nexthop <strict> (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
    static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])
    <except (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
        static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ]) >;
label-allocation (per-nexthop | per-table);
load-balance per-packet;
local-preference (preference | add number | subtract number);
map-to-interface (interface-name | self);
metric (metric-value | add number | igp <metric-offset> |
    minimum-igp <metric-offset> | subtract number | ... the following complex
    expression ...);
expression {
    metric (multiplier number | offset number | multiplier number offset number);
    metric2 (multiplier number | offset number | multiplier number offset number);
}
metric2 (metric-value | add number | subtract number);
metric3 (metric-value | add number | subtract number);
metric4 (metric-value | add number | subtract number);
next (policy | term);
next-hop (ip-address | discard | next-table routing-table-name | peer-address |
    reject | self);
origin (egp | igp | incomplete);
preference (preference | add number | subtract number);
preference2 (preference | add number | subtract number);
priority (high | low | medium);
source-class class-name;
ssm-source source;
tag (tag-number | add number | subtract number);
tag2 (tag-number | add number | subtract number);
trace;
}
protocol [ protocol-names ];
rib routing-table-name;
route-filter ip-prefix</prefix-length> (exact | longer | orlonger |
    through ip-prefix</prefix-length> | upto /prefix-length) {
    (accept | reject);
    as-path-expand (as-number | last-as) <count number>;
    as-path-prepend as-number;
    class class-name;
    color (preference | add number | subtract number);

```

```

color2 (preference | add number | subtract number);
community (add | delete | set | + | - | =) community-name;
cos-next-hop-map map-name;
damping list-name;
default-action (accept | reject);
destination-class class-name;
dynamic-db;
external {
    type (1 | 2);
}
forwarding-class class-name;
install-nexthop <strict> (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
    static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])
    <except (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
        static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])>;
label-allocation (per-nexthop | per-table);
load-balance per-packet;
local-preference (preference | add number | subtract number);
map-to-interface (interface-name | self);
metric (metric-value | add number | igp <metric-offset> |
    minimum-igp <metric-offset> | subtract number | ... the following complex
    expression ...);
expression {
    metric (multiplier number | offset number | multiplier number offset number);
    metric2 (multiplier number | offset number | multiplier number offset number);
}
metric2 (metric-value | add number | subtract number);
metric3 (metric-value | add number | subtract number);
metric4 (metric-value | add number | subtract number);
next (policy | term);
next-hop (ip-address | discard | next-table routing-table-name | peer-address |
    reject | self);
origin (egp | igp | incomplete);
preference (preference | add number | subtract number);
preference2 (preference | add number | subtract number);
priority (high | low | medium);
source-class class-name;
ssm-source source;
tag (tag-number | add number | subtract number);
tag2 (tag-number | add number | subtract number);
trace;
}
route-type (external | internal);
source-address-filter ip-prefix</prefix-length> (exact | longer | orlonger |
    through ip-prefix</prefix-length> | upto /prefix-length) {
    route-filter ip-prefix</prefix-length> (exact | longer | orlonger |
        through ip-prefix</prefix-length> | upto /prefix-length) {
        (accept | reject);
    }
}
as-path-expand (as-number | last-as) <count number>;
as-path-prepend as-number;
class class-name;
color (preference | add number | subtract number);
color2 (preference | add number | subtract number);
community (add | delete | set | + | - | =) community-name;
cos-next-hop-map map-name;
damping list-name;

```

```
default-action (accept | reject);
destination-class class-name;
dynamic-db;
external {
    type (1 | 2);
}
forwarding-class class-name;
install-nexthop <strict> (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
    static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])
    <except (lsp [ lsp-names ] | lsp-regex [ regular-expressions |
    static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])]>;
label-allocation (per-nexthop | per-table);
load-balance per-packet;
local-preference (preference | add number | subtract number);
map-to-interface (interface-name | self);
metric (metric-value | add number | igp <metric-offset> |
    minimum-igp <metric-offset> | subtract number | ... the following complex
    expression ...);
expression {
    metric (multiplier number | offset number | multiplier number offset number);
    metric2 (multiplier number | offset number | multiplier number offset number);
}
metric2 (metric-value | add number | subtract number);
metric3 (metric-value | add number | subtract number);
metric4 (metric-value | add number | subtract number);
next (policy | term);
next-hop (ip-address | discard | next-table routing-table-name | peer-address |
    reject | self);
origin (egp | igp | incomplete);
preference (preference | add number | subtract number);
preference2 (preference | add number | subtract number);
priority (high | low | medium);
source-class class-name;
ssm-source source;
tag (tag-number | add number | subtract number);
tag2 (tag-number | add number | subtract number);
trace;
}
tag [ tag-numbers ];
tag2 tag-number;
}
term term-name {
    from {
        aggregate-contributor;
        area area-id;
        as-path [ regular-expression-names ];
        as-path-group [ as-path-group-names ];
        color preference;
        color2 preference;
        community [ community-names ];
        community-count number;
        condition [ conditions ];
        external {
            type (1 | 2);
        }
        family family-name;
```

```

instance instance-name;
interface [ interface-names ];
level isis-level;
local-preference value;
metric metric-value;
metric2 metric-value;
metric3 metric-value;
metric4 metric-value;
multicast-scope (scope-value | global | link-local | node-local | organization-local
| site-local) (orhigher | orlower);
neighbor [ ip-addresses ];
next-hop [ ip-addresses ];
next-hop-type type;
origin (egp | igp | incomplete);
policy [ policy-names ];
preference preference;
preference2 preference;
prefix-list prefix-list-name;
prefix-list-filter prefix-list-name (exact | longer | orlonger) {
    (accept | reject);
    as-path-expand (as-number | last-as) <count number>;
    as-path-prepend as-number;
    class class-name;
    color (preference | add number | subtract number);
    color2 (preference | add number | subtract number);
    community (add | delete | set | + | - | =) community-name;
    cos-next-hop-map map-name;
    damping list-name;
    default-action (accept | reject);
    destination-class class-name;
    dynamic-db;
    external {
        type (1 | 2);
    }
    forwarding-class class-name;
    install-nexthop <strict> (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
        static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])
        <except (lsp [ lsp-names ] | lsp-regex [ regular-expressions |
            static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])]>;
    label-allocation (per-nexthop | per-table);
    load-balance per-packet;
    local-preference (preference | add number | subtract number);
    map-to-interface (interface-name | self);
    metric (metric-value | add number | igp <metric-offset> |
        minimum-igp <metric-offset> | subtract number | ... the following complex
        expression ...);
    expression {
        metric (multiplier number | offset number | multiplier number offset number);
        metric2 (multiplier number | offset number | multiplier number offset number);
    }
    metric2 (metric-value | add number | subtract number);
    metric3 (metric-value | add number | subtract number);
    metric4 (metric-value | add number | subtract number);
    next (policy | term);
    next-hop (ip-address | discard | next-table routing-table-name | peer-address |
        reject | self);

```

```

origin (egp | igp | incomplete);
preference (preference | add number | subtract number);
preference2 (preference | add number | subtract number);
priority (high | low | medium);
source-class class-name;
ssm-source source;
tag (tag-number | add number | subtract number);
tag2 (tag-number | add number | subtract number);
trace;
}
protocol [ protocol-names ];
rib routing-table-name;
source-address-filter ip-prefix</prefix-length> (exact | longer | orlonger |
  through ip-prefix</prefix-length> | upto /prefix-length) {
  route-filter ip-prefix</prefix-length> (exact | longer | orlonger |
    through ip-prefix</prefix-length> | upto /prefix-length) {
    (accept | reject);
  }
}
as-path-expand (as-number | last-as) <count number>;
as-path-prepend as-number;
class class-name;
color (preference | add number | subtract number);
color2 (preference | add number | subtract number);
community (add | delete | set | + | - | =) community-name;
cos-next-hop-map map-name;
damping list-name;
default-action (accept | reject);
destination-class class-name;
dynamic-db;
external {
  type (1 | 2);
}
forwarding-class class-name;
install-nexthop <strict> (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
  static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])
  <except (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
    static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])>;
label-allocation (per-nexthop | per-table);
load-balance per-packet;
local-preference (preference | add number | subtract number);
map-to-interface (interface-name | self);
metric (metric-value | add number | igp <metric-offset> |
  minimum-igp <metric-offset> | subtract number | ... the following complex
  expression ...);
expression {
  metric (multiplier number | offset number | multiplier number offset number);
  metric2 (multiplier number | offset number | multiplier number offset number);
}
metric2 (metric-value | add number | subtract number);
metric3 (metric-value | add number | subtract number);
metric4 (metric-value | add number | subtract number);
next (policy | term);
next-hop (ip-address | discard | next-table routing-table-name | peer-address |
  reject | self);
origin (egp | igp | incomplete);
preference (preference | add number | subtract number);
preference2 (preference | add number | subtract number);

```



```

priority (high | low | medium);
source-class class-name;
ssm-source source;
tag (tag-number | add number | subtract number);
tag2 (tag-number | add number | subtract number);
trace;
}
route-type (external | internal);
source-address-filter ip-prefix</prefix-length> (exact | longer | orlonger |
through ip-prefix</prefix-length> | upto /prefix-length) {
route-filter ip-prefix</prefix-length> (exact | longer | orlonger |
through ip-prefix</prefix-length> | upto /prefix-length) {
(accept | reject);
as-path-expand (as-number | last-as) <count number>;
as-path-prepend as-number;
class class-name;
color (preference | add number | subtract number);
color2 (preference | add number | subtract number);
community (add | delete | set | + | - | =) community-name;
cos-next-hop-map map-name;
damping list-name;
default-action (accept | reject);
destination-class class-name;
dynamic-db;
external {
type (1 | 2);
}
forwarding-class class-name;
install-nexthop <strict> (lsp [lsp-names] | lsp-regex [regular-expressions] |
static-lsp [lsp-names] | static-lsp-regex [regular-expressions])
<except (lsp [lsp-names] | lsp-regex [regular-expressions] |
static-lsp [lsp-names] | static-lsp-regex [regular-expressions])>;
label-allocation (per-nexthop | per-table);
load-balance per-packet;
local-preference (preference | add number | subtract number);
map-to-interface (interface-name | self);
metric (metric-value | add number | igp <metric-offset> |
minimum-igp <metric-offset> | subtract number | ... the following complex
expression ...);
expression {
metric (multiplier number | offset number | multiplier number offset number);
metric2 (multiplier number | offset number | multiplier number offset number);
}
metric2 (metric-value | add number | subtract number);
metric3 (metric-value | add number | subtract number);
metric4 (metric-value | add number | subtract number);
next (policy | term);
next-hop (ip-address | discard | next-table routing-table-name | peer-address |
reject | self);
origin (egp | igp | incomplete);
preference (preference | add number | subtract number);
preference2 (preference | add number | subtract number);
priority (high | low | medium);
source-class class-name;
ssm-source source;
tag (tag-number | add number | subtract number);

```

```

    tag2 (tag-number | add number | subtract number);
    trace;
}
state;
tag [ tag-numbers ];
tag2 tag-number;
}
then {
    (accept | reject);
    as-path-expand (as-number | last-as) <count number>;
    as-path-prepend as-number;
    class class-name;
    color (preference | add number | subtract number);
    color2 (preference | add number | subtract number);
    community (add | delete | set | + | - | =) community-name;
    cos-next-hop-map map-name;
    damping list-name;
    default-action (accept | reject);
    destination-class class-name;
    external {
        type (1 | 2);
    }
    forwarding-class class-name;
    install-nexthop <strict> (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
        static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])
        <except (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
            static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])>;
    label-allocation (per-nexthop | per-table);
    load-balance per-packet;
    local-preference (preference | add number | subtract number);
    map-to-interface (interface-name | self);
    metric (metric-value | add number | igp <metric-offset> |
        minimum-igp <metric-offset> | subtract number | ... the following complex
        expression ...);
    expression {
        metric (multiplier number | offset number | multiplier number offset number);
        metric2 (multiplier number | offset number | multiplier number offset number);
    }
    metric2 (metric-value | add number | subtract number);
    metric3 (metric-value | add number | subtract number);
    metric4 (metric-value | add number | subtract number);
    next (policy | term);
    next-hop (ip-address | discard | next-table routing-table-name | peer-address |
        reject | self);
    origin (egp | igp | incomplete);
    preference (preference | add number | subtract number);
    preference2 (preference | add number | subtract number);
    priority (high | low | medium);
    source-class class-name;
    ssm-source source;
    tag (tag-number | add number | subtract number);
    tag2 (tag-number | add number | subtract number);
    trace;
}
to {
    area area-id;

```

```

as-path [ regular-expression-names ];
as-path-group [ as-path-group-names ];
color preference;
color2 preference;
community [ community-names ];
external {
    type (1 | 2);
}
family family-name;
instance instance-name;
interface [ interface-names ];
level isis-level;
local-preference value;
metric metric-value;
metric2 metric-value;
metric3 metric-value;
metric4 metric-value;
neighbor [ ip-addresses ];
next-hop [ ip-addresses ];
origin (egp | igp | incomplete);
policy [ policy-names ];
preference preference;
preference2 preference;
protocol [ protocol-names ];
rib routing-table-name;
tag [ tag-numbers ];
tag2 tag-number;
}
}
then {
    (accept | reject);
    as-path-expand (as-number | last-as) <count number>;
    as-path-prepend as-number;
    class class-name;
    color (preference | add number | subtract number);
    color2 (preference | add number | subtract number);
    community (add | delete | set | + | - | =) community-name;
    cos-next-hop-map map-name;
    damping list-name;
    default-action (accept | reject);
    destination-class class-name;
    external {
        type (1 | 2);
    }
    forwarding-class class-name;
    install-nexthop <strict> (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
        static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])
        <except (lsp [ lsp-names ] | lsp-regex [ regular-expressions ]
        static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])>;
    label-allocation (per-nexthop | per-table);
    load-balance per-packet;
    local-preference (preference | add number | subtract number);
    map-to-interface (interface-name | self);
    metric (metric-value | add number | igp <metric-offset> |
        minimum-igp <metric-offset> | subtract number | ... the following complex expression
        ...);
}

```

```

    expression {
        metric (multiplier number | offset number | multiplier number offset number);
        metric2 (multiplier number | offset number | multiplier number offset number);
    }
    metric2 (metric-value | add number | subtract number);
    metric3 (metric-value | add number | subtract number);
    metric4 (metric-value | add number | subtract number);
    next (policy | term);
    next-hop (ip-address | discard | next-table routing-table-name | peer-address | reject |
        self);
    origin (egp | igp | incomplete);
    preference (preference | add number | subtract number);
    preference2 (preference | add number | subtract number);
    priority (high | low | medium);
    source-class class-name;
    ssm-source source;
    tag (tag-number | add number | subtract number);
    tag2 (tag-number | add number | subtract number);
    trace;
}
to {
    area area-id;
    as-path [ regular-expression-names ];
    as-path-group [ as-path-group-names ];
    color preference;
    color2 preference;
    community [ community-names ];
    external {
        type (1 | 2);
    }
    family family-name;
    instance instance-name;
    interface [ interface-names ];
    level isis-level;
    local-preference value;
    metric metric-value;
    metric2 metric-value;
    metric3 metric-value;
    metric4 metric-value;
    neighbor [ ip-addresses ];
    next-hop [ ip-addresses ];
    origin (egp | igp | incomplete);
    policy [ policy-names ];
    preference preference;
    preference2 preference;
    protocol [ protocol-names ];
    rib routing-table-name;
    tag [ tag-numbers ];
    tag2 tag-number;
}
}
prefix-list list-name {
    ip-prefix</prefix-length>;
    apply-path path;
    dynamic-db;
}

```

```

vsi-policy policy-name {
  from {
    vsi-manager identifier vsi-type identifier vsi-version version-number vsi-instance
      instance-name;
  }
  then {
    filter filter-name;
  }
}

```

### Unsupported Statements in the [edit policy-options] Hierarchy Level

All statements in the **[edit policy-options]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*

## [edit protocols] Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the **[edit protocols]** hierarchy:

- [\[edit protocols bfd\] Configuration Statement Hierarchy on EX Series Switches on page 391](#)
- [\[edit protocols bgp\] Configuration Statement Hierarchy on EX Series Switches on page 392](#)
- [\[edit protocols connections\] Configuration Statement Hierarchy on EX Series Switches on page 401](#)
- [\[edit protocols dcbx\] Configuration Statement Hierarchy on EX Series Switches on page 403](#)
- [\[edit protocols dot1x\] Configuration Statement Hierarchy on EX Series Switches on page 404](#)
- [\[edit protocols igmp\] Configuration Statement Hierarchy on EX Series Switches on page 406](#)
- [\[edit protocols igmp-snooping\] Configuration Statement Hierarchy on EX Series Switches on page 407](#)
- [\[edit protocols isis\] Configuration Statement Hierarchy on EX Series Switches on page 408](#)
- [\[edit protocols lacp\] Configuration Statement Hierarchy on EX Series Switches on page 411](#)
- [\[edit protocols link-management\] Configuration Statement Hierarchy on EX Series Switches on page 412](#)
- [\[edit protocols lldp\] Configuration Statement Hierarchy on EX Series Switches on page 413](#)

- [\[edit protocols lldp-med\] Configuration Statement Hierarchy on EX Series Switches on page 415](#)
- [\[edit protocols mld\] Configuration Statement Hierarchy on EX Series Switches on page 416](#)
- [\[edit protocols mld-snooping\] Configuration Statement Hierarchy on EX Series Switches on page 417](#)
- [\[edit protocols mpls\] Configuration Statement Hierarchy on EX Series Switches on page 418](#)
- [\[edit protocols msdp\] Configuration Statement Hierarchy on EX Series Switches on page 429](#)
- [\[edit protocols mstp\] Configuration Statement Hierarchy on EX Series Switches on page 431](#)
- [\[edit protocols mvrp\] Configuration Statement Hierarchy on EX Series Switches on page 433](#)
- [\[edit protocols neighbor-discovery\] Configuration Statement Hierarchy on EX Series Switches on page 434](#)
- [\[edit protocols oam\] Configuration Statement Hierarchy on EX Series Switches on page 435](#)
- [\[edit protocols ospf\] Configuration Statement Hierarchy on EX Series Switches on page 438](#)
- [\[edit protocols ospf3\] Configuration Statement Hierarchy on EX Series Switches on page 441](#)
- [\[edit protocols pim\] Configuration Statement Hierarchy on EX Series Switches on page 444](#)
- [\[edit protocols rip\] Configuration Statement Hierarchy on EX Series Switches on page 447](#)
- [\[edit protocols ripng\] Configuration Statement Hierarchy on EX Series Switches on page 450](#)
- [\[edit protocols router-advertisement\] Configuration Statement Hierarchy on EX Series Switches on page 451](#)
- [\[edit protocols router-discovery\] Configuration Statement Hierarchy on EX Series Switches on page 452](#)
- [\[edit protocols rstp\] Configuration Statement Hierarchy on EX Series Switches on page 453](#)
- [\[edit protocols rsvp\] Configuration Statement Hierarchy on EX Series Switches on page 455](#)
- [\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches on page 459](#)
- [\[edit protocols stp\] Configuration Statement Hierarchy on EX Series Switches on page 460](#)

- [\[edit protocols uplink-failure-detection\]](#) Configuration Statement Hierarchy on EX Series Switches on page 461
- [\[edit protocols vrrp\]](#) Configuration Statement Hierarchy on EX Series Switches on page 462
- [\[edit protocols vstp\]](#) Configuration Statement Hierarchy on EX Series Switches on page 463

**Related  
Documentation**

- [EX Series Switch Software Features Overview](#)
- [EX Series Virtual Chassis Software Features Overview](#)

## [\[edit protocols bfd\]](#) Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [\[edit protocols bfd\]](#) hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit protocols bfd\] Hierarchy Level on page 391](#)
- [Unsupported Statements in the \[edit protocols bfd\] Hierarchy Level on page 391](#)

### [Supported Statements in the \[edit protocols bfd\] Hierarchy Level](#)

The following hierarchy shows the [\[edit protocols bfd\]](#) configuration statements supported on EX Series switches:

```
protocols {
  bfd {
    no-issu-timer-negotiation;
    traceoptions {
      file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
      flag <flag>;
      no-remote-trace;
    }
  }
}
```

### [Unsupported Statements in the \[edit protocols bfd\] Hierarchy Level](#)

All statements in the [\[edit protocols bfd\]](#) hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

- Related Documentation**
- *Protocol-Independent Routing Properties Feature Guide for Routing Devices*
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [\[edit protocols bgp\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit protocols bgp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit protocols bgp\] Hierarchy Level on page 392](#)
- [Unsupported Statements in the \[edit protocols bgp\] Hierarchy Level on page 400](#)

### [Supported Statements in the \[edit protocols bgp\] Hierarchy Level](#)

---

The following hierarchy shows the **[edit protocols bgp]** configuration statements supported on EX Series switches:

```
protocols {
  bgp {
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    authentication-algorithm (hmac-sha-1-96 | md5);
    authentication-key key;
    bfd-liveness-detection {
      authentication {
        algorithm algorithm-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      hold-down-interval milliseconds;
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      session-mode (automatic | multihop | single-hop);
      transmit-interval {
        threshold milliseconds;
        minimum-interval milliseconds;
      }
    }
    version (1 | automatic);
```



```

}
cluster cluster-identifier;
damping;
description text-description;
disable;
export [ policy-names ];
family inet {
  any {
    loops number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
  }
  flow {
    loops number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
  }
  multicast {
    loops number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
  }
  unicast {
    loops number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
    topology name {
      community target identifier;
    }
  }
}
family inet6 {
  any {
    loops number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
  }
  labeled-unicast {
    aggregate-label {
      community community-name;
    }
  }
}

```

```
explicit-null connected-only;
loops number;
per-group-label;
prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
}
}
rib-group group-name;
traffic-statistics {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    interval seconds;
}
}
multicast {
    loops number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
}
unicast {
    loops number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
    topology name {
        community target identifier;
    }
}
} # end of [edit protocols bgp family]
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
group group-name {
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    allow [ all ip-prefix</prefix-length> ];
    as-override;
    authentication-algorithm (hmac-sha-1-96 | md5);
    authentication-key key;
    bfd-liveness-detection {
        authentication {
            algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                meticulous-keyed-sha-1 | simple-password);
            loose-check;
        }
        holddown-interval milliseconds;
        minimum-interval milliseconds;
    }
}
```

```

minimum-receive-interval milliseconds;
multiplier number;
no-adaptation;
session-mode (automatic | multihop | single-hop);
transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
}
version (1 | automatic);
}
cluster cluster-identifier;
damping;
description text-description;
export [ policy-names ];
family inet {
    any {
        loops number;
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        rib-group group-name;
    }
    flow {
        loops number;
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        rib-group group-name;
    }
    multicast {
        loops number;
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        rib-group group-name;
    }
    unicast {
        loops number;
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        rib-group group-name;
        topology name {
            community target identifier;
        }
    }
}
family inet6 {
    any {
        loops number;
        prefix-limit {
            maximum number;

```

```

        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
}
labeled-unicast {
    aggregate-label {
        community community-name;
    }
    explicit-null connected-only;
    loops number;
    per-group-label;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
    traffic-statistics {
        file filename <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
        interval seconds;
    }
}
multicast {
    loops number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
}
unicast {
    loops number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    rib-group group-name;
    topology name {
        community target identifier;
    }
}
} # end of [edit protocols bgp group family]
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
hold-time seconds;
idle-after-switch-over (seconds | forever);
import [ policy-names ];
include-mp-next-hop;
keep (all | none);
local-interface interface-name;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);

```

```

mtu-discovery;
multihop {
  no-nexthop-change;
  ttl tvl-value;
}
neighboraddress {
  advertise-external <conditional>;
  advertise-inactive;
  (advertise-peer-as | no-advertise-peer-as);
  as-override;
  authentication-algorithm algorithm;
  authentication-key key;
  bfd-liveness-detection {
    authentication {
      algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
        meticulous-keyed-sha-1 | simple-password);
      loose-check;
    }
    holddown-interval milliseconds;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    session-mode (automatic | multihop | single-hop);
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (1 | automatic);
  }
}
cluster cluster-identifier;
damping;
description text-description;
export [ policy-names ];
family {
  (inet | inet6 | inet-mvpn | inet6-mpvn | inet-vpn | inet6-vpn | iso-vpn | l2-vpn)
  {
    (any | flow | multicast | unicast | signaling) {
      accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      damping;
      prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      rib-group group-name;
      topology name {
        community {
          target identifier;
        }
      }
    }
  }
  flow {
    no-validate policy-name;
  }
}

```

```
}
labeled-unicast {
  accepted-prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  aggregate-label {
    community community-name;
  }
  explicit-null {
    connected-only;
  }
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  resolve-vpn;
  rib inet.3;
  rib-group group-name;
  topology name {
    community {
      target identifier;
    }
  }
}
}
route-target {
  advertise-default;
  external-paths number;
  accepted-prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
}
signaling {
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
}
}
graceful-restart {
  disable;
  restart-time seconds;
  stale-routes-time seconds;
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
```

```

local-interface interface-name;
local-preference preference;
log-updown;
metric-out (metric | minimum-igp <offset> | igp <offset>);
mtu-discovery;
multihop <ttl-value>;
multipath {
    multiple-as;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
type (external | internal);
}
hold-time seconds;
idle-after-switch-over (seconds | forever);
import [ policy-names ];
include-mp-next-hop;
keep (all | none);
local-address address;
local-as autonomous-system <loops number> < alias> <no-prepend-global-as>
    <private>;
local-interface interface-name;
local-preference local-preference;

```

```

log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl ttl-value;
}
multipath;
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
}
peer-as autonomous-system;
precision-timers | no-precision-timers;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
}

```

### Unsupported Statements in the [edit protocols bgp] Hierarchy Level

All statements in the [edit protocols bgp] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 30: Unsupported [edit protocols bgp] Configuration Statements on EX Series Switches**

Statement	Hierarchy Level
NOTE: Variables, such as <i>family-name</i> , are not shown in the statements or hierarchies.	
accepted-prefix-limit	[edit protocols bgp family]
add-path	[edit protocols bgp family]



Table 30: Unsupported [edit protocols bgp] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy Level
authentication-key-chain	[edit protocols bgp] [edit protocols bgp group] [edit protocols bgp group neighbor]
cisco-non-deterministic	[edit protocols bgp path-selection]
eternal-router-id	[edit protocols bgp path-selection]
igp-multiplier	[edit protocols bgp path-selection]
ipsec-sa	[edit protocols bgp] [edit protocols bgp group]
key-chain	[edit protocols bgp bfd-liveness-detection authentication] [edit protocols bgp group bfd-liveness-detection authentication] [edit protocols bgp group neighbor bfd-liveness-detection authentication]
maximum	[edit protocols bgp family accepted-prefix-limit]
med-multiplier	[edit protocols bgp path-selection]
med-plus-igp	[edit protocols bgp path-selection]
no-validate	[edit protocols bgp family inet flow]
path-count	[edit protocols bgp family add-path send]
prefix-policy	[edit protocols bgp family add-path send]
receive	[edit protocols bgp family add-path]
send	[edit protocols bgp family add-path]
teardown	[edit protocols bgp family accepted-prefix-limit]
vpn-apply-export	[edit protocols bgp] [edit protocols bgp group neighbor]

- Related Documentation**
- *BGP Feature Guide for Routing Devices*
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols connections] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols connections]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols connections\] Hierarchy Level on page 402](#)
- [Unsupported Statements in the \[edit protocols connections\] Hierarchy Level on page 402](#)

### Supported Statements in the [edit protocols connections] Hierarchy Level

The following hierarchy shows the **[edit protocols connections]** configuration statements supported on EX Series switches:

```
protocols {
  connections {
    interface-switch connection-name {
      interface interface-name.unit-number;
    }
    lsp-switch connection-name {
      receive-lsp label-switched-path;
      transmit-lsp label-switched-path;
    }
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      receive-lsp label-switched-path;
      transmit-lsp label-switched-path;
    }
  }
}
```

### Unsupported Statements in the [edit protocols connections] Hierarchy Level

All statements in the **[edit protocols connections]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 31: Unsupported [edit protocols connections] Configuration Statements on EX Series Switches**

Statement	Hierarchy
<i>NOTE:</i> Variables, such as <i>p2mp-receive-switch</i> , are not shown in the statements or hierarchies.	
p2mp-receive-switch	[edit protocols connections]
p2mp-transmit-switch	[edit protocols connections]

## Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [\[edit protocols dcbx\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit protocols dcbx]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols dcbx\] Hierarchy Level on page 403](#)
- [Unsupported Statements in the \[edit protocols dcbx\] Hierarchy Level on page 403](#)

### [Supported Statements in the \[edit protocols dcbx\] Hierarchy Level](#)

The following hierarchy shows the **[edit protocols dcbx]** configuration statements supported on EX Series switches:

```
protocols {
  dcbx {
    disable;
    interface (all | interface-name) {
      application-map application-map-name;
      applications {
        fcoe {
          no-auto-negotiation;
        }
      }
      disable ;
      priority-flow-control {
        no-auto-negotiation;
      }
    }
  }
}
```

### [Unsupported Statements in the \[edit protocols dcbx\] Hierarchy Level](#)

All statements in the **[edit protocols dcbx]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 32: Unsupported [edit protocols dcbx] Configuration Statements on EX Series Switches

Statement	Hierarchy
-----------	-----------

NOTE: Variables, such as *interface-name*, are not shown in the statements or hierarchies.

enhanced-transmission-selection	[edit protocols dcbx interface]
---------------------------------	---------------------------------

#### Related Documentation

- *Example: Configuring an FCoE Transit Switch*
- *Example: Configuring DCBX to Support an iSCSI Application*
- *Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)*
- *Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)*
- *Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches*
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols dot1x] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols dot1x] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols dot1x\] Hierarchy Level on page 404](#)
- [Unsupported Statements in the \[edit protocols dot1x\] Hierarchy Level on page 405](#)

### Supported Statements in the [edit protocols dot1x] Hierarchy Level

The following hierarchy shows the [edit protocols dot1x] configuration statements supported on EX Series switches:

```
protocols {
  dot1x {
    authenticator {
      authentication-profile-name access-profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan (vlan-id | vlan-name);
        mac-radius {
          flap-on-disconnect;
          restrict;
        }
      }
    }
  }
}
```

## Unsupported Statements in the [edit protocols dot1x] Hierarchy Level

- [Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an EX Series Switch on page 1847](#)
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to an EX Series Switch on page 1875](#)
- [Example: Configuring Fallback Options on EX Series Switches for EAP-TTLS Authentication and Odyssey Access Clients on page 1903](#)
- [Example: Configuring Static MAC Bypass of Authentication on an EX Series Switch on page 1872](#)

- [802.1X for EX Series Switches Overview on page 1821](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## **[edit protocols igmp] Configuration Statement Hierarchy on EX Series Switches**

This topic lists supported and unsupported configuration statements in the **[edit protocols igmp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols igmp\] Hierarchy Level on page 406](#)
- [Unsupported Statements in the \[edit protocols igmp\] Hierarchy Level on page 407](#)

### **Supported Statements in the [edit protocols igmp] Hierarchy Level**

---

The following hierarchy shows the **[edit protocols igmp]** configuration statements supported on EX Series switches:

```
protocols {
  igmp {
    accounting;
    interface interface-name {
      (accounting | no-accounting);
      disable;
      group-policy [ policy-names ];
      group-policy policy-name;
      immediate-leave;
      oif-map [ map-names ];
      passive <allow-receive> <send-general-query> <send-group-query>;
      promiscuous-mode;
      ssm-map ssm-map-name;
      static {
        group multicast-group-address {
          exclude;
          group-count number;
          group-increment increment;
          source ip-address {
            source-count number;
            source-increment increment;
          }
        }
      }
    }
  }
  version version;
```

```

    }
    maximum-transmit-rate packets-per-second;
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}

```

### Unsupported Statements in the [edit protocols igmp] Hierarchy Level

All statements in the [edit protocols igmp] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols igmp-snooping] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols igmp-snooping] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols igmp-snooping\] Hierarchy Level on page 407](#)
- [Unsupported Statements in the \[edit protocols igmp-snooping\] Hierarchy Level on page 408](#)

### Supported Statements in the [edit protocols igmp-snooping] Hierarchy Level

The following hierarchy shows the [edit protocols igmp-snooping] configuration statements supported on EX Series switches:

```

protocols {
  igmp-snooping {
    vlan vlan-identifier{
      immediate-leave;
      interface (all | interface-name) {
        group-limit <1..65535>
      }
    }
  }
}

```

```
host-only-interface
multicast-router-interface;
immediate-leave;
static {
    group multicast-ip-address; {
        source <>
    }
}
}
}
proxy {
    source-address ip-address;
}
query-interval number;
query-last-member-interval number;
query-response-interval number;
robust-count number;
traceoptions {
    file filename <files number> <no-stamp> <replace> <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
}
```

#### Unsupported Statements in the [edit protocols igmp-snooping] Hierarchy Level

All statements in the [edit protocols igmp-snooping] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

### [edit protocols isis] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols isis] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols isis\] Hierarchy Level on page 409](#)
- [Unsupported Statements in the \[edit protocols isis\] Hierarchy Level on page 411](#)



### Supported Statements in the [edit protocols isis] Hierarchy Level

The following hierarchy shows the [edit protocols isis] configuration statements supported on EX Series switches.

```

protocols {
  isis {
    disable;
    export [ policy-names ];
    graceful-restart {
      disable;
      helper-disable;
      restart-duration seconds;
    }
    ignore-attached-bit;
    interface interface-name {
      bfd-liveness-detection {
        authentication {
          algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
            meticulous-keyed-sha-1 | simple-password);
          loose-check;
        }
        detection-time {
          threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        version (0 | 1 | automatic);
      }
      checksum;
      csnp-interval (seconds | disable);
      disable;
      hello-padding (adaptive | loose | strict);
      disable;
    }
    level (1 | 2) {
      disable;
      hello-authentication-key key;
      hello-authentication-key-chain;
      hello-authentication-type authentication;
      hello-interval seconds;
      hold-time seconds;
      ipv4-multicast-metric number;
      ipv6-multicast-metric number;
      ipv6-unicast-metric number;
      metric metric;
      passive;
      priority number;
    }
  }
}

```

```
link-protection;
mesh-group (value | blocked);
no-adjacency-down-notification;
no-eligible-backup;
no-ipv4-multicast;
no-ipv6-multicast;
no-ipv6-unicast;
no-unicast-topology;
node-link-protection;
passive;
point-to-point;
}
level (1 | 2) {
  authentication-key key;
  authentication-type authentication;
  disable;
  external-preference preference;
  no-csnp-authentication;
  no-hello-authentication;
  no-psnp-authentication;
  preference preference;
  prefix-export-limit number;
  wide-metrics-only;
}
loose-authentication-check;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
no-ipv6-routing;
overload {
  advertise-high-metrics;
  allow-route-leaking;
  timeout seconds;
}
reference-bandwidth reference-bandwidth;
rib-group {
  inet group-name;
  inet6 group-name;
}
spf-options {
  delay milliseconds;
  holddown milliseconds;
  rapid-runs number;
}
topologies {
  ipv4-multicast;
  ipv6-multicast;
  ipv6-unicast;
}
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
}
```

```
}

```

Unsupported Statements in the [edit protocols isis] Hierarchy Level

All statements in the [edit protocols isis] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 33: Unsupported [edit protocols isis] Configuration Statements on EX Series Switches

Statement	Hierarchy
NOTE: Variables, such as <i>filename</i> , are not shown in the statements or hierarchies.	
authentication-key-chain	[edit protocols isis level<1   2>]
lsp-interval	[edit protocols isis interface]
lsp-lifetime	[edit protocols isis]
key-chain	[edit protocols isis interface bfd-liveness-detection authentication]

- Related Documentation
- *IS-IS Feature Guide for Routing Devices*
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

[edit protocols lacp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols lacp] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols lacp\] Hierarchy Level on page 411](#)
- [Unsupported Statements in the \[edit protocols lacp\] Hierarchy Level on page 412](#)

Supported Statements in the [edit protocols lacp] Hierarchy Level

The following hierarchy shows the [edit protocols lacp] configuration statements supported on EX Series switches:

```
protocols {
  lacp {

```

```
ppm {
  centralized
}
traceoptions {
  file <filename> <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
}
```

---

### Unsupported Statements in the [edit protocols lacp] Hierarchy Level

All statements in the [edit protocols lacp] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols link-management] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols link-management] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols link-management\] Hierarchy Level on page 412](#)
- [Unsupported Statements in the \[edit protocols link-management\] Hierarchy Level on page 413](#)

---

### Supported Statements in the [edit protocols link-management] Hierarchy Level

The following hierarchy shows the [edit protocols link-management] configuration statements supported on EX Series switches:

```
protocols {
  link-management {
    peer peer-name {
      address address;
      control-channel [ control-channel-interfaces ];
      lmp-control-channel interface-name {
        remote-address address;
      }
    }
  }
}
```

```

    }
    lmp-protocol {
        hello-dead-interval milliseconds;
        hello-interval milliseconds;
        passive;
        retransmission-interval milliseconds;
        retry-limit number;
    }
    te-link [ te-link-names ];
}
te-link te-link-name {
    disable;
    interface interface-name {
        disable;
        local-address address;
        remote-address address;
        remote-id id-number;
    }
    label-switched-path lsp-name {
        disable;
        local-address address;
        remote-address address;
        remote-id id-number;
    }
    local-address address;
    remote-address address;
    remote-id id-number;
    te-metric metric;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag;
}
}
}

```

### Unsupported Statements in the [edit protocols link-management] Hierarchy Level

All statements in the [edit protocols link-management] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols lldp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols lldp] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.

- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols lldp\] Hierarchy Level on page 414](#)
- [Unsupported Statements in the \[edit protocols lldp\] Hierarchy Level on page 414](#)

---

### Supported Statements in the [edit protocols lldp] Hierarchy Level

The following hierarchy shows the **[edit protocols lldp]** configuration statements supported on EX Series switches:

```
protocols {
  lldp {
    advertisement-interval seconds;
    disable;
    hold-multiplier seconds;
    interface (all | interface-name) {
      disable;
      power-negotiation {
        disable;
      }
    }
  }
  lldp-configuration-notification-interval seconds;
  management-address;
  netbios-snooping;
  ptopo-configuration-maximum-hold-time seconds;
  ptopo-configuration-trap-interval seconds;
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
  transmit-delay seconds;
}
```

---

### Unsupported Statements in the [edit protocols lldp] Hierarchy Level

All statements in the **[edit protocols lldp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols lldp-med] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols lldp-med]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols lldp-med\] Hierarchy Level on page 415](#)
- [Unsupported Statements in the \[edit protocols lldp-med\] Hierarchy Level on page 415](#)

### Supported Statements in the [edit protocols lldp-med] Hierarchy Level

The following hierarchy shows the **[edit protocols lldp-med]** configuration statements supported on EX Series switches:

```

protocols {
  lldp-med {
    disable;
    fast-start number;
    interface (all | interface-name) {
      disable;
      location {
        civic-based {
          ca-type {
            index {
              ca-value value;
            }
          }
          country-code code;
          what value;
        }
        elin number;
      }
    }
  }
}

```

### Unsupported Statements in the [edit protocols lldp-med] Hierarchy Level

All statements in the **[edit protocols lldp-med]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

- Related Documentation**
- [Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch](#)
  - [show lldp on page 2018](#)
  - [Configuring LLDP-MED \(CLI Procedure\) on page 1917](#)
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [\[edit protocols mld\]](#) Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [\[edit protocols mld\]](#) hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols mld\] Hierarchy Level on page 416](#)
- [Unsupported Statements in the \[edit protocols mld\] Hierarchy Level on page 417](#)

### [Supported Statements in the \[edit protocols mld\] Hierarchy Level](#)

---

The following hierarchy shows the [\[edit protocols mld\]](#) configuration statements supported on EX Series switches:

```
protocols {
  mld {
    accounting;
    interface interface-name {
      (accounting | no-accounting);
      disable;
      group-limit number;
      group-policy [ policy-names ];
      immediate-leave;
      oif-map [ map-names ];
      passive <allow-receive> <send-general-query> <send-group-query>;
      ssm-map ssm-map-name;
      ssm-map-policy policy-name;
      static {
        group multicast-group-address {
          exclude;
          group-count number;
          group-increment increment;
          source source-ip-address {
            source-count number;
            source-increment number;
          }
        }
      }
    }
  }
}
```



```

    }
  }
  version (1 | 2);
}
maximum-transmit-rate packets-per-second;
query-interval seconds;
query-last-member-interval seconds;
query-response-interval seconds;
robust-count number;
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
}
}

```

### Unsupported Statements in the [edit protocols mld] Hierarchy Level

All statements in the [edit protocols mld] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols mld-snooping] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols mld-snooping] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols mld-snooping\] Hierarchy Level on page 417](#)
- [Unsupported Statements in the \[edit protocols mld-snooping\] Hierarchy Level on page 418](#)

### Supported Statements in the [edit protocols mld-snooping] Hierarchy Level

The following hierarchy shows the [edit protocols mld-snooping] configuration statements supported on EX Series switches:

```

protocols {
  mld-snooping {
    traceoptions {

```

```
file filename <files number> <no-stamp> <replace> <size maximum-file-size>
  <world-readable | no-world-readable>;
flag flag <flag-modifier> <disable>;
}
vlan (all | vlan-identifier) {
  disable;
  immediate-leave;
  interface (all | interface-name) {
    multicast-router-interface;
    static {
      group multicast-ip-address;
    }
  }
  robust-count number;
  version number;
}
}
```

#### Unsupported Statements in the [edit protocols mld-snooping] Hierarchy Level

All statements in the [edit protocols mld-snooping] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- *Example: Configuring MLD Snooping*
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

### [edit protocols mpls] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols mpls] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols mpls\] Hierarchy Level on page 418](#)
- [Unsupported Statements in the \[edit protocols mpls\] Hierarchy Level on page 420](#)

#### Supported Statements in the [edit protocols mpls] Hierarchy Level

The following hierarchy shows the [edit protocols mpls] configuration statements supported on EX Series switches:

```
protocols {
```

```

mpls {
  class-of-service cos-value;
  disable;
  explicit-null;
  interface (interface-name | all) {
    disable;
  }
  ipv6-tunneling ;
  label-switched-path lsp-name {
    description text-string;
    disable;
    exclude-slr;
    from address;
    ldp-tunneling;
    no-cspf;
    no-decrement-ttl;
    oam {
      bfd-liveness-detection {
        detection-time {
          threshold milliseconds;
        }
        failure-action (make-before-break <teardown-timeout seconds> | teardown);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        version (1 | automatic);
      }
    }
  }
  to address;
}
no-cspf;
no-decrement-ttl;
no-propagate-ttl;
static-label-switched-path lsp-name {
  ingress {
    install {
      destination-prefix <active>;
    }
    next-hop (address | interface-name | address/interface-name);
    push out-label;
    to address;
  }
  transit incoming-label {
    description text-string;
    next-hop (address | interface-name | address/interface-name);
    pop;
    swap out-label;
  }
}
traffic-engineering (bgp | bgp-igp | bgp-igp-both-ribs | mpls-forwarding);
}
}

```

### Unsupported Statements in the [edit protocols mpls] Hierarchy Level

All statements in the [edit protocols mpls] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 34: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches**

Statement	Hierarchy
<b>NOTE:</b> Variables, such as <i>interface-name</i> , are not shown in the statements or hierarchies.	
active	[edit protocols mpls static-label-switched-path ingress install] [edit protocols mpls label-switched-path install]
adaptive	[edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
adjust-interval	[edit protocols mpls label-switched-path auto-bandwidth]
adjust-threshold	[edit protocols mpls label-switched-path auto-bandwidth]
adjust-threshold-overflow-limit	[edit protocols mpls label-switched-path auto-bandwidth]
adjust-threshold-underflow-limit	[edit protocols mpls label-switched-path auto-bandwidth]
admin-down	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
admin-group	[edit protocols mpls] [edit protocols mpls interface] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
admin-group-extended	[edit protocols mpls] [edit protocols mpls interface] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
admin-groups	[edit protocols mpls]
advertisement-hold-time	[edit protocols mpls]
allow-fragmentation	[edit protocols mpls path-mtu]
always-mark-connection-protection-tlv	[edit protocols mpls interface]
associate-backup-pe-groups	[edit protocols mpls label-switched-path]

Table 34: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
auto-bandwidth	[edit protocols mpls label-switched-path] [edit protocols mpls statistics]
auto-policing	[edit protocols mpls]
backup	[edit protocols mpls label-switched-path]
bandwidth	[edit protocols mpls] [edit protocols mpls bandwidth] [edit protocols mpls label-switched-path] [edit protocols mpls static-label-switched-path bypass] [edit protocols mpls label-switched-path fast-reroute] [edit protocols mpls static-label-switched-path ingress] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary] [edit protocols mpls static-label-switched-path transit]
bandwidth-model	[edit protocols mpls diffserv-te]
bandwidth-percent	[edit protocols mpls label-switched-path fast-reroute]
bfd-liveness-detection	[edit protocols mpls label-switched-path primary oam] [edit protocols mpls label-switched-path secondary oam] [edit protocols mpls oam]
bypass	[edit protocols mpls static-label-switched-path]
bypass-name	[edit protocols mpls static-label-switched-path ingress link-protection] [edit protocols mpls static-label-switched-path ingress node-protection] [edit protocols mpls static-label-switched-path transit link-protection] [edit protocols mpls static-label-switched-path transit node-protection]
class	[edit protocols mpls auto-policing]
class-of-service	[edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary] [edit protocols mpls static-label-switched-path ingress]
context-identifier	[edit protocols mpls egress-protection]
ct0	[edit protocols mpls bandwidth] [edit protocols mpls label-switched-path bandwidth] [edit protocols mpls label-switched-path primary bandwidth] [edit protocols mpls label-switched-path secondary bandwidth]

Table 34: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
ct1	[edit protocols mpls bandwidth] [edit protocols mpls label-switched-path bandwidth] [edit protocols mpls label-switched-path primary bandwidth] [edit protocols mpls label-switched-path secondary bandwidth]
ct2	[edit protocols mpls bandwidth] [edit protocols mpls label-switched-path bandwidth] [edit protocols mpls label-switched-path primary bandwidth] [edit protocols mpls label-switched-path secondary bandwidth]
ct3	[edit protocols mpls bandwidth] [edit protocols mpls label-switched-path bandwidth] [edit protocols mpls label-switched-path primary bandwidth] [edit protocols mpls label-switched-path secondary bandwidth]
description	[edit protocols mpls static-label-switched-path bypass] [edit protocols mpls static-label-switched-path ingress]
detection-time	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
diffserv-te	[edit protocols mpls]
drop	[edit protocols mpls auto-policing class]
egress-protection	[edit protocols mpls] [edit protocols mpls label-switched-path]
encoding-type	[edit protocols mpls label-switched-path lsp-attributes]
exclude	[edit protocols mpls admin-group] [edit protocols mpls label-switched-path admin-group] [edit protocols mpls label-switched-path admin-group-extended] [edit protocols mpls label-switched-path primary admin-group] [edit protocols mpls label-switched-path primary admin-group-extended] [edit protocols mpls label-switched-path secondary admin-group] [edit protocols mpls label-switched-path secondary admin-group-extended] [edit protocols mpls label-switched-path fast-reroute]
exclude-slrq	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]

Table 34: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
expand-loose-hop	[edit protocols mpls]
failure-action	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
fast-reroute	[edit protocols mpls label-switched-path]
file	[edit protocols mpls label-switched-path primary oam traceoptions] [edit protocols mpls label-switched-path secondary oam traceoptions] [edit protocols mpls label-switched-path traceoptions] [edit protocols mpls statistics] [edit protocols mpls traceoptions]
files	[edit protocols mpls statistics file]
filter	[edit protocols mpls static-label-switched-path ingress policing]
flag	[edit protocols mpls label-switched-path primary oam traceoptions] [edit protocols mpls label-switched-path secondary oam traceoptions] [edit protocols mpls label-switched-path traceoptions] [edit protocols mpls traceoptions]
gpid	[edit protocols mpls label-switched-path lsp-attributes]
hop-limit	[edit protocols mpls] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary] [edit protocols mpls label-switched-path fast-reroute] [edit protocols mpls label-switched-path]
icmp-tunneling	[edit protocols mpls]
include-all	[edit protocols mpls admin-group] [edit protocols mpls admin-group-extended] [edit protocols mpls label-switched-path admin-group] [edit protocols mpls label-switched-path admin-group-extended] [edit protocols mpls label-switched-path fast-reroute] [edit protocols mpls label-switched-path primary admin-group] [edit protocols mpls label-switched-path primary admin-group-extended] [edit protocols mpls label-switched-path secondary admin-group] [edit protocols mpls label-switched-path secondary admin-group-extended]

Table 34: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
include-any	[edit protocols mpls admin-group] [edit protocols mpls admin-group-extended] [edit protocols mpls label-switched-path admin-group] [edit protocols mpls label-switched-path admin-group-extended] [edit protocols mpls label-switched-path fast-reroute] [edit protocols mpls label-switched-path primary admin-group] [edit protocols mpls label-switched-path primary admin-group-extended] [edit protocols mpls label-switched-path secondary admin-group] [edit protocols mpls label-switched-path secondary admin-group-extended]
install	[edit protocols mpls label-switched-path]
inter-domain	[edit protocols mpls label-switched-path]
interval	[edit protocols mpls statistics]
least-fill	[edit protocols mpls label-switched-path]
link-protection	[edit protocols mpls label-switched-path] [edit protocols mpls static-label-switched-path ingress] [edit protocols mpls static-label-switched-path transit]
log-updown	[edit protocols mpls]
loss-priority-high	[edit protocols mpls auto-policing class]
loss-priority-low	[edit protocols mpls auto-policing class]
lsp-attributes	[edit protocols mpls label-switched-path]
make-before-break	[edit protocols mpls label-switched-path secondary oam bfd-liveness-detection failure-action]
maximum-bandwidth	[edit protocols mpls label-switched-path auto-bandwidth]
metric	[edit protocols mpls egress-protection context-identifier] [edit protocols mpls label-switched-path] [edit protocols mpls static-label-switched-path ingress]
mib-mpls-show-p2mp	[edit protocols mpls]
mimum-bandwidth	[edit protocols mpls label-switched-path auto-bandwidth]



Table 34: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
minimum-interval	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path primary oam bfd-liveness-detection transit-interval] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection transit-interval]
minimum-receive-interval	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
monitor-bandwidth	[edit protocols mpls label-switched-path auto-bandwidth]
most-fill	[edit protocols mpls label-switched-path]
mpls-lsp-traps	[edit protocols mpls log-updown no-trap]
mpls-tp-mode	[edit protocols mpls oam]
mtu-signaling	[edit protocols mpls path-mtu rsvp]
multiplier	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
next-hop	[edit protocols mpls path] [edit protocols mpls static-label-switched-path bypass]
next-next-label	[edit protocols mpls static-label-switched-path ingress node-protection] [edit protocols mpls static-label-switched-path transit node-protection]
no-adaptation	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
no-auto-policing	[edit protocols mpls static-label-switched-path ingress policing]
no-cspf	[edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
no-decrement-ttl	[edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]

Table 34: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
node-link-protection	[edit protocols mpls label-switched-path]
node-protection	[edit protocols mpls static-label-switched-path ingress] [edit protocols mpls static-label-switched-path transit]
no-exclude	[edit protocols mpls label-switched-path fast-reroute]
no-include-all	[edit protocols mpls label-switched-path fast-reroute]
no-include-any	[edit protocols mpls label-switched-path fast-reroute]
no-install-to-address	[edit protocols mpls label-switched-path] [edit protocols mpls static-label-switched-path ingress]
no-record	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
no-remote-trace	[edit protocols mpls label-switched-path oam traceoptions] [edit protocols mpls label-switched-path secondary oam traceoptions]
no-syslog	[edit protocols mpls log-updown]
no-trap	[edit protocols mpls log-updown]
no-world-readable	[edit protocols mpls statistics file]
number	[edit protocols mpls auto-policing class]
oam	[edit protocols mpls] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
optimize-aggressive	[edit protocols mpls]
optimize-hold-dead-delay	[edit protocols mpls]
optimize-switchover-delay	[edit protocols mpls]
optimize-timer	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
path	[edit protocols mpls]

Table 34: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
path-mtu	[edit protocols mpls]
policing	[edit protocols mpls static-label-switched-path ingress] [edit protocols mpls label-switched-path]
preference	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls static-label-switched-path ingress] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
primary	[edit protocols mpls egress-protection context-identifier]
priority	[edit protocols mpls] [edit protocols mpls diffserv-te te-class-matrix tex] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
protection-revert-time	[edit protocols mpls interface static]
protector	[edit protocols mpls egress-protection context-identifier]
push	[edit protocols mpls static-label-switched-path bypass]
random	[edit protocols mpls label-switched-path]
record	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
retry-limit	[edit protocols mpls label-switched-path]
retry-timer	[edit protocols mpls label-switched-path]
revert-timer	[edit protocols mpls] [edit protocols mpls label-switched-path]
rfc3812-traps	[edit protocols mpls log-updown no-trap]
rsvp	[edit protocols mpls path-mtu]
rsvp-error-hold-time	[edit protocols mpls]
select	[edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]

Table 34: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
signal-bandwidth	[edit protocols mpls label-switched-path lsp-attributes]
size	[edit protocols mpls statistics file]
smart-optimize-timer	[edit protocols mpls]
soft-preemption	[edit protocols mpls label-switched-path]
srlg	[edit protocols mpls interface]
standby	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
static	[edit protocols mpls interface]
statistics	[edit protocols mpls]
switch-away-lsps	[edit protocols mpls interface]
switching-type	[edit protocols mpls label-switched-path lsp-attributes]
syslog	[edit protocols mpls log-updown]
tex	[edit protocols mpls diffserv-te te-class-matrix]
teardown	[edit protocols mpls label-switched-path secondary oam bfd-liveness-detection failure-action]
te-class-matrix	[edit protocols mpls diffserv-te]
template	[edit protocols mpls label-switched-path]
threshold	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection transmit-interval] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection detection-time] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection transmit-interval]
to	[edit protocols mpls static-label-switched-path bypass]
traceoptions	[edit protocols mpls] [edit protocols mpls label-switched-path ] [edit protocols mpls label-switched-path oam] [edit protocols mpls label-switched-path primary oam] [edit protocols mpls label-switched-path secondary oam]

Table 34: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
traffic-class	[edit protocols mpls diffserv-te te-class-matrix tex]
transmit-interval	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
trap	[edit protocols mpls log-updown]
trap-path-down	[edit protocols mpls log-updown]
trap-path-up	[edit protocols mpls log-updown]
version	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
world-readable	[edit protocols mpls statistics file]

**Related Documentation**

- *Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure)*
- *Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)*
- *Configuring MPLS on Provider Switches (CLI Procedure)*
- *Junos OS MPLS for EX Series Switches Overview*
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols msdp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols msdp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols msdp\] Hierarchy Level on page 430](#)
- [Unsupported Statements in the \[edit protocols msdp\] Hierarchy Level on page 431](#)

### Supported Statements in the [edit protocols msdp] Hierarchy Level

The following hierarchy shows the [edit protocols msdp] configuration statements supported on EX Series switches:

```
protocols {
  msdp {
    active-source-limit {
      log-interval seconds;
      log-warning value;
      maximum number;
      threshold number;
    }
    data-encapsulation (disable | enable);
    disable;
    export [ policy-names ];
    group group-name {
      disable;
      export [ policy-names ];
      import [ policy-names ];
      local-address address;
      mode (mesh-group | standard);
      peer address {
        active-source-limit {
          log-interval seconds;
          log-warning value;
          maximum number;
          threshold number;
        }
        authentication-key peer-key;
        default-peer;
        disable;
        export [ policy-names ];
        hold-time seconds;
        import [ policy-names ];
        keep-alive seconds;
        local-address address;
        sa-hold-time seconds;
        traceoptions {
          file filename <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
          flag flag <flag-modifier> <disable>;
        }
      }
    }
    traceoptions {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
  hold-time seconds;
  import [ policy-names ];
  keep-alive seconds;
  local-address address;
  peer address {
```

```

    active-source-limit {
        log-interval seconds;
        log-warning value;
        maximum number;
        threshold number;
    }
    authentication-key peer-key;
    default-peer;
    disable;
    export [ policy-names ];
    hold-time seconds;
    import [ policy-names ];
    keep-alive seconds;
    local-address address;
    sa-hold-time seconds;
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
rib-group group-name;
sa-hold-time seconds;
source ip-prefix </prefix-length> {
    active-source-limit {
        log-interval seconds;
        log-warning value;
        maximum number;
        threshold number;
    }
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
}

```

### Unsupported Statements in the [edit protocols msdp] Hierarchy Level

All statements in the [edit protocols msdp] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

### [edit protocols mstp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols mstp] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.

- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols mstp\] Hierarchy Level on page 432](#)
- [Unsupported Statements in the \[edit protocols mstp\] Hierarchy Level on page 433](#)

### Supported Statements in the [edit protocols mstp] Hierarchy Level

The following hierarchy shows the **[edit protocols mstp]** configuration statements supported on EX Series switches:

```
protocols {
  mstp {
    bpdu-block-on-edge;
    bpdu-destination-mac-address provider-bridge-group;
    bridge-priority priority;
    configuration-name configuration-name;
    disable;
    forward-delay seconds;
    hello-time seconds;
    interface interface-name {
      access-trunk;
      bpdu-timeout-action {
        alarm;
        block;
      }
      cost cost;
      edge;
      mode (point-to-point | shared);
      no-root-port;
      priority interface-priority;
    }
    max-age seconds;
    max-hops hops;
    msti identifier {
      bridge-priority priority;
      interface interface-name {
        cost cost;
        priority interface-priority;
      }
      vlan [ vlan-ids ];
    }
    priority-hold-time seconds;
    revision-level revision-level;
    traceoptions {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      flag flag <disable>;
    }
  }
}
```



```
}
}
```

### Unsupported Statements in the [edit protocols mstp] Hierarchy Level

All statements in the **[edit protocols mstp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols mvrp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols mvrp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols mvrp\] Hierarchy Level on page 433](#)
- [Unsupported Statements in the \[edit protocols mvrp\] Hierarchy Level on page 434](#)

### Supported Statements in the [edit protocols mvrp] Hierarchy Level

The following hierarchy shows the **[edit protocols mvrp]** configuration statements supported on EX Series switches:

```
protocols {
  mvrp {
    interface (all | interface-name) {
      join-timer milliseconds;
      leave-timer milliseconds;
      leaveall-timer milliseconds;
      registration (forbidden | normal);
    }
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    no-dynamic-vlan;
    traceoptions {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      flag flag <disable>;
```

```

    }
  }
}

```

### Unsupported Statements in the [edit protocols mvrp] Hierarchy Level

All statements in the **[edit protocols mvrp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 35: Unsupported [edit protocols mvrp] Configuration Statements on EX Series Switches**

Statement	Hierarchy
bpdu-destination-mac-address	[edit protocols mvrp]
point-to-point	[edit protocols mvrp interface <i>interface-name</i> ]
registration, restricted option	[edit protocols mvrp interface <i>interface-name</i> ]

- Related Documentation**
- *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches*
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

### [edit protocols neighbor-discovery] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols neighbor-discovery]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols neighbor-discovery\] Hierarchy Level on page 434](#)
- [Unsupported Statements in the \[edit neighbor-discovery\] Hierarchy Level on page 435](#)

### Supported Statements in the [edit protocols neighbor-discovery] Hierarchy Level

The following hierarchy shows the **[edit protocols neighbor-discovery]** configuration statements supported on EX Series switches:

```

protocols {
  neighbor-discovery {
    no-dad-on-state-change ;
    onlink-subnet-only;
  }
}

```

```
}
```

### Unsupported Statements in the [edit neighbor-discovery] Hierarchy Level

All statements in the **[edit protocols neighbor-discovery]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols oam] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols oam]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols oam\] Hierarchy Level on page 435](#)
- [Unsupported Statements in the \[edit protocols oam\] Hierarchy Level on page 437](#)

### Supported Statements in the [edit protocols oam] Hierarchy Level

The following hierarchy shows the **[edit protocols oam]** configuration statements supported on EX Series switches:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        action-profile profile-name {
          action {
            interface-down;
          }
          default-actions {
            interface-down;
          }
          event {
            adjacency-loss;
          }
        }
      }
      esp-traceoptions {
        file filename <files number> <no-stamp> <replace> <size size> <world-readable>
          | no-world-readable;
        flag (all | error | esp | interface | krt | lib | normal | task | timer);
      }
    }
  }
}
```

```

}
linktrace {
  age (30m | 10m | 1m | 30s | 10s);
  path-database-size path-database-size;
}
maintenance-domain domain-name {
  level number;
  maintenance-association ma-name {
    continuity-check {
      hold-interval minutes;
      interface-status-tlv;
      interval (10m | 10s | 1m | 1s | 100ms);
      loss-threshold number;
      port-status-tlv;
    }
    mep mep-id {
      auto-discovery;
      direction down;
      interface interface-name {
        vlan-id identifier;
      }
      priority number;
      remote-mep mep-id {
        action-profile profile-name;
        sla-iterator-profile profile-name {
          data-tlv-size size;
          iteration-count count-value;
          priority priority-value;
        }
      }
    }
  }
  short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
}
mip-half-function (none | default | explicit);
name-format (character-string | none | dns | mac+2oct);
vlan-name name;
}
performance-monitoring {
  no-delegate-processing;
  sla-iterator-profiles {
    profile-name {
      calculation-weight {
        delay delay-value;
        delay-variation delay-variation-value;
      }
      cycle-time cycle-time-value;
      iteration-period iteration-period-value;
      measurement-type two-way-delay;
      passive;
    }
  }
}
}
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
}

```

```

        no-remote-trace;
    }
}
link-fault-management {
    action-profile profile-name {
        action {
            link-down;
            syslog;
        }
        event {
            link-adjacency-loss;
            link-event-rate {
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
            }
        }
    }
}
interface interface-name {
    apply-action-profile profile-name;
    event-thresholds {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
    }
    link-discovery (active | passive);
    negotiation-options {
        allow-remote-loopback;
        no-allow-link-events;
    }
    pdu-interval interval;
    pdu-threshold threshold-value;
    remote-loopback;
}
traceoptions {
    file filename <files number> <match regex> <size size> <world-readable |
        no-world-readable>;
    flag flag;
    no-remote-trace;
}
}
}
}
}

```

### Unsupported Statements in the [edit protocols oam] Hierarchy Level

All statements in the [edit protocols oam] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols ospf] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols ospf]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols ospf\] Hierarchy Level on page 438](#)
- [Unsupported Statements in the \[edit protocols ospf\] Hierarchy Level on page 440](#)

### Supported Statements in the [edit protocols ospf] Hierarchy Level

The following hierarchy shows the **[edit protocols ospf]** configuration statements supported on EX Series switches:

```
protocols {
  ospf {
    area area-id {
      area-range ip-prefix</prefix-length> <exact> <override-metric metric> <restrict>;
      context-identifier
      interface interface-name {
        authentication {
          md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
          simple-password key-string;
        }
        bandwidth-based-metrics {
          bandwidth value metric number;
        }
        bfd-liveness-detection {
          authentication {
            algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
              meticulous-keyed-sha-1 | simple-password);
            loose-check;
          }
          detection-time {
            threshold milliseconds;
          }
          full-neighbors-only;
          minimum-interval milliseconds;
          minimum-receive-interval milliseconds;
          multiplier number;
          no-adaptation;
          transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
          }
        }
      }
    }
  }
}
```

```

    }
    version (0|1 | automatic);
}
dead-interval seconds;
disable;
dynamic-neighbors;
flood-reduction;
hello-interval seconds;
interface-type (nbma | p2mp | p2p);
ipsec-sa sa-name;
(link-protection | node-link-protection);
metric metric;
no-eligible-backup;
no-interface-state-traps;
no-neighbor-down-notification;
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
poll-interval seconds;
priority number;
retransmit-interval seconds;
secondary;
te-metric metric;
transit-delay seconds;
}
network-summary-export [ policy-names ];
network-summary-import [ policy-names ];
no-context-identifier-advertisement;
nssa {
    area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
    default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
    }
    (summaries | no-summaries);
}
stub <default-metric metric> <summaries | no-summaries>;
virtual-link neighbor-id router-id transit-area area-id;
}
backup-spf-options
    disable;
    downstream-paths-only;
    no-install;
}
database-protection {
    ignore-count number;
    ignore-time seconds;
    maximum-lsa number;
    reset-time seconds;
    warning-only;
    warning-threshold percent;
}
disable;

```

```

export [ policy-names ];
external-preference preference;
graceful-restart {
    disable;
    helper-disable <both | restart-signaling | standard>;
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
import [ policy-names ];
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
reference-bandwidth reference-bandwidth;
rib-group group-name;
spf-options {
    delay milliseconds;
    holddown milliseconds;
    no-ignore-our-externals;
    rapid-runs number;
}
topology {
    disable;
    rib-group group-name;
    topology-id number;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    advertise-unnumbered-interfaces;
    credibility-protocol-preference;
    ignore-lsp-metrics;
    multicast-rpf-routes;
    no-topology;
    shortcuts <lsp-metric-into-summary>;
}
}
}

```

### Unsupported Statements in the [edit protocols ospf] Hierarchy Level

All statements in the [edit protocols ospf] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 36: Unsupported [edit protocols ospf] Configuration Statements on EX Series Switches**

Statement	Hierarchy
-----------	-----------

**NOTE:** Variables, such as *interface-name*, are not shown in the statements or hierarchies.



Table 36: Unsupported `[edit protocols ospf]` Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
<code>backup-spf-options</code>	<code>[edit protocols ospf topology]</code>
<code>key-chain</code>	<code>[edit protocols ospf area interface bfd-liveness-detection authentication]</code>
<code>overload</code>	<code>[edit protocols ospf topology]</code>
<code>prefix-export-limit</code>	<code>[edit protocols ospf topology]</code>
<code>spf-options</code>	<code>[edit protocols ospf topology]</code>
<code>topology</code>	<code>[edit protocols ospf area interface]</code>

- Related Documentation**
- *OSPF Feature Guide for Routing Devices*
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## `[edit protocols ospf3]` Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the `[edit protocols ospf3]` hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the `\[edit protocols ospf3\]` Hierarchy Level on page 441](#)
- [Unsupported Statements in the `\[edit protocols ospf3\]` Hierarchy Level on page 444](#)

### Supported Statements in the `[edit protocols ospf3]` Hierarchy Level

The following hierarchy shows the `[edit protocols ospf3]` configuration statements supported on EX Series switches:

```
protocols {
  ospf3 {
    area area-id {
      area-range ip-prefix</prefix-length> <exact> <override-metric metric> <restrict>;
      context-identifier
      inter-area-prefix-export [ policy-names ];
      inter-area-prefix-import [ policy-names ];
```

```

interface interface-name {
  bandwidth-based-metrics {
    bandwidth value metric number;
  }
  bfd-liveness-detection {
    authentication {
      algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
        meticulous-keyed-sha-1 | simple-password);
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    full-neighbors-only;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (0|1 | automatic);
  }
  dead-interval seconds;
  disable;
  flood-reduction;
  hello-interval seconds;
  interface-type (p2mp-over-lan | p2p);
  ipsec-sa sa-name;
  (link-protection | node-link-protection);
  metric metric;
  no-eligible-backup;
  passive {
    traffic-engineering {
      remote-node-id address;
    }
  }
  priority number;
  retransmit-interval seconds;
  secondary;
  transit-delay seconds;
}
no-context-identifier-advertisement;
nssa {
  area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
  default-lsa {
    default-metric metric;
    metric-type type;
    type-7;
  }
  (summaries | no-summaries);
}
stub <default-metric metric> <summaries | no-summaries>;
}
backup-spf-options (disable | downstream-paths-only | no-install);

```

```

database-protection {
    ignore-count number;
    ignore-time seconds;
    maximum-lsa number;
    reset-time seconds;
    warning-only;
    warning-threshold percent;
}
disable;
export [ policy-names ];
external-preference preference;
graceful-restart {
    disable;
    helper-disable;
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
import [ policy-names ];
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
realm (ipv4-multicast| ipv6-multicast) {
    ... same statements as at the [edit protocols ospf3] hierarchy level, EXCEPT FOR ...
    area area-id {
        interface interface-name {
            no-eligible-backup; # NOT valid at this level
        }
    }
    backup-spf-options { ... } # NOT valid at this level
    realm realm-identifier { ... } # NOT valid at this level
    traffic-engineering { ... } # NOT valid at this level
}
realm ipv4-unicast {
    ... same statements as at the [edit protocols ospf3] hierarchy level, PLUS ...
    area area-id {
        interface interface-name {
        }
    }
}
realm ipv6-unicast {
    ... same statements as at the [edit protocols ospf3] hierarchy level, PLUS ...
    disable;
    backup-spf-options {
        disable;
        downstream-paths-only;
        no-install;
    }
}
reference-bandwidth reference-bandwidth;
rib-group group-name;
spf-options {
    delay milliseconds;
    holddown milliseconds;
    no-ignore-our-externals;
}

```

```

        rapid-runs number;
    }
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    traffic-engineering {
        ignore-lsp-metrics;
        shortcuts <lsp-metric-into-summary>;
    }
}

```

### Unsupported Statements in the [edit protocols ospf3] Hierarchy Level

All statements in the [edit protocols ospf3] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exception:

**Table 37: Unsupported [edit protocols ospf 3] Configuration Statements on EX Series Switches**

Statement	Hierarchy
NOTE: Variables, such as <i>interface-name</i> , are not shown in the statements or hierarchies.	
key-chain	[edit protocols ospf3 area interface bfd-liveness-detection authentication]

- Related Documentation**
- *OSPF Feature Guide for Routing Devices*
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

### [edit protocols pim] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols pim] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols pim\] Hierarchy Level on page 445](#)
- [Unsupported Statements in the \[edit protocols pim\] Hierarchy Level on page 447](#)

### Supported Statements in the [edit protocols pim] Hierarchy Level

The following hierarchy shows the [edit protocols pim] configuration statements supported on EX Series switches:

```

protocols {
  pim {
    assert-timeout seconds;
    default-vpn-source {
      interface-name interface-name;
    }
    dense-groups {
      address <announce | reject>;
    }
    disable;
    dr-election-on-p2p;
    export [ policy-names ];
    family (inet | inet6) {
      disable;
    }
    graceful-restart {
      disable;
      restart-duration seconds;
    }
    import [ policy-names ];
    interface interface-name {
      accept-remote-source;
      disable;
      bfd-liveness-detection {
        authentication {
          algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
            meticulous-keyed-sha-1 | simple-password);
          key-chain key-chain-name;
          loose-check;
        }
        detection-time {
          threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        version (1 | automatic);
      }
      disable;
      family (inet | inet6) {
        disable;
      }
      hello-interval seconds;
      mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse |
        sparse-dense);
    }
  }
}

```

```
neighbor-policy [ policy-names ];
override-interval milliseconds;
priority number;
propagation-delay milliseconds;
reset-tracking-bit;
version (1 | 2);
}
join-load-balance;
join-prune-timeout seconds;
mpls-internet-multicast;
nexthop-hold-time time;
no-wildcard-register-stop;
nonstop-routing {
    disable;
}
override-interval milliseconds;
propagation-delay milliseconds;
reset-tracking-bit;
rib-group {
    inet group-name;
    inet6 group-name;
}
rp {
    auto-rp {
        (announce | discovery | mapping);
        (mapping-agent-election | no-mapping-agent-election);
    }
    bootstrap {
        family (inet | inet6) {
            export [ policy-names ];
            import [ policy-names ];
            priority number;
        }
    }
    bootstrap-export [ policy-names ];
    bootstrap-import [ policy-names ];
    bootstrap-priority number;
    dr-register-policy [ policy-names ];
    embedded-rp {
        group-ranges {
            ip-prefix </prefix-length>;
        }
        maximum-rps limit;
    }
    local {
        address address;
        disable;
        family (inet | inet6) {
            address address;
            anycast-pim {
                local-address address;
                rp-set {
                    address address <forward-msdp-sa>;
                }
            }
        }
        disable;
    }
}
```

```

    group-ranges {
        ip-prefix </prefix-length>;
    }
    hold-time seconds;
    override;
    priority number;
}
group-ranges {
    ip-prefix </prefix-length>;
}
hold-time seconds;
override;
priority number;
}
register-probe-time time;
rp-register-policy [ policy-names ];
static {
    address address {
        group-ranges {
            ip-prefix </prefix-length>;
        }
        override;
        version (1 | 2);
    }
}
}
spt-threshold {
    infinity [ policy-names ];
}
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
no-world-readable>;
    flag flag <flag-modifier> <disable>;
    flag (route | state) <flag-modifier> <disable> <filter <match-on prefix>
<policy [ policy-names ]>>;
}
}
}

```

### Unsupported Statements in the [edit protocols pim] Hierarchy Level

All statements in the [edit protocols pim] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- *Multicast Protocols Feature Guide for Routing Devices*
- [edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 389

## [edit protocols rip] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols rip] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols rip\] Hierarchy Level on page 448](#)
- [Unsupported Statements in the \[edit protocols rip\] Hierarchy Level on page 449](#)

---

### Supported Statements in the [edit protocols rip] Hierarchy Level

---

The following hierarchy shows the [edit protocols rip] configuration statements supported on EX Series switches:

```
protocols {
  rip {
    authentication-key password;
    authentication-type type;
    (check-zero | no-check-zero);
    group group-name {
      bfd-liveness-detection {
        authentication {
          algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
            meticulous-keyed-sha-1 | simple-password);
          loose-check;
        }
        detection-time {
          threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        version (1 | automatic);
      }
    }
    export [ policy-names ];
    import [ policy-names ];
    metric-out metric;
    neighbor neighbor-name {
      any-sender;
      authentication-key password;
      authentication-type type;
      bfd-liveness-detection {
        authentication {
```



```

        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
        meticulous-keyed-sha-1 | simple-password);
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
(check-zero | no-check-zero);
import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
route-timeout seconds;
send (broadcast | multicast | none | version-1);
update-interval seconds;
}
preference preference;
route-timeout seconds;
update-interval seconds;
}
graceful-restart {
    disable;
    restart-time seconds;
}
holddown seconds;
import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
rib-group group-name;
route-timeout seconds;
send (broadcast | multicast | none | version-1);
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
update-interval seconds;
}
}

```

### Unsupported Statements in the [edit protocols rip] Hierarchy Level

All statements in the [edit protocols rip] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 38: Unsupported [edit protocols-rip] Configuration Statements on EX Series Switches

Statement	Hierarchy
demand-circuit	[edit protocols rip group] [edit protocols rip group neighbor]
key-chain	[edit protocols rip group bfd-liveness-detection authentication] [edit protocols rip group neighbor bfd-liveness-detection authentication]
max-retrans-time	[edit protocols rip group] [edit protocols rip group neighbor]

NOTE: Variables, such as *group-name*, are not shown in the statements or hierarchies.

- Related Documentation**
- [RIP Feature Guide for Routing Devices](#)
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols ripng] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols ripng]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols ripng\] Hierarchy Level on page 450](#)
- [Unsupported Statements in the \[edit protocols ripng\] Hierarchy Level on page 451](#)

### Supported Statements in the [edit protocols ripng] Hierarchy Level

The following hierarchy shows the **[edit protocols ripng]** configuration statements supported on EX Series switches:

```

protocols {
  ripng {
    graceful-restart {
      disable;
      restart-time seconds;
    }
    group group-name {
      export [ policy-names ];
      import [ policy-names ];
      metric-out metric;
    }
  }
}

```

```

neighbor neighbor-name {
  import [ policy-names ];
  metric-in metric;
  receive <none>;
  route-timeout seconds;
  send <none>;
  update-interval seconds;
}
preference number;
route-timeout seconds;
update-interval seconds;
}
holddown seconds;
import [ policy-names ];
metric-in metric;
receive <none>;
route-timeout seconds;
send <none>;
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
update-interval seconds;
}
}

```

### Unsupported Statements in the [edit protocols ripng] Hierarchy Level

All statements in the [edit protocols ripng] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- *RIPng Feature Guide for Routing Devices*
- [edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 389

## [edit protocols router-advertisement] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols router-advertisement] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols router-advertisement\] Hierarchy Level on page 452](#)
- [Unsupported Statements in the \[edit protocols router-advertisement\] Hierarchy Level on page 452](#)

### [Supported Statements in the \[edit protocols router-advertisement\] Hierarchy Level](#)

---

The following hierarchy shows the **[edit protocols router-advertisement]** configuration statements supported on EX Series switches:

```
protocols {
  router-advertisement {
    interface interface-name {
      current-hop-limit number;
      default-lifetime seconds;
      (link-mtu | no-link-mtu);
      (managed-configuration | no-managed-configuration);
      max-advertisement-interval seconds;
      min-advertisement-interval seconds;
      (other-stateful-configuration | no-other-stateful-configuration);
      prefix prefix {
        (autonomous | no-autonomous);
        (on-link | no-on-link);
        preferred-lifetime seconds;
        valid-lifetime seconds;
      }
      reachable-time milliseconds;
      retransmit-timer milliseconds;
      virtual-router-only;
    }
    traceoptions {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      flag flag;
    }
  }
}
```

### [Unsupported Statements in the \[edit protocols router-advertisement\] Hierarchy Level](#)

---

All statements in the **[edit protocols router-advertisement]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [\[edit protocols router-discovery\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit protocols router-discovery]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols router-discovery\] Hierarchy Level on page 453](#)
- [Unsupported Statements in the \[edit protocols router-discovery\] Hierarchy Level on page 453](#)

### [Supported Statements in the \[edit protocols router-discovery\] Hierarchy Level](#)

The following hierarchy shows the **[edit protocols router-discovery]** configuration statements supported on EX Series switches:

```
protocols {
  router-discovery {
    address address {
      (advertise | (broadcast | multicast ) |ignore);
      (ineligible | priority number);
    }
    disable;
    interface interface-name {
      lifetime seconds;
      max-advertisement-interval seconds;
      min-advertisement-interval seconds;
    }
    traceoptions {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      flag flag;
    }
  }
}
```

### [Unsupported Statements in the \[edit protocols router-discovery\] Hierarchy Level](#)

All statements in the **[edit protocols router-discovery]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [\[edit protocols rstp\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit protocols rstp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols rstp\] Hierarchy Level on page 454](#)
- [Unsupported Statements in the \[edit protocols rstp\] Hierarchy Level on page 454](#)

---

### Supported Statements in the [edit protocols rstp] Hierarchy Level

The following hierarchy shows the **[edit protocols rstp]** configuration statements supported on EX Series switches:

```
protocols {
  rstp {
    bpdu-block-on-edge;
    bridge-priority priority;
    disable;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      arp-on-stp;
      bpdu-timeout-action {
        block;
        log;
      }
      cost cost;
      disable;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    traceoptions {
      file filename <files number > <size size > <no-stamp | no-world-readable |
        world-readable>;
      flag flag;
    }
  }
}
```

---

### Unsupported Statements in the [edit protocols rstp] Hierarchy Level

All statements in the **[edit protocols rstp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

- Related Documentation**
- [Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919](#)
  - [Understanding RSTP for EX Series Switches on page 4906](#)
  - [show spanning-tree bridge on page 5035](#)
  - [show spanning-tree interface on page 5045](#)
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [\[edit protocols rsvp\]](#) Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [\[edit protocols rsvp\]](#) hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols rsvp\] Hierarchy Level on page 455](#)
- [Unsupported Statements in the \[edit protocols rsvp\] Hierarchy Level on page 456](#)

### [Supported Statements in the \[edit protocols rsvp\] Hierarchy Level](#)

The following hierarchy shows the [\[edit protocols rsvp\]](#) configuration statements supported on EX Series switches.

```
protocols {
  rsvp {
    disable;
    hello-acknowledgements;
    interface interface-name {
      (aggregate | no-aggregate);
      authentication-key key;
      disable;
      hello-interval seconds;
      (reliable | no-reliable);
    }
    keep-multiplier number;
    load-balance bandwidth;
    no-interface-hello;
    no-local-reversion;
    no-p2mp-sublsp;
    node-hello;
    refresh-time seconds;
    setup-protection;
```

```

    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}

```

### Unsupported Statements in the [edit protocols rsvp] Hierarchy Level

All statements in the [edit protocols rsvp] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 39: Unsupported [edit protocols rsvp] Configuration Statements on EX Series Switches**

Statement	Hierarchy
<b>NOTE:</b> Variables, such as <i>interface-name</i> , are not shown in the statements or hierarchies.	
aggregate	[edit protocols rsvp peer-interface]
aggressive	[edit protocols rsvp preemption]
admin-group	[edit protocols rsvp interface link-protection] [edit protocols rsvp interface link-protection bypass]
authentication-key	[edit protocols rsvp peer-interface]
bandwidth	[edit protocols rsvp interface] [edit protocols rsvp interface link-protection] [edit protocols rsvp interface link-protection bypass] [edit protocols rsvp interface subscription] [edit protocols rsvp load-balance]
bypass	[edit protocols rsvp interface link-protection]
class-of-service	[edit protocols rsvp interface link-protection] [edit protocols rsvp interface link-protection bypass]
cleanup-timer	[edit protocols rsvp preemption soft-preemption]
ct0	[edit protocols rsvp interface link-protection bandwidth] [edit protocols rsvp interface link-protection bypass bandwidth] [edit protocols rsvp interface subscription]
ct1	[edit protocols rsvp interface link-protection bandwidth] [edit protocols rsvp interface link-protection bypass bandwidth] [edit protocols rsvp interface subscription]
ct2	[edit protocols rsvp interface link-protection bandwidth] [edit protocols rsvp interface link-protection bypass bandwidth] [edit protocols rsvp interface subscription]



Table 39: Unsupported [edit protocols rsvp] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
ct3	[edit protocols rsvp interface link-protection bandwidth] [edit protocols rsvp interface link-protection bypass bandwidth] [edit protocols rsvp interface subscription]
description	[edit protocols rsvp interface link-protection bypass]
devices	[edit protocols rsvp tunnel-services]
disable	[edit protocols rsvp graceful-restart] [edit protocols rsvp interface link-protection] [edit protocols rsvp peer-interface] [edit protocols rsvp preemption]
exclude	[edit protocols rsvp interface link-protection admin-group] [edit protocols rsvp interface link-protection bypass admin-group]
exclude-srlg	[edit protocols rsvp interface link-protection bypass] [edit protocols rsvp interface link-protection]
fast-reroute	[edit protocols rsvp]
graceful-deletion-timeout	[edit protocols rsvp]
graceful-restart	[edit protocols rsvp ]
hello-interval	[edit protocols rsvp peer-interface]
helper-disable	[edit protocols rsvp graceful-restart]
hop-limit	[edit protocols rsvp interface link-protection] [edit protocols rsvp interface link-protection bypass]
include-all	[edit protocols rsvp interface link-protection admin-group] [edit protocols rsvp interface link-protection bypass admin-group]
include-any	[edit protocols rsvp interface link-protection admin-group] [edit protocols rsvp interface link-protection bypass admin-group]
link-protection	[edit protocols rsvp interface]
load-balance	[edit protocols rsvp]
loose	[edit protocols rsvp interface link-protection bypass path] [edit protocols rsvp interface link-protection path]

Table 39: Unsupported [edit protocols rsvp] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
max-bypasses	[edit protocols rsvp interface link-protection]
maximum-helper-recovery-time	[edit protocols rsvp graceful-restart]
maximum-helper-restart-time	[edit protocols rsvp graceful-restart]
no-aggregate	[edit protocols rsvp peer-interface]
no-cspf	[edit protocols rsvp interface link-protection] [edit protocols rsvp interface link-protection bypass]
no-node-id-subobject	[edit protocols rsvp]
no-node-protection	[edit protocols rsvp interface link-protection]
no-reliable	[edit protocols rsvp peer-interface]
normal	[edit protocols rsvp preemption]
optimize-timer	[edit protocols rsvp fast-reroute] [edit protocols rsvp interface link-protection]
path	[edit protocols rsvp interface link-protection] [edit protocols rsvp interface link-protection bypass]
peer-interface	[edit protocols rsvp]
preemption	[edit protocols rsvp]
priority	[edit protocols rsvp interface link-protection] [edit protocols rsvp interface link-protection bypass]
reliable	[edit protocols rsvp peer-interface ]
soft-preemption	[edit protocols rsvp preemption]
strict	[edit protocols rsvp interface link-protection bypass path] [edit protocols rsvp interface link-protection path]
subscription	[edit protocols rsvp interface] [edit protocols rsvp interface link-protection]
to	[edit protocols rsvp interface link-protection bypass]
tunnel-services	[edit protocols rsvp]
update-threshold	[edit protocols rsvp interface]

- Related Documentation**
- *RSVP Feature Guide for Routing Devices*
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit protocols sflow]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols sflow\] Hierarchy Level on page 459](#)
- [Unsupported Statements in the \[edit sflow\] Hierarchy Level on page 460](#)

### [Supported Statements in the \[edit protocols sflow\] Hierarchy Level](#)

The following hierarchy shows the **[edit protocols sflow]** configuration statements supported on EX Series switches:

```
sflow {
  agent-id;
  collector {
    ip-address;
    udp-port port-number;
  }
  interfaces interface-name {
    polling-interval seconds;
    sample-rate {
      egress number;
      ingress number;
    }
  }
  polling-interval seconds;
  sample-rate {
    egress number;
    ingress number;
  }
  source-ip;
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
    no-world-readable>;
  flag (all | client-server | configuration | interface | rtsock);
}
```

### Unsupported Statements in the [edit sflow] Hierarchy Level

---

All statements in the **[edit protocols sflow]** hierarchy level that are displayed in the command-line interface (CLI) on the EX Series switch are supported on the switch and operate as documented.

- Related Documentation**
- [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\) on page 4049](#)
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols stp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols stp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols stp\] Hierarchy Level on page 460](#)
- [Unsupported Statements in the \[edit protocols stp\] Hierarchy Level on page 461](#)

### Supported Statements in the [edit protocols stp] Hierarchy Level

---

The following hierarchy shows the **[edit protocols stp]** configuration statements supported on EX Series switches:

```
protocols {
  stp {
    bpdv-block-on-edge;
    bridge-priority priority;
    disable;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      arp-on-stp;
      bpdv-timeout-action {
        block;
        log;
      }
    }
    cost cost;
    disable;
    edge;
    mode mode;
    no-root-port;
    priority priority;
```

```

}
max-age seconds;
traceoptions {
  file filename <files number > <size size> <no-stamp | world-readable |
    no-world-readable>;
  flag flag;
}
}

```

### Unsupported Statements in the [edit protocols stp] Hierarchy Level

All statements in the **[edit protocols stp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4956](#)
- [Configuring STP \(CLI Procedure\)](#)
- [Understanding STP for EX Series Switches on page 4910](#)
- [show spanning-tree bridge on page 5035](#)
- [show spanning-tree interface on page 5045](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

### [edit protocols uplink-failure-detection] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols uplink-failure-detection]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols uplink-failure-detection\] Hierarchy Level on page 462](#)
- [Unsupported Statements in the \[edit protocols uplink-failure-detection\] Hierarchy Level on page 462](#)

### Supported Statements in the [edit protocols uplink-failure-detection] Hierarchy Level

---

The following hierarchy shows the [edit protocols uplink-failure-detection] configuration statements supported on EX Series switches:

```
protocols {
  uplink-failure-detection {
    action {
      log;
    }
    group {
      group-name {
        link-to-monitor {
          interface-name;
        }
        link-to-disable {
          interface-name;
        }
      }
    }
    traceoptions {
      file filename <files number> <no-stamp> <replace> <size size> <world-readable |
        no-world-readable>;
      flag (all | dcd | groups | interface | parse );
    }
  }
}
```

### Unsupported Statements in the [edit protocols uplink-failure-detection] Hierarchy Level

---

All statements in the [edit protocols uplink-failure-detection] hierarchy level that are displayed in the command-line interface (CLI) on the EX Series switch are supported on the switch and operate as documented.

#### Related Documentation

- [Configuring Interfaces for Uplink Failure Detection \(CLI Procedure\) on page 4064](#)
- [Understanding Uplink Failure Detection on page 3977](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [edit protocols vrrp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols vrrp] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.

- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols vrrp\] Hierarchy Level on page 463](#)
- [Unsupported Statements in the \[edit protocols vrrp\] Hierarchy Level on page 463](#)

### [Supported Statements in the \[edit protocols vrrp\] Hierarchy Level](#)

The following hierarchy shows the **[edit protocols vrrp]** configuration statements supported on EX Series switches:

```
protocols {
  vrrp {
    failover-delay milliseconds;
    startup-silent-period seconds;
    traceoptions {
      file <filename> <files number> <match regular-expression> <microsecond-stamp>
        <size maximum-file-size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
  }
}
```

### [Unsupported Statements in the \[edit protocols vrrp\] Hierarchy Level](#)

All statements in the **[edit protocols vrrp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [Ethernet Interfaces Feature Guide for Routing Devices](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## [\[edit protocols vstp\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit protocols vstp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols vstp\] Hierarchy Level on page 464](#)
- [Unsupported Statements in the \[edit protocols vstp\] Hierarchy Level on page 464](#)

---

### Supported Statements in the [edit protocols vstp] Hierarchy Level

The following hierarchy shows the **[edit protocols vstp]** configuration statements supported on EX Series switches:

```
protocols {
  vstp {
    bpdu-block-on-edge;
    disable;
    force-version stp;
    vlan (all | vlan-id | vlan-name) {
      bridge-priority priority;
      forward-delay seconds;
      hello-time seconds;
      interface (all | interface-name) {
        arp-on-stp;
        bpdu-timeout-action {
          block;
          log;
        }
        cost cost;
        disable;
        edge;
        mode mode;
        no-root-port;
        priority priority;
      }
      max-age seconds;
      traceoptions {
        file filename <files number > <size size> <no-stamp | no-world-readable |
          world-readable>;
        flag flag;
      }
    }
  }
}
```

---

### Unsupported Statements in the [edit protocols vstp] Hierarchy Level

All statements in the **[edit protocols vstp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [Configuring VLAN Spanning-Tree Protocol](#)
- [Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912](#)
- [show spanning-tree bridge on page 5035](#)
- [show spanning-tree interface on page 5045](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)



## [edit redundant-power-system] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit redundant-power-system]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit redundant-power-system\] Hierarchy Level on page 465](#)
- [Unsupported Statements in the \[edit redundant-power-system\] Hierarchy Level on page 465](#)

### Supported Statements in the [edit redundant-power-system] Hierarchy Level

The following hierarchy shows the **[edit redundant-power-system]** configuration statements supported on EX Series switches:

```
redundant-power-system {
  member member-number {
    priority (0|1|2|3|4|5|6);
  }
}
```

### Unsupported Statements in the [edit redundant-power-system] Hierarchy Level

All statements in the **[edit redundant-power-system]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- *EX Series Redundant Power System (RPS) Documentation*
- *Understanding How Power Priority Is Determined and Set for Switches Connected to the EX Series Redundant Power System*
- *EX Series Redundant Power System Hardware Overview*

## [edit routing-options] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit routing-options]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.

- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit routing-options\] Hierarchy Level on page 466](#)
- [Unsupported Statements in the \[edit routing-options\] Hierarchy Level on page 482](#)

### Supported Statements in the [edit routing-options] Hierarchy Level

The following hierarchy shows the **[edit routing-options]** configuration statements supported on EX Series switches:

```
routing-options {
  access {
    route ip-prefix < / prefix-length > {
      metric route-cost;
      next-hop next-hop;
      preference route-distance;
      qualified-next-hop address {
        bfd-liveness-detection {
          authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
          }
          detection-time {
            threshold milliseconds;
          }
          holddown-interval milliseconds;
          local address address;
          minimum-interval milliseconds;
          minimum-receive-interval milliseconds;
          minimum-receive-ttl milliseconds;
          multiplier number;
          neighbor address;
          no-adaptation;
          transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
          }
          version (1 | automatic);
        }
        interface interface-name;
        mac-address mac-address;
        metric metric;
        preference preference-value;
      }
      tag route-tag;
    }
  } # end of [edit routing-options access]
```

```

access-internal {
  route ip-prefix </prefix-length> {
    next-hop [ addresses ];
    qualified-next-hop address {
      bfd-liveness-detection {
        authentication {
          algorithm algorithm-name;
          key-chain key-chain-name;
          loose-check;
        }
        detection-time {
          threshold milliseconds;
        }
        holddown-interval milliseconds;
        local address address;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-receive-ttl milliseconds;
        multiplier number;
        neighbor address;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        version (1 | automatic);
      }
      interface interface-name;
      mac-address mac-address;
      metric metric;
      preference preference-value;
    }
  }
} # end of [edit routing-options access-internal]
admin-groups-extended group-name {
  group-value group-identifier;
}
admin-groups-extended-range {
  maximum maximum-number;
  minimum minimum-number;
}
aggregate {
  defaults {
    (active | passive);
  }
  as-path {
    aggregator as-number address;
    atomic-aggregate;
    origin (egp | igp | incomplete);
    path path-identifier;
  }
  brief;
  color metric <type metric-type>;
  color2 metric <type metric-type>;
  community [ community-id no-advertise no-export no-export-subconfed ];
  discard;
  full;
}

```

```

metric metric <type metric-type>;
metric2 metric <type metric-type>;
metric3 metric <type metric-type>;
metric4 metric <type metric-type>;
preference preference-value <type metric-type>;
preference2 preference-value <type metric-type>;
tag metric <type metric-type>;
tag2 metric <type metric-type>;
}
route {
  (active | passive);
  as-path {
    aggregator as-number address;
    atomic-aggregate;
    origin (egp | igp | incomplete);
    path path-identifier;
  }
  brief;
  color metric <type metric-type>;
  color2 metric <type metric-type>;
  community [ community-id no-advertise no-export no-export-subconfed ];
  discard;
  full;
  metric metric <type metric-type>;
  metric2 metric <type metric-type>;
  metric3 metric <type metric-type>;
  metric4 metric <type metric-type>;
  policy [policy-names];
  preference preference-value <type metric-type>;
  preference2 preference-value <type metric-type>;
  tag metric <type metric-type>;
  tag2 metric <type metric-type>;
} # end of [edit routing-options aggregate]
}
bgp-orf-cisco-mode;
bmp {
  memory-limit bytes;
  station-address (ip-address | name);
  station-port-number port-number;
  statistics-timeout seconds;
}
confederation as-number members [ as-numbers ];
dynamic-tunnels {
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag (all | kernel | task | tunnel);
  }
}
fate-sharing {
  group group-name {
    cost value;
    from {
      address <to address>;
    }
  }
}

```

```

}
flow {
  firewall-install-disable {
    route;
    term-order;
    validation;
  }
  route name {
    match {
      destination address;
      destination-port [ afs bgp biff bootpc bootps cmd cvspserver dhcp domain eklogin
        ekshell exec finger ftp ftp-data http https ident imap kerberos-sec klogin kpasswd
        krb-prop krbupdate kshell ldap ldp login mobileip-agent mobilip-mn msdp
        netbios-dgm netbios-ns netbios-ssn nfsd nntp ntalk ntp pop3 pptp printer radacct
        radius rip rkinit smtp snmp snmptrap snpp socks ssh sunrpc syslog tacacs
        tacacs-ds talk telnet tftp timed who xdmcp ];
      dscp [ code-points ];
      fragment [ don't-fragment first-fragment is-fragment last-fragment
        not-a-fragment ];
      icmp-code [ communication-prohibited-by-filtering destination-host-prohibited
        destination-host-unknown fragmentation-needed host-precedence-violation
        host-unreachable host-unreachable-for-tos ip-header-bad network-unreachable
        network-unreachable-for-tos port-unreachable precedence-cutoff-in-effect
        protocol-unreachable redirect-for-host redirect-for-network
        redirect-for-tos-and-host redirect-for-tos-and-net required-option-missing
        source-host-isolated source-route-failed ttl-eq-zero-during-reassembly
        ttl-eq-zero-during-transit ];
      icmp-type [ echo-reply echo-request info-reply info-request mask-reply
        mask-request parameter-problem redirect router-advertisement router-solicit
        source-quench time-exceeded timestamp timestamp-reply unreachable ];
      packet-length [ values ];
      port [ ... same values as for the preceding destination-port statement ... ];
      protocol [ ah esp gre icmp igmp ipip ospf pim rsvp sctp tcp udp ];
      source address;
      source-port [ ... same values as for the preceding destination-port statement ... ];
      tcp-flags [ ack fin push rst syn urgent ];
    }
    then {
      (accept | discard);
      community community-name;
      next-term;
      rate-limit value;
      routing-instance routing-instance-name;
      sample;
    }
  }
}
term-order;
validation {
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
} # end of [edit routing-options flow]
forwarding-table {

```

```
export [ policy-names ];
indexed-next-hop;
(indirect-next-hop | no-indirect-next-hop);
(indirect-next-hop-change-acknowledgements |
  no-indirect-next-hop-change-acknowledgements);
krt-nexthop-ack-timeout;
unicast-reverse-path;
}
generate {
  defaults {
    (active | passive);
    as-path {
      aggregator as-number address;
      atomic-aggregate;
      origin (egp | igp | incomplete);
      path path-identifier;
    }
    brief;
    color metric <type metric-type>;
    color2 metric <type metric-type>;
    community [ community-id no-advertise no-export no-export-subconfed ];
    discard;
    full;
    metric metric <type metric-type>;
    metric2 metric <type metric-type>;
    metric3 metric <type metric-type>;
    metric4 metric <type metric-type>;
    preference preference-value <type metric-type>;
    preference2 preference-value <type metric-type>;
    tag metric <type metric-type>;
    tag2 metric <type metric-type>;
  }
  route {
    (active | passive);
    as-path {
      aggregator as-number address;
      atomic-aggregate;
      origin (egp | igp | incomplete);
      path path-identifier;
    }
    brief;
    color metric <type metric-type>;
    color2 metric <type metric-type>;
    community [ community-id no-advertise no-export no-export-subconfed ];
    discard;
    full;
    metric metric <type metric-type>;
    metric2 metric <type metric-type>;
    metric3 metric <type metric-type>;
    metric4 metric <type metric-type>;
    preference preference-value <type metric-type>;
    policy [ policy-names ];
    preference2 preference-value <type metric-type>;
    tag metric <type metric-type>;
    tag2 metric <type metric-type>;
  }
}
```

```

} # end of [edit routing-options generate]
graceful-restart {
  disable;
  restart-duration seconds;
}
host-fast-reroute {
  global-arp-prefix-limit;
  global-supplementary-blackout-timer;
}
instance-export [ policy-names ];
instance-import [ policy-names ];
interface-routes {
  family (inet | inet6) {
    export {
      lan;
      point-to-point;
    }
    import [ policy-names ];
  }
  rib-group {
    inet group-name;
    inet6 group-name;
  }
}
logical-system-mux {
  traceoptions {
    file {
      <file name>;
      files;
      no-world-readable;
      size;
      world-readable;
    }
    flag {
      all;
      debug;
      general;
      normal;
      parse;
      policy;
      route;
      state;
      task;
      timer;
    }
  }
}
martians {
  ip-prefix</prefix-length> (exact | longer | orlonger |
  prefix-length-range /minimum-prefix-length-/maximum-prefix-length |
  through ip-prefix</prefix-length> | upto /prefix-length) <allow>;
}
maximum-paths path-limit <log-only | threshold value> <log-interval seconds>;
maximum-prefixes prefix-limit <log-only | threshold value> <log-interval seconds>;
med-igp-update-interval minutes;
multicast {

```

```
asm-override-ssm;
backup-pe-group group-name {
    backups [ addresses ];
    local-address address;
}
flow-map flow-map-name {
    bandwidth <bps> <adaptive>;
    forwarding-cache {
        timeout (never <non-discard-entry-only> | minutes);
    }
    policy [ policy-names ];
    redundant-sources [ addresses ];
}
forwarding-cache {
    family {
        inet {
            threshold {
                log-warning;
                reuse;
                suppress;
            }
        }
        inet6 {
            threshold {
                log-warning;
                reuse;
                suppress;
            }
        }
    }
    threshold {
        log-warning;
        reuse;
        suppress;
    }
    timeout;
}
interface interface-name {
    maximum-bandwidth bps;
}
pim-to-igmp-proxy {
    upstream-interface [ interface-names ];
}
pim-to-mld-proxy {
    upstream-interface [ interface-names ];
}
rpf-check-policy [ policy-names ];
scope name {
    interface interface;
    prefix prefix;
}
scope-policy [ policy-names ];
ssm-groups value;
ssm-map ssm-map-name {
    policy [ policy-names ];
    source [ addresses ];
}
```



```

}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <disable>;
}
} # end of [edit routing-options multicast]
nonstop-routing;
no-bfd-triggered-local-repair;
options {
    mark seconds;
    syslog {
        level level;
        upto level;
    }
}
ppm {
    no-delegate-processing;
}
resolution {
    rib routing-table-name {
        import [ policy-names ];
        resolution-ribs [ routing-table-names ];
    }
    tracefilter [ filter-policy-names ];
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
rib routing-table-name {
    access {
        route ip-prefix </prefix-length> {
            metric route-cost;
            next-hop next-hop;
            preference route-distance;
            qualified-next-hop address {
                bfd-liveness-detection {
                    authentication {
                        algorithm algorithm-name;
                        key-chain key-chain-name;
                        loose-check;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                }
                holddown-interval milliseconds;
                local address address;
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                minimum-receive-ttl milliseconds;
                multiplier number;
                neighbor address;
                no-adaptation;
                transmit-interval {

```

```
        minimum-interval milliseconds;  
        threshold milliseconds;  
    }  
    version (1 | automatic);  
}  
interface interface-name;  
mac-address mac-address;  
metric metric;  
preference preference-value;  
}  
tag route-tag;  
}  
} # end of [edit routing-options rib access]  
access-internal {  
    route ip-prefix</prefix-length> {  
        next-hop [ addresses ];  
        qualified-next-hop address {  
            bfd-liveness-detection {  
                authentication {  
                    algorithm algorithm-name;  
                    key-chain key-chain-name;  
                    loose-check;  
                }  
                detection-time {  
                    threshold milliseconds;  
                }  
                holddown-interval milliseconds;  
                local address address;  
                minimum-interval milliseconds;  
                minimum-receive-interval milliseconds;  
                minimum-receive-ttl milliseconds;  
                multiplier number;  
                neighbor address;  
                no-adaptation;  
                transmit-interval {  
                    minimum-interval milliseconds;  
                    threshold milliseconds;  
                }  
                version (1 | automatic);  
            }  
        }  
        interface interface-name;  
        mac-address mac-address;  
        metric metric;  
        preference preference-value;  
    }  
    tag route-tag;  
}  
}  
} # end of [edit routing-options rib access-internal]  
aggregate {  
    defaults {  
        (active | passive);  
        as-path {  
            aggregator as-number address;  
            atomic-aggregate;  
            origin (egp | igp | incomplete);  
            path path-identifier;
```

```

}
brief;
color metric <type metric-type>;
color2 metric <type metric-type>;
community [ community-id no-advertise no-export no-export-subconfed ];
discard;
full;
metric metric <type metric-type>;
metric2 metric <type metric-type>;
metric3 metric <type metric-type>;
metric4 metric <type metric-type>;
preference preference-value <type metric-type>;
preference2 preference-value <type metric-type>;
tag metric <type metric-type>;
tag2 metric <type metric-type>;
}
route {
  (active | passive);
  as-path {
    aggregator as-number address;
    atomic-aggregate;
    origin (egp | igp | incomplete);
    path path-identifier;
  }
  brief;
  color metric <type metric-type>;
  color2 metric <type metric-type>;
  community [ community-id no-advertise no-export no-export-subconfed ];
  discard;
  full;
  metric metric <type metric-type>;
  metric2 metric <type metric-type>;
  metric3 metric <type metric-type>;
  metric4 metric <type metric-type>;
  policy [policy-names];
  preference preference-value <type metric-type>;
  preference2 preference-value <type metric-type>;
  tag metric <type metric-type>;
  tag2 metric <type metric-type>;
}
} # end of [edit routing-options rib aggregate]
generate {
  defaults {
    (active | passive);
    as-path {
      aggregator as-number address;
      atomic-aggregate;
      origin (egp | igp | incomplete);
      path path-identifier;
    }
    brief;
    color metric <type metric-type>;
    color2 metric <type metric-type>;
    community [ community-id no-advertise no-export no-export-subconfed ];
    discard;
    full;

```

```

metric metric <type metric-type>;
metric2 metric <type metric-type>;
metric3 metric <type metric-type>;
metric4 metric <type metric-type>;
policy value;
preference preference-value <type metric-type>;
preference2 preference-value <type metric-type>;
tag metric <type metric-type>;
tag2 metric <type metric-type>;
}
route {
  (active | passive);
  as-path {
    aggregator as-number address;
    atomic-aggregate;
    origin (egp | igp | incomplete);
    path path-identifier;
  }
  brief;
  color metric <type metric-type>;
  color2 metric <type metric-type>;
  community [ community-id no-advertise no-export no-export-subconfed ];
  discard;
  full;
  metric metric <type metric-type>;
  metric2 metric <type metric-type>;
  metric3 metric <type metric-type>;
  metric4 metric <type metric-type>;
  policy [policy-names];
  preference preference-value <type metric-type>;
  preference2 preference-value <type metric-type>;
  tag metric <type metric-type>;
  tag2 metric <type metric-type>;
}
} # end of [edit routing-options rib generate]
martians {
  ip-prefix</prefix-length> (exact | longer | orlonger |
    prefix-length-range /minimum-prefix-length-/maximum-prefix-length |
    through ip-prefix</prefix-length> | upto /prefix-length) <allow>;
}
maximum-paths path-limit <log-only | threshold value> <log-interval seconds>;
maximum-prefixes prefix-limit <log-only | threshold value> <log-interval seconds>;
static {
  defaults {
    (active | passive);
    as-path {
      aggregator as-number address;
      atomic-aggregate;
      origin (egp | igp | incomplete);
      path path-identifier;
    }
    color metric <type metric-type>;
    color2 metric <type metric-type>;
    community [ community-id no-advertise no-export no-export-subconfed ];
    (install | no-install);
    metric metric <type metric-type>;

```

```

metric2 metric <type metric-type>;
metric3 metric <type metric-type>;
metric4 metric <type metric-type>;
preference preference-value <type metric-type>;
preference2 preference-value <type metric-type>;
(readvertise | no-readvertise);
(resolve | no-resolve);
(retain | no-retain);
tag metric <type metric-type>;
tag2 metric <type metric-type>;
}
rib-group group-name;
route {
  (active | passive);
  as-path {
    aggregator as-number address;
    atomic-aggregate;
    origin (egp | igp | incomplete);
    path path-identifier;
  }
  backup-pe-group backup-pe-group;
  bfd-liveness-detection {
    authentication {
      algorithm algorithm-name;
      key-chain key-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    holddown-interval milliseconds;
    local address address;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-receive-ttl milliseconds;
    multiplier number;
    neighbor address;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (1 | automatic);
  }
  color metric <type metric-type>;
  color2 metric <type metric-type>;
  community [ community-id no-advertise no-export no-export-subconfed ];
  discard;
  (install | no-install);
  lsp-next-hop next-hop-address {
    metric metric;
    preference preference-value;
  }
  metric metric <type metric-type>;
  metric2 metric <type metric-type>;
  metric3 metric <type metric-type>;

```

```
metric4 metric <type metric-type>;
next-hop next-hop-address;
p2mp-lsp-next-hop next-hop-address {
    metric metric;
    preference preference-value;
}
preference preference-value <type metric-type>;
preference2 preference-value <type metric-type>;
qualified-next-hop address {
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        holddown-interval milliseconds;
        local address address;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-receive-ttl milliseconds;
        multiplier number;
        neighbor address;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (1 | automatic);
    }
    interface interface-name;
    mac-address mac-address;
    metric metric;
    preference preference-value;
}
(readvertise | no-readvertise);
receive;
reject;
(resolve | no-resolve);
(retain | no-retain);
static-lsp-nexthop next-hop-address {
    metric metric;
    preference preference-value;
}
tag metric <type metric-type>;
tag2 metric <type metric-type>;
}
} # end of [edit routing-options rib static]
} # end of [edit routing-options rib]
rib-groups {
    group-name {
        export-rib table-name;
        import-policy [ policy-names ];
        import-rib [ table-names ];
    }
}
```

```

    }
}
route-distinguisher-id address;
route-record;
router-id address;
source-routing {
    ip;
    ipv6;
}
srlg group-name {
    srlg-cost number;
    srlg-value number;
}
static {
    defaults {
        (active | passive);
        as-path {
            aggregator as-number address;
            atomic-aggregate;
            origin (egp | igp | incomplete);
            path path-identifier;
        }
        color metric <type metric-type>;
        color2 metric <type metric-type>;
        community [ community-id no-advertise no-export no-export-subconfed ];
        (install | no-install);
        metric metric <type metric-type>;
        metric2 metric <type metric-type>;
        metric3 metric <type metric-type>;
        metric4 metric <type metric-type>;
        preference preference-value <type metric-type>;
        preference2 preference-value <type metric-type>;
        (readvertise | no-readvertise);
        (resolve | no-resolve);
        (retain | no-retain);
        tag metric <type metric-type>;
        tag2 metric <type metric-type>;
    }
}
rib-group group-name;
route {
    (active | passive);
    as-path {
        aggregator as-number address;
        atomic-aggregate;
        origin (egp | igp | incomplete);
        path path-identifier;
    }
}
backup-pe-group backup-pe-group;
bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
}

```

```
    }
    holddown-interval milliseconds;
    local address address;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-receive-ttl milliseconds;
    multiplier number;
    neighbor address;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
color metric <type metric-type>;
color2 metric <type metric-type>;
community [ community-id no-advertise no-export no-export-subconfed ];
discard;
(install | no-install);
lsp-next-hop next-hop-address {
    metric metric;
    preference preference-value;
}
metric metric <type metric-type>;
metric2 metric <type metric-type>;
metric3 metric <type metric-type>;
metric4 metric <type metric-type>;
next-hop next-hop-address;
p2mp-lsp-next-hop next-hop-address {
    metric metric;
    preference preference-value;
}
preference preference-value <type metric-type>;
preference2 preference-value <type metric-type>;
qualified-next-hop address {
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        holddown-interval milliseconds;
        local address address;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-receive-ttl milliseconds;
        multiplier number;
        neighbor address;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
    }
}
```



```

    }
    version (1 | automatic);
  }
  interface interface-name;
  mac-address mac-address;
  metric metric;
  preference preference-value;
}
(readvertise | no-readvertise);
receive;
reject;
(resolve | no-resolve);
(retain | no-retain);
static-lsp-next-hop next-hop-address {
  metric metric;
  preference preference-value;
}
tag metric <type metric-type>;
tag2 metric <type metric-type>;
}
} # end of [edit routing-options static]
topologies {
  family {
    inet {
      topology name;
    }
    inet6 {
      topology name;
    }
  }
}
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
  no-world-readable>;
  flag (all | condition-manager | config-internal | general | graceful-restart | hfrf-fsm |
  hfrf-route | normal | nsr-synchronization | parse | policy | regex-parse | route | session
  | state | task | timer);
}
validation {
  group group-name {
    max-sessions;
    session address {
      hold-time;
      local-address;
      port;
      preference;
      record-lifetime;
      refresh-time;
      traceoptions;
    }
  }
}
notification-rib value;
static {
  record destination {
    maximum-length prefix-length;
  }
}

```

```

    }
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag (all | error | general | keepalive | nsr-synchronization | packets | policy | state |
        task | timer | update);
    }
}
}

```

### Unsupported Statements in the [edit routing-options] Hierarchy Level

All statements in the [edit routing-options] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 40: Unsupported [edit routing-options] Configuration Statements on EX Series Switches**

Statement	Hierarchy
<b>NOTE:</b> Variables, such as <i>interface-name</i> , are not shown in the statements or hierarchies.	
forwarding-cache	[edit routing-options multicast]
scope	[edit routing-options multicast]
ssm-groups	[edit routing-options multicast]
threshold	[edit routing-options multicast forwarding-cache]
timeout	[edit routing-options multicast forwarding-cache]

**Related Documentation**

- *Junos OS Configuration Statements and Commands*

### [edit security] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit security] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit security\] Hierarchy Level on page 483](#)
- [Unsupported Statements in the \[edit security\] Hierarchy Level on page 485](#)

### Supported Statements in the [edit security] Hierarchy Level

The following hierarchy shows the **[edit security]** configuration statements supported on EX Series switches:

```
security {
  alarms {
    potential-violation {
      authentication failures;
      cryptographic-self-test;
      key-generation-self-test;
      non-cryptographic-self-test;
      policy number per (minute | second);
      replay-attacks {
        threshold value;
      }
      security-log-percent-full;
    }
  }
  certificates {
    cache-size bytes;
    cache-timeout-negative seconds;
    certification-authority ca-profile-name {
      ca-name certificate-authority-name;
      crl filename;
      encoding (binary | pem);
      enrollment-url url;
      file certificate-filename;
      ldap-url url-name;
    }
    enrollment-retry number;
    local certificate-name {
      certificate-key-string;
      load-key-file URL-or-path;
    }
    maximum-certificates number;
    path-length bytes;
  }
  ipsec {
    security-association sa-name {
      description text-description;
      manual {
        direction (bidirectional | inbound | outbound) {
        }
        mode (transport | tunnel);
      }
    }
  }
  log {
    cache {
      exclude name {
```

```
        destination-address;
        destination-port;
        event-id;
        failure;
        interface-name;
        policy-name;
        process;
        source-address;
        source-port;
        success;
        username;
    }
    limit number;
}
}
macsec {
    connectivity-association connectivity-association-name {
        exclude-protocol protocol-name;
        include-sci;
        mka {
            must-secure;
            key-server-priority priority-number;
            transmit-interval interval;
        }
        no-encryption;
        offset (0|30|50);
        pre-shared-key {
            cak hexadecimal-number;
            ckn hexadecimal-number;
        }
        replay-protect {
            replay-window-size number-of-packets;
        }
        secure-channel secure-channel-name {
            direction (inbound | outbound);
            encryption;
            id {
                mac-address mac-address;
                port-id port-id-number;
            }
            offset (0|30|50);
            security-association security-association-number {
                key key-string;
            }
        }
        security-mode security-mode;
    }
    interfaces interface-name {
        connectivity-association connectivity-association-name;
    }
}
pki {
    auto-re-enrollment {
        certificate-id certificate-id {
            ca-profile-name profile-name;
            challenge-password password;
        }
    }
}
```

```
        re-enroll-trigger-time-percentage percentage;  
        re-generate-keypair;  
    }  
}  
traceoptions {  
    file <filename> <files number> <match regular-expression> <size maximum-file-size>  
    <world-readable | no-world-readable>;  
    flag flag;  
}  
}  
ssh-known-hosts {  
    fetch-from-server (hostname | address);  
    host (hostname | address) {  
        dsa-key key;  
        ecdsa-sha2-nistp256-key key;  
        ecdsa-sha2-nistp384-key key;  
        ecdsa-sha2-nistp521-key key;  
        rsa-key key;  
        rsa1-key key;  
    }  
    load-key-file filename;  
}  
traceoptions {  
    file <filename> <files number> <match regular-expression> <size maximum-file-size>  
    <world-readable | no-world-readable>;  
    flag flag;  
    level level;  
    no-remote-trace;  
    rate-limit rate;  
}  
}
```

Unsupported Statements in the [edit security] Hierarchy Level

All statements in the [edit security] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 41: Unsupported [edit security] Configuration Statements on EX Series Switches

Statement	Hierarchy
NOTE: Variables, such as <i>filename</i> , are not shown in the statements or hierarchies.	
audible	[edit security alarms]
continuous	[edit security alarms audible]

[edit services] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit services] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switches, see *EX Series Switch Software Features Overview*.
- [Supported Statements in the \[edit services\] Hierarchy Level on page 486](#)
- [Unsupported Statements in the \[edit services\] Hierarchy Level on page 488](#)

---

### Supported Statements in the [edit services] Hierarchy Level

The following hierarchy shows the **[edit services]** configuration statements supported on EX Series switches:

```
services {
  captive-portal {
    authentication-profile-name authentication-profile-name;
    custom-options {
      banner-message string;
      footer-bgcolor color;
      footer-message string;
      footer-text-color color;
      form-header-bgcolor color;
      form-header-message string;
      footer-header-text-color color;
      form-reset-label label-name;
      form-submit-label label-name;
      header-bgcolor color;
      header-logo filename;
      header-message string;
      header-text-color color0;
      post-authentication-url url;
    }
  }
  interface (all | interface-name) {
    quiet-period seconds;
    retries number-of-retries;
    server-timeout seconds;
    session-expiry seconds;
    supplicant (multiple | single | single-secure);
  }
  secure-authentication (http | https);
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
}
rpm {
  bgp {
    data-fill data;
    data-size size;
```

```

destination-port port;
history-size size;
moving-average-size number-of-samples;
probe-count count;
probe-interval seconds;
probe-type type;
routing-instances {
    routing-instance-name;
}
test-interval seconds;
}
probe owner {
    test test-name {
        data-fill data;
        data-size size;
        destination-port port;
        dscp-code-point dscp-bits;
        hardware-timestamp;
        history-size size;
        moving-average-size number;
        one-way-hardware-timestamp;
        probe-count count;
        probe-interval seconds;
        probe-type type;
        routing-instance instance-name;
        source-address address;
        target (address address | url url);
        test-interval interval;
        thresholds {
            egress-time microseconds;
            ingress-time microseconds;
            jitter-egress microseconds;
            jitter-ingress microseconds;
            jitter-rtt microseconds;
            rtt microseconds;
            std-dev-egress microseconds;
            std-dev-ingress microseconds;
            std-dev-rtt microseconds;
            successive-loss count;
            total-loss count;
        }
        traps [ trap-names ];
    }
}
probe-limit number;
probe-server {
    tcp {
        port port-number;
    }
    udp {
        port port-number;
    }
}
}
unified-access-control {
    certificate-verification (optional | required | warning);
}

```

```

infranet-controllerhostname {
    address ip-address;
    interface interface-name;
    password password;
    port port-number;
}
interval seconds;
timeout seconds;
timeout-action (close | no-change);
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <disable>;
}
}

```

### Unsupported Statements in the [edit services] Hierarchy Level

All statements in the [edit services] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 42: Unsupported [edit services] Configuration Statements on EX Series Switches**

Statement	Hierarchy
<b>NOTE:</b> Variables, such as <i>interface-name</i> , are not shown in the statements or hierarchies.	
ca-profile	[edit services unified-access-control infranet-controller]
interface	[edit services interface-pools] [edit services service-device-pools pool]
pool	[edit services interface-pools] [edit services service-device-pools]
server-certificate-subject	[edit services unified-access-control infranet-controller]
service-device-pools	[edit services]
service-interface-pools	[edit services]

**Related Documentation**

- *Junos OS Configuration Statements and Commands*

### [edit snmp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit snmp] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.



- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*

This topic lists:

- [Supported Statements in the \[edit snmp\] Hierarchy Level on page 489](#)
- [Unsupported Statements in the \[edit snmp\] Hierarchy Level on page 493](#)

### Supported Statements in the [edit snmp] Hierarchy Level

The following hierarchy shows the **[edit snmp]** configuration statements supported on EX Series switches:

```
snmp {
  client-list list-name {
    address {
      restrict;
    }
  }
  community community-name {
    authorization (read-only | read-write);
    client-list-name list-name;
    clients {
      address <restrict>;
    }
    routing-instance instance-name;
    routing-instance instance-name {
      client-list-name list-name;
      clients {
        address <restrict>;
      }
    }
    view view-name;
  }
  contact contact-information;
  description description;
  engine-id {
    (local engine-id | use-default-ip-address | use-mac-address);
  }
  filter-duplicates;
  filter-interfaces {
    interfaces
    all-internal-interfaces;
    interface 1;
    interface 2;
  }
  health-monitor {
    falling-threshold percentage;
    idp {
      falling-threshold;
      interval seconds;
    }
  }
}
```

```
        rising-threshold;
    }
    interval seconds;
    rising-threshold percentage;
}
interface [ interface-names ];
location location;
name system-name;
nonvolatile {
    commit-delay seconds;
}
rmon {
    alarm index {
        description description;
        falling-event-index index;
        falling-threshold integer;
        falling-threshold-interval seconds;
        interval seconds;
        request-type (get-next-request | get-request | walk-request);
        rising-event-index index;
        rising-threshold integer;
        sample-type (absolute-value | delta-value);
        startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
        syslog-subtag text-string;
        variable oid-variable;
    }
    event index {
        community community-name;
        description description;
        type (log | log-and-trap | none | snmptrap);
    }
    history index {
        bucket-size number;
        interface interface-name;
        interval seconds;
        owner owner-name;
    }
}
routing-instance-access {
    access-list {
        routing-instance-name <restrict>;
    }
}
traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
trap-group group-name {
    categories {
        authentication;
        chassis;
        configuration;
        link;
        otn-alarms {
```

```

        alarm-name;
    }
    remote-operations;
    rmon-alarm;
    routing;
    services;
    sonet-alarms {
        alarm-name;
    }
    startup;
    vrrp-events;
}
destination-port port-number;
routing-instance instance-name;
routing-instance instance-name;
targets {
    address;
}
version (all | v1 | v2);
}
trap-options {
    agent-address outgoing-interface;
    enterprise-oid;
    routing-instance instance-name;
    routing-instance instance-name {
        source-address (address | lo0);
    }
    source-address address;
}
v3 {
    ... the v3 subhierarchy appears after the main [edit snmp] hierarchy level ...
}
view view-name {
    oid object-identifier <exclude | include>;
}
}

snmp {
    v3 {
        notify name {
            tag tag-name;
            type (inform | trap);
        }
        notify-filter profile-name {
            oid oid <exclude | include>;
        }
        snmp-community community-index {
            community-name community-name;
            context context-name;
            security-name security-name;
            tag tag-name;
        }
        target-address target-address-name {
            address address;
            address-mask address-mask;
            routing-instance routing-instance-name;
        }
    }
}

```

```
port port-number;
retry-count number;
routing-instance routing-instance-name;
tag-list tag-list;
target-parameters parameter-name;
timeout seconds;
}
target-parameters parameter-name {
  notify-filter profile-name;
  parameters {
    message-processing-model (v1 | v2c | v3);
    security-level (authentication | none | privacy);
    security-model (usm | v1 | v2c);
    security-name security-name;
  }
}
usm {
  local-engine {
    user username {
      authentication-md5 {
        authentication-key password;
        authentication-password password;
      }
      authentication-none;
      authentication-sha {
        authentication-key password;
        authentication-password password;
      }
      privacy-3des {
        privacy-password password;
      }
      privacy-aes128 {
        privacy-password password;
      }
      privacy-des {
        privacy-password password;
      }
      privacy-none;
    }
  }
  remote-engine engine-id {
    user username {
      authentication-md5 {
        authentication-key password;
        authentication-password password;
      }
      authentication-none;
      authentication-sha {
        authentication-key
        authentication-password password;
      }
      privacy-3des {
        privacy-password password;
      }
      privacy-aes128 {
        privacy-password password;
      }
    }
  }
}
```

```
    }
    privacy-des {
        privacy-password password;
    }
    privacy-none;
}
}
}
vacm {
    access {
        group group-name {
            context-prefix prefix {
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        context-match (exact | prefix);
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
    default-context-prefix prefix {
        security-model (any | usm | v1 | v2c) {
            security-level (authentication | none | privacy) {
                context-match (exact | prefix);
                notify-view view-name;
                read-view view-name;
                write-view view-name;
            }
        }
    }
}
}
}
}
security-to-group {
    security-model (usm | v1 | v2c) {
        security-name security-name {
            group group-name;
        }
    }
}
}
}
}
```

Unsupported Statements in the [edit snmp] Hierarchy Level

All statements in the [edit snmp] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 43: Unsupported [edit snmp] Configuration Statements on EX Series Switches

Statement	Hierarchy
-----------	-----------

NOTE: Variables, such as *community-name*, are not shown in the statements or hierarchies.

Table 43: Unsupported [edit snmp] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
logical-system	[edit snmp community] [edit snmp trap-group]

Table 43: Unsupported [edit snmp] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
	[edit snmp trap-options] [edit snmp v3 target-address]
logical-systems-trap-filter	[edit snmp]

- Related Documentation**
- [Configuring SNMP \(J-Web Procedure\) on page 4040](#)
  - [Network Management Administration Guide for Routing Devices](#)

## [edit system] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit system] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit system\] Hierarchy Level on page 495](#)
- [Unsupported Statements in the \[edit system\] Hierarchy Level on page 507](#)

### Supported Statements in the [edit system] Hierarchy Level

The following hierarchy shows the [edit system] configuration statements supported on EX Series switches.

```

system {
  accounting {
    destination {
      radius {
        server {
          server-address {
            accounting-port port-number;
            port port-number;
            retry number;
            secret password;
            source-address address;
            timeout seconds;
          }
        }
      }
    }
  }
  tacplus {

```

```

server {
    server-address {
        port port-number;
        secret password;
        single-connection;
        timeout seconds;
    }
}
}
events
traceoptions {
    file;
    flag;
    no-remote-trace;
}
}
allow-v4-mapped-packets;
archival {
    configuration {
        archive-sites {
            ftp://<username>:<password>@<host>:<port>/<url-path>;
            scp://<username>:<password>@<host>:<port>/<url-path>;
        }
        transfer-interval interval;
        transfer-on-commit;
    }
}
arp {
    aging-timer minutes;
    gratuitous-arp-delay;
    gratuitous-arp-on-ifup;
    interfaces interface-name {
        aging-timer minutes;
    }
    passive-learning;
    purging;
}
authentication-order [ authentication-methods ];
autoinstallation {
    configuration-servers {
        server-url <password password>;
    }
    interfaces {
        interface-name {
            bootp;
            rarp;
        }
    }
}
}
backup-router address <destination [ destination-addresses ]>;
commit {
    synchronize (and-quit | force);
}
(compress-configuration-files | no-compress-configuration-files);
default-address-selection;

```



```

domain-name domain-name;
domain-search [ domain-list ];
extensions {
  providers {
    provider-id {
      license-type license deployment-scope [ deployments ];
    }
  }
  resource-limits {
    package package-name {
      resources {
        cpu {
          priority number;
          time seconds;
        }
        file {
          core-size bytes;
          open number;
          size bytes;
        }
        memory {
          data-size bytes;
          locked-in bytes;
          resident-set-size bytes;
          socket-buffers bytes;
          stack-size bytes;
        }
      }
    }
  }
  process process-ui-name {
    resources {
      cpu {
        priority number;
        time seconds;
      }
      file {
        core-size bytes;
        open number;
        size bytes;
      }
      memory {
        data-size bytes;
        locked-in bytes;
        resident-set-size bytes;
        socket-buffers bytes;
        stack-size bytes;
      }
    }
  }
}
internet-options {
  (gre-path-mtu-discovery | no-gre-path-mtu-discovery);
  icmpv4-rate-limit bucket-size number packet-rate rate;
  icmpv6-rate-limit bucket-size number packet-rate rate;
  (ipip-path-mtu-discovery | no-ipip-path-mtu-discovery);
}

```

```
ipv6-duplicate-addr-detection-transmits;
(ipv6-path-mtu-discovery | noipv6-path-mtu-discovery);
ipv6-path-mtu-discovery-timeout;
ipv6-reject-zero-hop-limit | no-ipv6-reject-zero-hop-limit;
no-tcp-reset;
no-tcp-rfc1323-paws;
no-tcp-rfc1323;
(path-mtu-discovery | no-path-mtu-discovery);
source-port upper-limit port-number;
(source-quench | no-source-quench);
tcp-drop-synfin-set;
}
kernel-replication;
license {
  autoupdate {
    url url{
      password password;
    }
  }
  renew {
    before-expiration days;
    interval hours;
  }
  traceoptions {
    file <filename> <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
location {
  altitude feet;
  building name;
  country-code code;
  floor number;
  hcoord horizontal-coordinate;
  lata service-area;
  latitude degrees;
  longitude degrees;
  npa-nxx number;
  postal-code postal-code;
  rack number;
  vcoord vertical-coordinate;
}
login {
  announcement "text";
  class class-name {
    access-end "hh:mm:ss>>";
    access-start "hh:mm:ss>>";
    allow-commands "regular-expression";
    allow-configuration-regexps "regular-expression";
    allowed-days [ sunday monday tuesday wednesday thursday friday saturday ];
    deny-commands "regular-expression";
    deny-configuration-regexps "regular-expression";
    idle-timeout minutes;
    login-alarms;
```

```

login-script script-name;
login-tip;
permissions [ permissions ];
security-role [ security-role ];
}
deny-sources {
    address;
}
message "text";
password {
    change-type (character-sets | set-transitions);
    format (des | md5 | sha1);
    maximum-length number;
    minimum-changes number;
    minimum-length number;
}
retry-options {
    backoff-factor number;
    backoff-threshold number;
    lockout-period number;
    maximum-time number;
    minimum-time number;
    tries-before-disconnect number;
}
user username {
    authentication {
        (encrypted-password "password" | plain-text-password);
        load-key-file filename;
        ssh-dsa "public-key" <from hostname>;
        ssh-rsa "public-key" <from hostname>;
    }
    class class-name;
    full-name "complete-name";
    uid uid-value;
}
}
max-configurations-on-flash number;
name-server {
    address;
}
nd-maxmcast-solicit;
nd-retransmit-timer;
no-multicast-echo;
no-neighbor-learn;
no-ping-record-route;
no-ping-time-stamp;
}
ntp {
    authentication-key key-number type md5 value password;
    boot-server address;
    broadcast <address> <key key-number> <ttl value> <version value>;
    broadcast-client;
    multicast-client <address>;
    peer address <key key-number> <prefer> <version value>;
    server address <key key-number> <prefer> <version value>;
    source-address source-address;
}

```

```
    trusted-key [ key-numbers ];
  }
  ports {
    auxiliary {
      disable;
      insecure;
      port-type (mini-usb | rj45);
      type (ansi | small-xterm | vt100 | xterm);
    }
    console {
      disable;
      insecure;
      log-out-on-disconnect;
      type (ansi | small-xterm | vt100 | xterm);
    }
  }
  radius-options {
    attributes {
      nas-ip-address address;
    }
    password-protocol mschap-v2;
  }
  radius-server {
    server-address {
      accounting-port port-number;
      port port-number;
      retry number;
      secret password;
      source-address source-address;
      timeout seconds;
    }
  }
  root-authentication {
    (encrypted-password "password" | plain-text-password);
    load-key-file filename;
    ssh-dsa "public-key" <from hostname>;
    ssh-rsa "public-key" <from hostname>;
  }
  (saved-core-context | no-saved-core-context);
  saved-core-files number;
  scripts {
    commit {
      allow-transients;
      direct-access;
      file filename.xml {
        checksum (md5 | sha-256 | sha2) hash;
        optional;
        refresh;
        refresh-from url;
        source url;
      }
      refresh;
      refresh-from url;
      traceoptions {
        file <filename> <files number> <size maximum-file-size> <world-readable |
          no-world-readable>;
```

```

        flag flag;
        no-remote-trace;
    }
}
load-scripts-from-flash;
op {
    file filename.xml {
        arguments {
            argument-name {
                description descriptive-text;
            }
        }
        checksum (md5 | sha-256 | sha2) hash;
        command filename-alias;
        description descriptive-text;
        refresh;
        refresh-from url;
        source url;
    }
    no-allow-url;
    refresh;
    refresh-from url;
    traceoptions {
        file <filename> <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
}
services {
    database-replication {
        traceoptions {
            file <filename> <files number> <match regular-expression>
                <size maximum-file-size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
}
dhcp {
    boot-file filename;
    boot-server (address | hostname);
    default-lease-time (seconds | infinite);
    domain-name domain-name;
    domain-search {
        domain-suffix;
    }
    maximum-lease-time (seconds | infinite);
    name-server {
        address;
    }
    next-server address;
    option option-index (array type-name [ type-values ] | byte 8-bit-value | flag (false |
        off | on | true) | integer signed-32-bit-value | ip-address address |
        short signed-16-bit-value | string text-string | unsigned-integer 32-bit-value |
        unsigned-short 16-bit-value);

```

```
pool ip-prefix/prefix-length {
  address-range low address high address;
  boot-file filename;
  boot-server (address | hostname);
  default-lease-time (seconds | infinite);
  domain-name domain-name;
  domain-search {
    domain-suffix;
  }
  exclude-address {
    ipv4-address;
  }
  maximum-lease-time (seconds | infinite);
  name-server {
    address;
  }
  next-server address;
  option option-index (array type-name type-values ] | byte 8-bit-value | flag (false |
    off | on | true) | integer signed-32-bit-value | ip-address address |
    short signed-16-bit-value | string text-string | unsigned-integer 32-bit-value |
    unsigned-short 16-bit-value);
  propagate-settings interface-name;
  router {
    address;
  }
  server-identifier identifier;
  sip-server {
    address {
      address;
    }
    name {
      name;
    }
  }
  wins-server {
    address;
  }
}

router {
  address;
}

server-identifier identifier;
sip-server {
  address {
    address;
  }
  name {
    name;
  }
}

static-binding mac-address {
  boot-file filename;
  boot-server (address | hostname);
  client-identifier (ascii ascii-text | hexadecimal hexadecimal-value);
  domain-name domain-name;
  domain-search {
```

```

    domain-suffix;
}
fixed-address {
    ipv4-address;
}
host-name hostname;
name-server {
    address;
}
next-server address;
option option-index (array type-name type-values ] | byte 8-bit-value | flag (false |
    off | on | true) | integer signed-32-bit-value | ip-address address |
    short signed-16-bit-value | string text-string | unsigned-integer 32-bit-value |
    unsigned-short 16-bit-value);
router {
    address;
}
server-identifier identifier;
sip-server {
    address {
        address;
    }
    name {
        name;
    }
}
wins-server {
    address;
}
}
traceoptions {
    file <filename> <files number> <match regular-expression>
        <size maximum-file-size> <world-readable | no-world-readable>;
    flag flag;
    level severity;
    no-remote-trace;
}
wins-server {
    address;
}
}
dhcp-local-server {
    group group-name {
        interface interface-name {
            exclude;
            overrides {
                client-discover-match <option60-and-option82>;
                interface-client-limit number;
                no-arp;
                process-inform {
                    pool pool-name;
                }
            }
            trace;
            upto upto-interface-name;
        }
    }
}

```

```
overrides {
  client-discover-match <option60-and-option82>;
  interface-client-limit number;
  OBSOLETE - no-arp;
  process-inform {
    pool pool-name;
  }
}
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}
}
overrides {
  client-discover-match <option60-and-option82>;
  interface-client-limit number;
  OBSOLETE - no-arp;
  process-inform {
    pool pool-name;
  }
}
pool-match-order {
  external-authority;
  ip-address-first;
  option-82;
}
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}
}
finger {
  connection-limit limit;
  rate-limit limit;
}
ftp {
  connection-limit limit;
  rate-limit limit;
}
netconf {
  ssh {
    connection-limit limit;
    port number;
    rate-limit limit;
  }
}
```



```

outbound-ssh {
  client client-id {
    address {
      port port-number;
      retry number;
      timeout seconds;
    }
    device-id device-id;
    keep-alive {
      retry number;
      timeout seconds;
    }
    reconnect-strategy (in-order | sticky);
    secret secret;
    services netconf;
  }
  traceoptions {
    file <filename> <files number> <match regular-expression>
      <size maximum-file-size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
service-deployment {
  local-certificate certificate-name;
  servers {
    server-address {
      port port-number;
      security-options {
        (ssl3 | tls);
      }
      user username;
    }
  }
  source-address source-address;
  traceoptions {
    file <filename> <files number> <match regular-expression>
      <size maximum-file-size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
ssh {
  ciphers;
  connection-limit limit;
  hostkey-algorithm {
    ssh-dss | no-ssh-dss;
    ssh-ecdsa | no-ssh-ecdsa;
    ssh-rsa | no-ssh-rsh;
  }
  key-exchange;
  macs;
  protocol-version [ v1 v2 ];
  rate-limit limit;
  root-login (allow | deny | deny-password);
}

```

```
subscriber-management{
  gres-route-flush-delay;
  maintain-subscriber {
    interface-delete;
  }
  traceoptions {
    file filename <files number> <match regular-expression> <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
telnet {
  connection-limit limit;
  rate-limit limit;
}
web-management {
  control {
    max-threads number;
  }
  http {
    interface [ interface-names ];
    port port-number;
  }
  https {
    interface [ interface-names ];
    (local-certificate certificate-name | pki-local-certificate certificate-name |
      system-generated-certificate);
    port port-number;
  }
  management-url url;
  session {
    idle-timeout minutes;
    session-limit number;
  }
}
xnm-clear-text {
  connection-limit limit;
  rate-limit limit;
}
xnm-ssl {
  connection-limit limit;
  local-certificate certificate-name;
  rate-limit limit;
}
}
static-host-mapping {
  hostname {
    alias [ aliases ];
    inet [ addresses ];
    inet6 [ addresses ];
    sysid system-identifier;
  }
}
syslog {
  allow-duplicates;
```

```

archive <files number> <size size> <world-readable | no-world-readable>;
console {
    facility severity;
}
file filename {
    allow-duplicates;
    facility severity;
    archive <archive-sites {ftp-url <password password>}> <files number> <size size>
        <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
        no-world-readable>;
    explicit-priority;
    match "regular-expression";
    structured-data {
        brief;
    }
}
host (hostname | other-routing-engine) {
    facility severity;
    explicit-priority;
    facility-override facility;
    log-prefix string;
    match "regular-expression";
}
log-rotate-frequency;
time-format (year | millisecond | year millisecond);
user (username | *) {
    facility severity;
    explicit-priority;
    match "regular-expression";
}
}
tacplus-options {
    (exclude-cmd-attribute | no-cmd-attribute-value);
    service-name service-name;
}
tacplus-server {
    server-address {
        port port-number;
        secret password;
        single-connection;
        source-address source-address;
        timeout seconds;
    }
}
time-zone (GMT | GMT+hour-offset | GMT-hour-offset | zone-name);
tracing destination-override syslog host address;
use-imported-time-zones;
}

```

### Unsupported Statements in the [edit system] Hierarchy Level

All statements in the [edit system] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 44: Unsupported [edit system] Configuration Statements on EX Series Switches

Statement	Hierarchy
NOTE: Variables, such as <i>interface-name</i> , are not shown in the statements or hierarchies.	
mirror-flash-on-disk	[edit system]
processes	[edit system]

**Related Documentation**

- [Configuration File Management on EX Series Switches](#)
- [EX Series Switches Hardware and CLI Terminology Mapping on page 315](#)

## [edit virtual-chassis] Configuration Statement Hierarchy

This topic lists supported and unsupported configuration statements in the **[edit virtual-chassis]** hierarchy level on EX Series and QFX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms.

For detailed information about feature support on specific EX Series or QFX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit virtual-chassis\] Hierarchy Level on page 508](#)
- [Unsupported Statements in the \[edit virtual-chassis\] Hierarchy Level on page 509](#)

### Supported Statements in the [edit virtual-chassis] Hierarchy Level

The following hierarchy shows the **[edit virtual-chassis]** configuration statements supported on EX Series or QFX Series switches:

```
virtual-chassis {
  aliases {
    serial-number serial-number {
      alias-name alias-name;
    }
  }
  auto-provisioned;
  auto-sw-update {
    (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
    package-name package-name;
  }
  fast-failover (ge | vcp disable | xe);
}
```

```

graceful-restart {
  disable;
}
id id;
mac-persistence-timer [minutes | disable];;
member member-id {
  location location;
  mastership-priority number;
  no-management-vlan;
  role (line-card | routing-engine);
  serial-number;
}
no-split-detection;
preprovisioned;
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regex>;
  flag flag ;
}
vc-port {
  lag-hash (packet-based | source-port-based);
}
vcp-no-hold-time;
}

```

### Unsupported Statements in the [edit virtual-chassis] Hierarchy Level

All statements in the [edit virtual-chassis] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4300 Virtual Chassis \(CLI Procedure\) on page 5097](#)
- [Configuring an EX2200 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX3300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches \(CLI Procedure\)](#)
- [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#)

### [edit vlans] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit vlans] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.

- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit vlans\] Hierarchy Level on page 510](#)
- [Unsupported Statements in the \[edit vlans\] Hierarchy Level on page 511](#)

---

### Supported Statements in the [edit vlans] Hierarchy Level

The following hierarchy shows the **[edit vlans]** configuration statements supported on one or more of the EX Series switches:

```
vlans {  
  vlan-name {  
    description text-description;  
    dot1q-tunneling {  
      customer-vlans (id | native | range);  
      layer2-protocol-tunneling all | protocol-name {  
        drop-threshold number;  
        shutdown-threshold number;  
      }  
    }  
  }  
  filter {  
    input filter-name  
    output filter-name;  
  }  
  interface interface-name {  
    egress;  
    ingress;  
    mapping (native (push | swap) | policy | tag (push | swap));  
    pvlan-trunk;  
  }  
  isolation-id id-number;  
  l3-interface vlan.logical-interface-number;  
  l3-interface-ingress-counting layer-3-interface-name;  
  mac-limit limit action action;  
  mac-table-aging-time seconds;  
  no-local-switching;  
  no-mac-learning;  
  primary-vlan vlan-name;  
  vlan-id number;  
  vlan-prune;  
  vlan-range vlan-id-low-vlan-id-high;  
}
```

Unsupported Statements in the [edit vlans] Hierarchy Level

All statements in the [edit vlans] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 45: Unsupported [edit vlans] Configuration Statements on EX Series Switches

Statement	Hierarchy Level
-----------	-----------------

NOTE: Variables, such as *filename*, are not shown in the statements or hierarchies.

udid	[edit vlans dot1q-tunneling layer2-protocol-tunneling]
------	--

- Related Documentation
- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
  - *Example: Connecting an Access Switch to a Distribution Switch*
  - *Example: Setting Up Q-in-Q Tunneling on EX Series Switches*
  - *Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches*
  - *Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches*
  - *Example: Configuring a Private VLAN Spanning Multiple EX Series Switches*
  - *Creating a Private VLAN on a Single EX Series Switch (CLI Procedure)*





## CHAPTER 8

# Administration

- [Operational Commands on page 513](#)

### Operational Commands

---

- [set cli directory](#)
- [set cli idle-timeout](#)
- [set cli prompt](#)
- [set cli restart-on-upgrade](#)
- [set cli screen-length](#)
- [set cli screen-width](#)
- [set cli timestamp](#)
- [start shell](#)

## set cli directory

---

<b>Syntax</b>	set cli directory <i>directory</i>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Set the current working directory.
<b>Options</b>	<i>directory</i> —Pathname of the working directory.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>CLI User Interface Overview</i></li><li>• <i>show cli directory</i></li></ul>
<b>List of Sample Output</b>	<a href="#">set cli directory on page 514</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### set cli directory

```
user@host> set cli directory /var/home/regress
Current directory: /var/home/regress
```

## set cli idle-timeout

---

<b>Syntax</b>	set cli idle-timeout < <i>minutes</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Set the maximum time that an individual session can be idle before the user is logged off the router or switch.
<b>Options</b>	<i>minutes</i> —(Optional) Maximum idle time. The range of values, in minutes, is 0 through 100,000. If you do not issue this command, and the user's login class does not specify this value, the user is never forced off the system after extended idle times. Setting the value to 0 disables the timeout.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>CLI User Interface Overview</i></li> <li>• <i>show cli</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">set cli idle-timeout on page 515</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### set cli idle-timeout

```
user@host> set cli idle-timeout 60
Idle timeout set to 60 minutes
```

## set cli prompt

---

<b>Syntax</b>	set cli prompt <i>string</i>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Set the prompt so that it is displayed within the CLI.
<b>Options</b>	<i>string</i> —CLI prompt string. To include spaces in the prompt, enclose the string in quotation marks. By default, the string is <i>username@hostname</i> .
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>CLI User Interface Overview</i></li><li>• <i>show cli</i></li></ul>
<b>List of Sample Output</b>	<a href="#">set cli prompt on page 516</a>
<b>Output Fields</b>	When you enter this command, the new CLI prompt is displayed.

## Sample Output

### set cli prompt

```
user@host> set cli prompt lab1-router>
lab1-router>
```

## set cli restart-on-upgrade

---

<b>Syntax</b>	set cli restart-on-upgrade string (off   on)
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	For an individual session, set the CLI to prompt you to restart the router or switch after upgrading the software.
<b>Options</b>	<b>off</b> —Disables the prompt. <b>on</b> —Enables the prompt.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>CLI User Interface Overview</i></li> <li>• <i>show cli</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">set cli restart-on-upgrade on page 517</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### set cli restart-on-upgrade

```
user@host> set cli restart-on-upgrade on
Enabling restart-on-upgrade
```

## set cli screen-length

---

<b>Syntax</b>	<code>set cli screen-length <i>length</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Set terminal screen length.
<b>Options</b>	<i>length</i> —Number of lines of text that the terminal screen displays (0 through 10,000). The default is 24.
<b>Additional Information</b>	The point at which the ---( <b>more</b> )--- prompt appears on the screen is a function of this setting and the settings for the <code>set cli screen-width</code> and <code>set cli terminal</code> commands.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>CLI User Interface Overview</i></li><li>• <a href="#">set cli screen-width on page 519</a></li><li>• <i>set cli terminal</i></li><li>• <i>show cli</i></li></ul>
<b>List of Sample Output</b>	<a href="#">set cli screen-length on page 518</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### set cli screen-length

```
user@host> set cli screen-length 75
Screen length set to 75
```

## set cli screen-width

---

<b>Syntax</b>	<code>set cli screen-width <i>width</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Set the terminal screen width.
<b>Options</b>	<i>width</i> —Number of characters (0 through 1024) in a line. The default is 80.
<b>Additional Information</b>	The point at which the ---( <b>more</b> )--- prompt appears on the screen is a function of this setting and the settings for the <code>set cli screen-length</code> and <code>set cli terminal</code> commands.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>CLI User Interface Overview</i></li> <li>• <a href="#">set cli screen-length on page 518</a></li> <li>• <i>set cli terminal</i></li> <li>• <i>show cli</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">set cli screen-width on page 519</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.


### Sample Output

#### set cli screen-width

```
user@host> set cli screen-width
Screen width set to 132
```

## set cli timestamp

---

<b>Syntax</b>	set cli timestamp (format <i>timestamp-format</i>   disable)
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Set a timestamp for CLI output.
<b>Options</b>	<p><b>format <i>timestamp-format</i></b>—Set the date and time format for the timestamp. The timestamp format you specify can include the following placeholders in any order:</p> <ul style="list-style-type: none"><li>• <b>%m</b>—Two-digit month</li><li>• <b>%d</b>—Two-digit date</li><li>• <b>%T</b>—Six-digit hour, minute, and seconds</li></ul> <p><b>disable</b>—Remove the timestamp from the CLI.</p>
<div> <b>NOTE:</b> A timestamp is displayed by default when no command output is generated.</div>	
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>CLI User Interface Overview</i></li><li>• <i>show cli</i></li></ul>
<b>List of Sample Output</b>	<a href="#">set cli timestamp on page 520</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.


## Sample Output

### set cli timestamp

```
user@host> set cli timestamp format '%m-%d-%T'
'04-21-17:39:13'
CLI timestamp set to: '%m-%d-%T'
```



## start shell

<b>Syntax</b>	start shell (csh   sh) <user <i>username</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Exit from the CLI environment and create a UNIX-level shell. To return to the CLI, type <b>exit</b> from the shell.
<div>  <b>NOTE:</b> <ul style="list-style-type: none"> <li>To issue this command, the user must have the required login access privileges configured by including the <b>permissions</b> statement at the [edit system login class <i>class-name</i>] hierarchy level.</li> <li>UNIX wheel group membership or permissions are no longer required to issue this command.</li> </ul> </div>	
<b>Options</b>	<b>csh</b> —Create a UNIX C shell.  <b>sh</b> —Create a UNIX Bourne shell.  <b>user <i>username</i></b> —(Optional) Start the shell as another user.
<b>Additional Information</b>	When you are in the shell, the shell prompt has the following format:  <i>username@hostname</i> %  An example of the prompt is:  root@host%
<b>Required Privilege Level</b>	shell and maintenance
<b>List of Sample Output</b>	<a href="#">start shell csh on page 521</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### start shell csh

```

user@host> start shell csh
%

exit
%
```

```
username@hostname% start shell sh
%

exit
user@host>
```

## PART 4

# Access and User Management

- [Overview on page 525](#)
- [Configuration on page 529](#)
- [Administration on page 579](#)
- [Troubleshooting Procedures on page 601](#)



## CHAPTER 9

# Overview

- [Software Overview on page 525](#)

## Software Overview

---

- [Understanding Software Infrastructure and Processes on page 525](#)

### Understanding Software Infrastructure and Processes

Each switch runs the Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches on its general-purpose processors. Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the chassis.

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

With the J-Web interface and the command-line interface (CLI) to the Junos OS, you configure switching features and routing protocols and set the properties of network interfaces on your switch. After activating a software configuration, use either the J-Web or CLI user interface to monitor the switch, manage operations, and diagnose protocol and network connectivity problems.

- [Routing Engine and Packet Forwarding Engine on page 525](#)
- [Junos OS Processes on page 526](#)

### Routing Engine and Packet Forwarding Engine

---

A switch has two primary software processing components:

- **Packet Forwarding Engine**—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.
- **Routing Engine**—Provides three main functions:
  - Creates the packet forwarding switch fabric for the switch, providing route lookup, filtering, and switching on incoming data packets, then directing outbound packets to the appropriate interface for transmission to the network

- Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
- Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

### Junos OS Processes

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

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

Table 6 on page 30 describes the primary Junos OS processes.

**Table 46: Junos OS Processes**

Process	Name	Description
Chassis process	chassisd	<p>Detects hardware on the system that is used to configure network interfaces.</p> <p>Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered.</p> <p>Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.</p>
Ethernet switching process	eswd	<p>Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree protocol and access port security. The process is also responsible for managing Ethernet switching interfaces, VLANs, and VLAN interfaces.</p> <p>Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.</p>
Forwarding process	pfem	<p>Defines how routing protocols operate on the switch. The overall performance of the switch is largely determined by the effectiveness of the forwarding process.</p>
Interface process	dcd	<p>Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.</p>
Management process	mgd	<p>Provides communication between the other processes and an interface to the configuration database.</p> <p>Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured.</p> <p>Interacts with the other processes when commands are issued through one of the user interfaces on the switch.</p> <p>If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.</p>
Routing protocol process	rpd	<p>Defines how routing protocols such as RIP, OSPF, and BGP operate on the device, including selecting routes and maintaining forwarding tables.</p>

- Related Documentation**
- For more information about processes, see *Junos OS Network Operations Guide*
  - For more information about basic system parameters, supported protocols, and software processes, see *Junos OS System Basics Configuration Guide*





## CHAPTER 10

# Configuration

- [Configuration Tasks on page 529](#)
- [Configuration Statements on page 533](#)

### Configuration Tasks

---

- [Configuring Management Access for the EX Series Switch \(J-Web Procedure\) on page 529](#)
- [Generating SSL Certificates to Be Used for Secure Web Access on page 532](#)
- [Configuring MS-CHAPv2 to Provide Password-Change Support \(CLI Procedure\) on page 533](#)

### Configuring Management Access for the EX Series Switch (J-Web Procedure)

You can manage an EX Series switch remotely through the J-Web interface. To communicate with the switch, the J-Web interface uses HTTP. HTTP enables easy Web access, but uses no encryption. The data that is transmitted between the Web browser and the switch by means of HTTP is vulnerable to interception and attack. To enable secure Web access the switch supports HTTPS. You can enable HTTP or HTTPS access on specific interfaces and ports as needed.

Navigate to the Secure Access Configuration page by selecting **Configure > System Properties > Management Access**. On this page, you can enable HTTP and HTTPS access on interfaces for managing the EX Series switch through the J-Web interface. You can also install SSL certificates and enable Junos XML management protocol over SSL with the Secure Access page.

1. Click **Edit** to modify the configuration. Enter information into the Management Access Configuration page as described in [Table 47 on page 530](#).
2. To verify that Web access is enabled correctly, connect to the switch using the appropriate method:
  - For HTTP access—In your Web browser, type **http://URL** or **http://IP address**.
  - For HTTPS access—In your Web browser, type **https://URL** or **https://IP address**.
  - For SSL Junos XML management protocol access—To use this option, you must have a Junos XML management protocol client such as Junos Scope. For information about how to log in to Junos Scope, see the *Junos Scope Software User Guide*.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

**Table 47: Secure Management Access Configuration Summary**

Field	Function	Your Action
<b>Management Access tab</b>		
Management Port IP/Management Port IPv6	<p>Specifies the management port IP address. The software supports both IPv4 ( displayed as IP) and IPv6 address.</p> <p><b>NOTE:</b> IPv6 is not supported on EX2200 and EX 4500 switches.</p>	<p>To specify an IPv4 address:</p> <ol style="list-style-type: none"> <li>1. Select the check box <b>IPv4 address</b>.</li> <li>2. Type an IP address—for example: <b>10.10.10.10</b>.</li> <li>3. Enter the subnet mask or address prefix. For example, 24 bits represents <b>255.255.255.0</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> <p>To specify an IPv6 address:</p> <ol style="list-style-type: none"> <li>1. Select the check box <b>IPv6 address</b>.</li> <li>2. Type an IP address—for example: <b>2001:ab8:85a3::8a2e:370:7334</b>.</li> <li>3. Enter the subnet mask or address prefix.</li> <li>4. Click <b>OK</b>.</li> </ol>
Default Gateway	Defines a default gateway through which to direct packets addressed to networks that are not explicitly listed in the bridge table constructed by the switch.	For IPv4 address type a 32-bit IP address, in dotted decimal notation. Type a 128-bit IP address for IPv6 address type.
Loopback address	Specifies the IP address of the loopback interface.	Type an IP address.
Subnet Mask	Specifies the subnet mask for the loopback interface.	Enter the subnet mask or address prefix.
<b>Services tab</b>		
Services	Specifies services to be enabled: telnet and SSH.	Select to enable the required services.
Enable Junos XML management protocol over Clear Text	Enables clear text access to the Junos XML management protocol XML scripting API.	To enable clear text access, select the <b>Enable Junos XML management protocol over Clear Text</b> check box.
Enable Junos XML protocol over SSL	Enables secure SSL access to the Junos XML management protocol XML scripting API.	To enable SSL access, select the <b>Enable Junos XML management protocol over SSL</b> check box.

Table 47: Secure Management Access Configuration Summary (*continued*)

Field	Function	Your Action
Junos XML management protocol Certificate	Specifies SSL certificates to be used for encryption.  This field is available only after you create at least one SSL certificate.	To enable an SSL certificate, select a certificate from the Junos XML management protocol SSL Certificate list—for example, <b>new</b> .
Enable HTTP	Enables HTTP access on interfaces.	To enable HTTP access, select the <b>Enable HTTP access</b> check box.  Select and clear interfaces by clicking the direction arrows:  <ul style="list-style-type: none"> <li>To enable HTTP access on an interface, add the interface to the HTTP Interfaces list. You can either select either all interfaces or specific interfaces.</li> </ul>
Enable HTTPS	Enables HTTPS access on interfaces.	To enable HTTPS access, select the <b>Enable HTTPS access</b> check box.  Select and deselect interfaces by clicking the direction arrows:  <ul style="list-style-type: none"> <li>To enable HTTPS access on an interface, add the interface to the HTTPS Interfaces list. You can either select either all interfaces or specific interfaces.</li> </ul> <p><b>NOTE:</b> Specify the certificate to be used for HTTPS access.</p>

**Certificates tab**

Certificates	Displays digital certificates required for SSL access to the switch.  Allows you to add and delete SSL certificates.	To add a certificate:  <ol style="list-style-type: none"> <li>Have a general SSL certificate available. See Generating SSL Certificates for more information.</li> <li>Click <b>Add</b>. The Add a Local Certificate page opens.</li> <li>Type a name in the Certificate Name box—for example, <b>new</b>.</li> <li>Open the certificate file and copy its contents.</li> <li>Paste the generated certificate and RSA private key in the Certificate box.</li> </ol> <p>To edit a certificate, select it and click <b>Edit</b>.</p> <p>To delete a certificate, select it and click <b>Delete</b>.</p>
--------------	--	--

**Related Documentation** • [Security Features for EX Series Switches Overview on page 4693](#)

- [Understanding J-Web User Interface Sessions](#)
- [Enabling HTTPS and XNM-SSL Services on Switches Using Self-Signed Certificates \(CLI Procedure\) on page 1417](#)

## Generating SSL Certificates to Be Used for Secure Web Access

You can set up secure Web access for an EX Series switch. To enable secure Web access, you must generate a digital Secure Sockets Layer (SSL) certificate and then enable HTTPS access on the switch.

To generate an SSL certificate:

1. Enter the following **openssl** command in your SSH command-line interface on a BSD or Linux system on which **openssl** is installed. The **openssl** command generates a self-signed SSL certificate in the privacy-enhanced mail (PEM) format. It writes the certificate and an unencrypted 1024-bit RSA private key to the specified file.

```
% openssl req -x509 -nodes -newkey rsa:1024 -keyout filename.pem -out filename.pem
```

where **filename** is the name of a file in which you want the SSL certificate to be written—for example, **my-certificate**.

2. When prompted, type the appropriate information in the identification form. For example, type **US** for the country name.
3. Display the contents of the file that you created.

```
cat my-certificate.pem
```

You can use the J-Web Configuration page to install the SSL certificate on the switch. To do this, copy the file containing the certificate from the BSD or Linux system to the switch. Then open the file, copy its contents, and paste them into the Certificate box on the J-Web Secure Access Configuration page.

You can also use the following CLI statement to install the SSL certificate on the switch:

```
[edit]  
user@switch# set security certificates local my-signed-cert load-key-file my-certificate.pem
```

### Related Documentation

- [Configuring Management Access for the EX Series Switch \(J-Web Procedure\) on page 529](#)
- [Security Features for EX Series Switches Overview on page 4693](#)

## Configuring MS-CHAPv2 to Provide Password-Change Support (CLI Procedure)

Junos OS for EX Series switches enables you to configure the Microsoft Corporation implementation of the Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) on the switch to provide password-change support. Configuring MS-CHAPv2 on the switch provides users accessing a switch the option of changing the password when the password expires, is reset, or is configured to be changed at next login.

See RFC 2433, *Microsoft PPP CHAP Extensions*, for information about MS-CHAP.

Before you configure MS-CHAPv2 to provide password-change support, ensure that you have:

- Configured RADIUS server authentication. Configure users on the authentication server and set the first-tried option in the authentication order to radius. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).

To configure MS-CHAPv2, specify the following:

```
[edit system radius-options]
user@switch# set password-protocol mschap-v2
```

You must have the required access permission on the switch in order to change your password.

### Related Documentation

- [Managing Users \(J-Web Procedure\) on page 579](#)
- For more about configuring user access, see the [Junos OS Access Privilege Configuration Guide](#).

## Configuration Statements

- [allow-commands on page 535](#)
- [allow-configuration on page 536](#)
- [announcement on page 537](#)
- [archive-sites on page 537](#)
- [authentication \(Login\) on page 538](#)
- [authentication-order on page 539](#)
- [change-type on page 540](#)
- [class \(Assigning a Class to an Individual User\) on page 540](#)
- [class \(Defining Login Classes\) on page 541](#)
- [class-usage-profile on page 542](#)
- [counters on page 543](#)
- [deny-commands on page 543](#)
- [deny-configuration on page 544](#)
- [destination-classes on page 545](#)

- [fields \(for Interface Profiles\) on page 546](#)
- [file \(Associating with a Profile\) on page 547](#)
- [file \(Configuring a Log File\) on page 548](#)
- [files on page 549](#)
- [filter-profile on page 550](#)
- [format on page 551](#)
- [full-name on page 552](#)
- [idle-timeout \(System-Login\) on page 552](#)
- [interface-profile on page 553](#)
- [interval \(Accounting Options\) on page 554](#)
- [login on page 555](#)
- [login-alarms on page 556](#)
- [login-tip on page 556](#)
- [maximum-length on page 557](#)
- [message on page 558](#)
- [mib-profile on page 559](#)
- [minimum-changes on page 560](#)
- [minimum-length on page 561](#)
- [object-names on page 562](#)
- [operation on page 562](#)
- [password \(Login\) on page 563](#)
- [permissions on page 564](#)
- [radius-options \(edit system\) on page 565](#)
- [retry-options on page 566](#)
- [root-authentication on page 567](#)
- [root-login on page 568](#)
- [routing-engine-profile on page 569](#)
- [size on page 570](#)
- [source-classes on page 570](#)
- [start-time \(Log File Transfer\) on page 571](#)
- [tacplus-options on page 572](#)
- [tacplus-server on page 573](#)
- [traceoptions \(Address-Assignment Pool\) on page 574](#)
- [transfer-interval on page 575](#)
- [uid on page 576](#)
- [user \(Access\) on page 577](#)

## allow-commands

---

<b>Syntax</b>	<code>allow-commands "regular-expression";</code>
<b>Hierarchy Level</b>	[edit system login <a href="#">class</a> <i>class-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the operational mode commands that members of a login class can use.
<b>Default</b>	If you omit this statement and the <b>deny-commands</b> statement, users can issue only those commands for which they have access privileges through the <b>permissions</b> statement.
<b>Options</b>	<b>regular-expression</b> —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Specifying Access Privileges for Junos OS Operational Mode Commands</i></li> <li>• <a href="#">deny-commands on page 543</a></li> <li>• <a href="#">user on page 577</a></li> </ul>

## allow-configuration

---

<b>Syntax</b>	<code>allow-configuration "regular-expression";</code>
<b>Hierarchy Level</b>	[edit system login <a href="#">class</a> <i>class-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Explicitly allow configuration access to the specified levels in the hierarchy even if the permissions set with the <b>permissions</b> statement do not grant such access by default.
<b>Default</b>	If you omit this statement and the <b>deny-configuration</b> statement, users can edit only those commands for which they have access privileges through the <b>permissions</b> statement.
<b>Options</b>	<b>regular-expression</b> —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Specifying Access Privileges Using allow/deny-configuration Statements</i></li><li>• <i>Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies</i></li><li>• <a href="#">deny-configuration on page 544</a></li><li>• <a href="#">user on page 577</a></li></ul>



## announcement

---

<b>Syntax</b>	<code>announcement text;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">login</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure a system login announcement. This announcement appears after a user logs in.
<b>Options</b>	<b>text</b> —Text of the announcement. If the text contains any spaces, enclose it in quotation marks.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Junos OS to Display a System Login Announcement</i></li> <li>• <a href="#">message on page 558</a></li> </ul>

## archive-sites

---

<b>Syntax</b>	<code>archive-sites {     site-name; }</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">file filename</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure an archive site. If more than one site name is configured, an ordered list of archive sites for the accounting-data log files is created. When a file is archived, the router or switch attempts to transfer the file to the first URL in the list, moving to the next site only if the transfer does not succeed. The log file is stored at the archive site with a filename of the format <b>router-name_log-filename_timestamp</b> .
<b>Options</b>	<b>site-name</b> —Any valid FTP URL to a destination. For information about specifying valid FTP URLs, see the <i>Getting Started Guide for Routing Devices</i> .
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Archive Sites</i></li> </ul>

## authentication (Login)

---

Syntax	<pre>authentication {   (encrypted-password "password"   plain-text-password);   load-key-file URL filename;   ssh-dsa "public-key";   ssh-ecdsa "public-key";   ssh-rsa "public-key"; }</pre>
Hierarchy Level	[edit system login <b>user</b> <i>username</i> ]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Authentication methods that a user can use to log in to the router or switch. You can assign multiple authentication methods to a single user.
Options	<p><b>encrypted-password "password"</b>—Message Digest 5 (MD5) or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.</p> <p>You cannot configure a blank password for <b>encrypted-password</b> using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.</p> <p><b>load-key-file URL filename</b>—Load previously-generated RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys from a named file at a specified URL location. The file contains one or more SSH keys.</p> <p><b>plain-text-password</b>—When using this option, the command-line interface (CLI) prompts you for the password and then encrypts it.</p> <p><b>ssh-dsa "public-key"</b>—SSH version 2 authentication. Specify the DSA public key. You can specify one or more public keys for each user.</p> <p><b>ssh-ecdsa "public-key"</b>—SSH version 2 authentication. Specify the ECDSA public key. You can specify one or more public keys for each user.</p> <p><b>ssh-rsa "public-key"</b>—SSH version 1 and SSH version 2 authentication. Specify the RSA public key. You can specify one or more public keys for each user.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS User Accounts</i></li><li>• <a href="#">root-authentication on page 567</a></li></ul>

## authentication-order

---

<b>Syntax</b>	<code>authentication-order [ <i>authentication-methods</i> ];</code>
<b>Hierarchy Level</b>	<code>[edit system]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the order in which the software tries different user authentication methods when attempting to authenticate a user. For each login attempt, the software tries the authentication methods in order, starting with the first one, until the password matches.
<b>Default</b>	If you do not include the <b>authentication-order</b> statement, users are verified based on their configured passwords.
<b>Options</b>	<p><b><i>authentication-methods</i></b>—One or more authentication methods, listed in the order in which they should be tried. The method can be one or more of the following:</p> <ul style="list-style-type: none"> <li>• <b>password</b>—Use the password configured for the user with the <b>authentication</b> statement at the <code>[edit system login user]</code> hierarchy level.</li> <li>• <b>radius</b>—Use RADIUS authentication services.</li> <li>• <b>tacplus</b>—Use TACACS+ authentication services.</li> </ul>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication</i></li> <li>• <a href="#">authentication on page 538</a></li> </ul>

## change-type

---

<b>Syntax</b>	<code>change-type (character-sets   set-transitions);</code>
<b>Hierarchy Level</b>	[edit system login <a href="#">password</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Set requirements for using character sets in plain-text passwords. When you combine this statement with the <b>minimum-changes</b> statement, you can check for the total number of character sets included in the password or for the total number of character-set changes in the password. Newly created passwords must meet these requirements.
<b>Options</b>	Specify one of the following: <ul style="list-style-type: none"><li>• <b>character-sets</b>—The number of character sets in the password. Valid character sets include uppercase letters, lowercase letters, numbers, punctuation, and other special characters.</li><li>• <b>set-transitions</b>—The number of transitions between character sets.</li></ul>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Special Requirements for Junos OS Plain-Text Passwords</i></li><li>• <a href="#">minimum-changes on page 560</a></li></ul>

## class (Assigning a Class to an Individual User)

---

<b>Syntax</b>	<code>class class-name;</code>
<b>Hierarchy Level</b>	[edit system login <a href="#">user</a> <i>username</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Assign a user to a login class. You must assign each user to a login class.
<b>Options</b>	<b>class-name</b> —One of the classes defined at the [edit system login class] hierarchy level.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS User Accounts</i></li></ul>

## class (Defining Login Classes)

<b>Syntax</b>	<pre> class <i>class-name</i> {   <b>allow-commands</b> "<i>regular-expression</i>";   ( <b>allow-configuration</b>   allow-configuration-regexps ) "<i>regular expression 1</i>" "<i>regular expression 2</i>";   configuration-breadcrumbs;   <b>deny-commands</b> "<i>regular-expression</i>";   ( <b>deny-configuration</b>   deny-configuration-regexps ) "<i>regular expression 1</i>" "<i>regular expression 2</i>";   <b>idle-timeout</b> <i>minutes</i>;   login-script <i>filename</i>;   <b>login-tip</b>;   <b>permissions</b> [ <i>permissions</i> ]; }</pre>
<b>Hierarchy Level</b>	[edit system <a href="#">login</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Define a login class.
<b>Options</b>	<p><b><i>class-name</i></b>—A name you choose for the login class.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Defining Junos OS Login Classes</li> <li><a href="#">user on page 577</a></li> </ul>

## class-usage-profile

---

Syntax	<pre>class-usage-profile <i>profile-name</i> {     <i>file filename</i>;     <i>interval minutes</i>;     source-classes {         <i>source-class-name</i>;     }     destination-classes {         <i>destination-class-name</i>;     } }</pre>
Hierarchy Level	[edit accounting-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Create a class usage profile, which is used to log class usage statistics to a file in the <code>/var/log</code> directory. The class usage profile logs class usage statistics for the configured source classes on every interface that has <b>destination-class-usage</b> configured.</p> <p>For information about configuring source classes, see the <a href="#">Junos OS Routing Protocols Library for Routing Devices</a>. For information about configuring source class usage, see <a href="#">Network Management Administration Guide for Routing Devices</a>.</p>
Options	<p><b>profile-name</b>—Name of the destination class profile.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring Class Usage Profiles</a></li></ul>

## counters

---

<b>Syntax</b>	<code>counters {     counter-name; }</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">filter-profile</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Names of counters for which filter profile statistics are collected. The packet and byte counts for the counters are logged to a file in the <code>/var/log</code> directory.
<b>Options</b>	<i>counter-name</i> —Name of the counter.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Counters</i></li> </ul>

## deny-commands

---

<b>Syntax</b>	<code>deny-commands "regular-expression";</code>
<b>Hierarchy Level</b>	[edit system login <a href="#">class</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the operational mode commands that the user is denied permission to issue even though the permissions set with the <b>permissions</b> statement would allow it.
<b>Default</b>	If you omit this statement and the <b>allow-commands</b> statement, users can issue only those commands for which they have access privileges through the <b>permissions</b> statement.
<b>Options</b>	<i>regular-expression</i> —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Specifying Access Privileges for Junos OS Operational Mode Commands</i></li> <li>• <a href="#">allow-commands on page 535</a></li> <li>• <a href="#">user on page 577</a></li> </ul>

## deny-configuration

---

<b>Syntax</b>	<code>deny-configuration "regular-expression";</code>
<b>Hierarchy Level</b>	[edit system login <a href="#">class</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Explicitly deny configuration access to the specified levels in the hierarchy even if the permissions set with the <b>permissions</b> statement grant such access by default. Note that the user cannot view a particular hierarchy if configuration access is denied for that hierarchy.
<b>Default</b>	If you omit this statement and the <b>allow-configuration</b> statement, users can edit those levels in the configuration hierarchy for which they have access privileges through the <b>permissions</b> statement.
<b>Options</b>	<b>regular-expression</b> —Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Specifying Access Privileges Using allow/deny-configuration Statements</i></li><li>• <a href="#">allow-configuration on page 536</a></li><li>• <a href="#">user on page 577</a></li></ul>



---

## destination-classes

---

<b>Syntax</b>	<code>destination-classes {     <i>destination-class-name</i>; }</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">class-usage-profile</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the destination classes for which statistics are collected.
<b>Options</b>	<b><i>destination-class-name</i></b> —Name of the destination class to include in the source class usage profile.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring a Class Usage Profile</i></li></ul>

## fields (for Interface Profiles)

---

<b>Syntax</b>	<pre>fields {     field-name; }</pre>
<b>Hierarchy Level</b>	[edit accounting-options <b>interface-profile</b> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Statistics to collect in an accounting-data log file for an interface.
<b>Options</b>	<p><i>field-name</i>—Name of the field:</p> <ul style="list-style-type: none"><li>• <b>input-bytes</b>—Input bytes</li><li>• <b>input-errors</b>—Generic input error packets</li><li>• <b>input-multicast</b>—Input packets arriving by multicast</li><li>• <b>input-packets</b>—Input packets</li><li>• <b>input-unicast</b>—Input unicast packets</li><li>• <b>output-bytes</b>—Output bytes</li><li>• <b>output-errors</b>—Generic output error packets</li><li>• <b>output-multicast</b>—Output packets sent by multicast</li><li>• <b>output-packets</b>—Output packets</li><li>• <b>output-unicast</b>—Output unicast packets</li></ul>
<b>Required Privilege Level</b>	<p><b>interface</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Interface Profile</i></li></ul>

## file (Associating with a Profile)

<b>Syntax</b>	<code>file <i>filename</i>;</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">class-usage-profile</a> <i>profile-name</i> ], [edit accounting-options <a href="#">filter-profile</a> <i>profile-name</i> ], [edit accounting-options <a href="#">interface-profile</a> <i>profile-name</i> ], [edit accounting-options <a href="#">mib-profile</a> <i>profile-name</i> ], [edit accounting-options <a href="#">routing-engine-profile</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. The [edit accounting-options <a href="#">mib-profile</a> <i>profile-name</i> ] hierarchy added in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series Switches.
<b>Description</b>	Specify the accounting log file associated with the profile.
<b>Options</b>	<b><i>filename</i></b> —Name of the log file. You must specify a filename already configured in the <b>file</b> statement at the [edit accounting-options] hierarchy level.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Interface Profile</a></li> <li>• <a href="#">Configuring the Filter Profile</a></li> <li>• <a href="#">Configuring the MIB Profile</a></li> <li>• <a href="#">Configuring the Routing Engine Profile</a></li> </ul>

## file (Configuring a Log File)

---

<b>Syntax</b>	<pre>file <i>filename</i> {     archive-sites {         <i>site-name</i>;     }     files <i>number</i>;     nonpersistent;     size <i>bytes</i>;     source-classes <i>time</i>;     transfer-interval <i>minutes</i>; }</pre>
<b>Hierarchy Level</b>	[edit accounting-options]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify a log file to be used for accounting data.
<b>Options</b>	<i>filename</i> —Name of the file in which to write accounting data.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>Configuring Accounting-Data Log Files</li></ul>

---

## files

---

<b>Syntax</b>	<code>files <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">file</a> <i>filename</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the maximum number of log files to be used for accounting data.
<b>Options</b>	<i>number</i> —The maximum number of files. When a log file (for example, <b>profilelog</b> ) reaches its maximum size, it is renamed <b>profilelog.0</b> , then <b>profilelog.1</b> , and so on, until the maximum number of log files is reached. Then the oldest log file is overwritten. The minimum value for <i>number</i> is 3 and the default value is 10.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Accounting-Data Log Files</i></li></ul>

## filter-profile

---

<b>Syntax</b>	<pre>filter-profile <i>profile-name</i> {     counters {         counter-name;     }     file <i>filename</i>;     interval <i>minutes</i>; }</pre>
<b>Hierarchy Level</b>	[edit accounting-options]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Create a profile to filter and collect packet and byte count statistics and write them to a file in the <code>/var/log</code> directory. To apply the profile to a firewall filter, you include the <b>accounting-profile</b> statement at the [edit firewall filter <i>filter-name</i> ] hierarchy level. For more information about firewall filters, see <a href="#">Firewall Filters Feature Guide for Routing Devices</a> .
<b>Options</b>	<p><i>profile-name</i>—Name of the filter profile.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Filter Profile</a></li></ul>

## format

---

<b>Syntax</b>	<code>format ( md5   sha1   sha256   sha512 );</code>
<b>Hierarchy Level</b>	[edit system login password]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the authentication algorithm for plain-text passwords.
<b>Default</b>	For Junos OS, the default encryption format is <b>md5</b> . For Junos-FIPS software, the default encryption format is <b>sha1</b> .
<b>Options</b>	The hash algorithm that authenticates the password can be one of these algorithms: <ul style="list-style-type: none"> <li>• <b>md5</b>—Produces a 128-bit digest.</li> <li>• <b>sha1</b>—Produces a 160-bit digest.</li> <li>• <b>sha256</b>—Produces a 256-bit digest.</li> <li>• <b>sha512</b>—Produces a 512-bit digest.</li> </ul>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Special Requirements for Junos OS Plain-Text Passwords</i></li> </ul>

## full-name

---

<b>Syntax</b>	<code>full-name</code> <i>complete-name</i> ;
<b>Hierarchy Level</b>	[edit system login user]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the complete name of a user.
<b>Options</b>	<i>complete-name</i> —Full name of the user. If the name contains spaces, enclose it in quotation marks.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS User Accounts</i></li><li>• <a href="#">user on page 577</a></li><li>• <i>user</i></li></ul>

## idle-timeout (System-Login)

---

<b>Syntax</b>	<code>idle-timeout</code> <i>minutes</i> ;
<b>Hierarchy Level</b>	[edit system login <a href="#">class</a> <i>class-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	For a login class, configure the maximum time that a session can be idle before the user is logged out of the router or switch. The session times out after remaining at the CLI operational mode prompt for the specified time.
<b>Default</b>	If you omit this statement, a user is never forced off the system after extended idle times.
<b>Options</b>	<i>minutes</i> —Maximum idle time. <b>Range:</b> 0 through 4294967295 minutes
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Timeout Value for Idle Login Sessions</i></li><li>• <a href="#">user on page 577</a></li></ul>



## interface-profile

---

<b>Syntax</b>	<pre>interface-profile <i>profile-name</i> {     <b>fields</b> {         <i>field-name</i>;     }     <b>file</b> <i>filename</i>;     <b>interval</b> <i>minutes</i>; }</pre>
<b>Hierarchy Level</b>	[edit accounting-options]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Create a profile to filter and collect error and packet statistics and write them to a file in the <code>/var/log</code> directory. You can specify an interface profile for either a physical or a logical interface.
<b>Options</b>	<p><b><i>profile-name</i></b>—Name of the interface profile.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Interface Profile</i></li> </ul>

## interval (Accounting Options)

---

<b>Syntax</b>	<code>interval minutes;</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">class-usage-profile</a> <i>profile-name</i> ], [edit accounting-options <a href="#">filter-profile</a> <i>profile-name</i> ], [edit accounting-options <a href="#">interface-profile</a> <i>profile-name</i> ], [edit accounting-options <a href="#">mib-profile</a> <i>profile-name</i> ], [edit accounting-options <a href="#">routing-engine-profile</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. The [edit accounting-options <a href="#">mib-profile</a> <i>profile-name</i> ] hierarchy level added in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify how often statistics are collected for the accounting profile.
<b>Options</b>	<b>minutes</b> —Length of time between each collection of statistics. <b>Range:</b> 1 through 2880 minutes <b>Default:</b> 30 minutes
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Interface Profile</i></li><li>• <i>Configuring the Filter Profile</i></li><li>• <i>Configuring the MIB Profile</i></li><li>• <i>Configuring the Routing Engine Profile</i></li></ul>

## login

```
Syntax login {
    announcement text;
    class class-name {
        allow-commands "regular-expression";
        allow-configuration-regexps "regular expression 1" "regular expression 2";
        configuration-breadcrumbs;
        deny-commands "regular-expression";
        ( deny-configuration | deny-configuration-regexps ) "regular expression 1" "regular
            expression 2 ";
        idle-timeout minutes;
        login-script filename;
        login-tip;
        permissions [ permissions ];
    }
    message text;
    password {
        change-type (set-transitions | character-set);
        format (md5 | sha1 | des);
        maximum-length length;
        minimum-changes number;
        minimum-length length;
    }
    retry-options {
        backoff-threshold number;
        backoff-factor seconds;
        minimum-time seconds;
        tries-before-disconnect number;
    }
    user username {
        full-name complete-name;
        uid uid-value;
        class class-name;
        authentication authentication;
        (encrypted-password "password" | plain-text-password);
        ssh-rsa "public-key";
        ssh-dsa "public-key";
    }
}
```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure user access to the router or switch.



**NOTE:** The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- *Defining Junos OS Login Classes*

## login-alarms

---

**Syntax** login-alarms;

**Hierarchy Level** [edit system login class *class-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Show system alarms automatically when an **admin** user logs in to the router or switch.

**Options** *class-name*—Login class name.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring System Alarms to Appear Automatically Upon Login*

## login-tip

---

**Syntax** login-tip;

**Hierarchy Level** [edit system login class *class-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Enable CLI tips at login.

**Default** Disabled.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring Login Tips*

---

## maximum-length

---

<b>Syntax</b>	maximum-length <i>length</i> ;
<b>Hierarchy Level</b>	[edit system login passwords]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the maximum number of characters allowed in plain-text passwords. Newly created passwords must meet this requirement.
<b>Default</b>	For Junos-FIPS software, the maximum number of characters for plain-text passwords is 20. For Junos OS, no maximum is set.
<b>Options</b>	<b>length</b> —The maximum number of characters the password can include. <b>Range:</b> 1 to 64 characters
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Special Requirements for Junos OS Plain-Text Passwords</i></li><li>• <i>Example: Changing the Requirements for Junos OS Plain-Text Passwords</i></li><li>• <a href="#">password (Login) on page 563</a></li></ul>

## message

---

<b>Syntax</b>	<code>message <i>text</i>;</code>
<b>Hierarchy Level</b>	[edit system login]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	<p>Configure a system login message. This message appears before a user logs in.</p> <p>You can format the message using the following special characters:</p> <ul style="list-style-type: none"><li>• \n—New line</li><li>• \t—Horizontal tab</li><li>• \'—Single quotation mark</li><li>• \"—Double quotation mark</li><li>• \\—Backslash</li></ul>
<b>Options</b>	<i>text</i> —Text of the message.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Junos OS to Display a System Login Message</i></li><li>• <a href="#">announcement on page 537</a></li></ul>

## mib-profile

**Syntax** `mib-profile profile-name {  
     file filename;  
     interval minutes;  
     object-names {  
         mib-object-name;  
     }  
     operation operation-name;  
 }`

**Hierarchy Level** [edit accounting-options]

**Release Information** Statement introduced in Junos OS Release 8.2.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Create a MIB profile to collect selected MIB statistics and write them to a file in the `/var/log` directory.



**NOTE:** Do not configure MIB objects related to interface octets or packets for a MIB profile, because it can cause the SNMP walk or a CLI show command to time out.

**Options** *profile-name*—Name of the MIB statistics profile.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
 interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring the MIB Profile*

## minimum-changes

---

<b>Syntax</b>	<code>minimum-changes <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit system login passwords]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	<p>Specify the minimum number of character sets (or character set changes) required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement is used in combination with the <b>change-type</b> statement. If the change-type is <b>character-sets</b>, then the number of character sets included in the password is checked against the specified minimum. If change-type is <b>set-transitions</b>, then the number of character set changes in the password is checked against the specified minimum.</p>
<b>Default</b>	For Junos OS, the minimum number of changes is 1. For Junos-FIPS Software, the minimum number of changes is 3.
<b>Options</b>	<i>number</i> —The minimum number of character sets (or character set changes) required for the password.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Special Requirements for Junos OS Plain-Text Passwords</i></li><li>• <a href="#">change-type on page 540</a></li></ul>



## minimum-length

<b>Syntax</b>	minimum-length <i>length</i> ;
<b>Hierarchy Level</b>	[edit system login passwords]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	<p>Specify the minimum number of characters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as <b>minimum-upper-cases</b>, <b>minimum-punctuations</b>, <b>minimum-lower-cases</b>, and so on.</p> <p>Using several password minimum requirement options will cause the <b>minimum-length</b> to be reset if the total sum of the required minimums exceeds the <b>minimum-length</b> setting.</p>
<b>Default</b>	For Junos OS, the minimum number of characters for plain-text passwords is six. For Junos-FIPS software, the minimum number of characters for plain-text passwords is 10.
<b>Options</b>	<b>length</b> —The minimum number of characters the password must include. <b>Range:</b> 6 to 20 characters
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Special Requirements for Junos OS Plain-Text Passwords</i></li> <li>• <i>Example: Changing the Requirements for Junos OS Plain-Text Passwords</i></li> <li>• <a href="#">maximum-length on page 557</a></li> </ul>

## object-names

---

<b>Syntax</b>	<code>object-names {     <i>mib-object-name</i>; }</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">mib-profile</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the name of each MIB object for which MIB statistics are collected for an accounting-data log file.
<b>Options</b>	<i>mib-object-name</i> —Name of a MIB object. You can specify more than one MIB object name.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the MIB Profile</i></li></ul>

## operation

---

<b>Syntax</b>	<code>operation <i>operation-name</i>;</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">mib-profile</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the name of the operation used to collect MIB statistics for an accounting-data log file.
<b>Options</b>	<i>operation-name</i> —Name of the operation to use. You can specify a <b>get</b> , <b>get-next</b> , or <b>walk</b> operation. <b>Default:</b> walk
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the MIB Profile</i></li></ul>

## password (Login)


<b>Syntax</b>	<pre>password {   change-type (set-transitions   character-set);   format (md5   sha1   sha256   sha512);   maximum-length length;   minimum-changes number;   minimum-length length;   minimum-lower-cases number;   minimum-numeric number;   minimum-punctuations number;   minimum-upper-cases number; }</pre>
<b>Hierarchy Level</b>	[edit system <a href="#">login</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Configure special requirements such as character length and encryption format for plain-text passwords. Newly created passwords must meet these requirements.</p> <p>Using several password minimum requirement options will cause the <b>minimum-length</b> to be reset if the total sum of the required minimums exceeds the <b>minimum-length</b> setting.</p> <p>The individual statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Special Requirements for Junos OS Plain-Text Passwords</i></li> <li><i>Example: Changing the Requirements for Junos OS Plain-Text Passwords</i></li> </ul>

## permissions

---

<b>Syntax</b>	<code>permissions [ <i>permissions</i> ];</code>
<b>Hierarchy Level</b>	[edit system login <a href="#">class</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the login access privileges to be provided on the router or switch.
<b>Options</b>	<i>permissions</i> —Privilege type. For a list of permission flag types, see <i>Understanding Junos OS Access Privilege Levels</i> .
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Access Privilege Levels</i></li><li>• <a href="#">user on page 577</a></li></ul>

## radius-options (edit system)

<b>Syntax</b>	<pre>radius-options {   attributes {     nas-ip-address <i>ip-address</i>;   }   enhanced-accounting;   password-protocol <i>mschap-v2</i>; }</pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>MS-CHAPv2 password protocol configuration option introduced in Junos OS Release 9.2.</p> <p>MS-CHAPv2 password protocol configuration option introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<div>  <b>NOTE:</b> The <code>radius-options</code> statement is not available on QFabric systems. </div>	
<p><b>enhanced-accounting</b> statement introduced in Junos OS Release 14.1.</p>	
<b>Description</b>	Configure RADIUS options for the NAS-IP address for outgoing RADIUS packets and password protocol used in RADIUS packets.
<b>Options</b>	<p><b>enhanced-accounting</b>—View the attribute values of a logged in user.</p> <p><b>nas-ip-address <i>ip-address</i></b>—IP address of the network access server (NAS) that requests user authentication.</p> <p><b>password-protocol <i>mschap-v2</i></b>—Protocol MS-CHAPv2, used for password authentication and password changing.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring MS-CHAPv2 for Password-Change Support</i></li> <li>• <i>Configuring RADIUS Authentication (QFX Series)</i></li> <li>• <i>Configuring RADIUS System Accounting</i></li> <li>• <i>enhanced-accounting</i></li> </ul>

## retry-options

---

<b>Syntax</b>	<pre>retry-options {     backoff-factor <i>seconds</i>;     backoff-threshold <i>number</i>;     maximum-time <i>seconds</i>;     minimum-time <i>seconds</i>;     tries-before-disconnect <i>number</i>; }</pre>
<b>Hierarchy Level</b>	[edit system <a href="#">login</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>maximum-time</b> option introduced in Junos OS Release 9.6.</p> <p><b>maximum-time</b> option introduced in Junos OS Release 9.6 for EX Series switches.</p>
<b>Description</b>	Maximum number of times a user can attempt to enter a password while logging in through SSH or Telnet before being disconnected.
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Limiting the Number of User Login Attempts for SSH and Telnet Sessions</i></li><li>• <a href="#">rate-limit on page 1668</a></li></ul>

## root-authentication

<b>Syntax</b>	<pre> root-authentication {   (encrypted-password "password"   plain-text-password);   load-key-file URL filename;   ssh-dsa "public-key";   ssh-ecdsa "public-key";   ssh-rsa "public-key"; } </pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Configure the authentication methods for the root-level user, whose username is <b>root</b>.</p>

You can use the **load-key-file** *URL filename* statement to load an SSH key file that was previously generated using **ssh-keygen**. The *URL filename* is the path to the file's location and name. When using this option, the contents of the key file are copied into the configuration immediately after entering the **load-key-file** *URL* statement. This command loads RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys.

Optionally, you can use the **ssh-dsa**, **ssh-ecdsa**, or **ssh-rsa** statements to directly configure SSH RSA, DSA, or ECDSA keys to authenticate root logins. You can configure more than one public key for SSH authentication of root logins as well as for user accounts. When a user logs in as root, the public keys are referenced to determine whether the private key matches any of them.

To view the SSH keys entries, use the configuration mode **show** command. For example:

```

[edit system]
user@host# set root-authentication load-key-file my-host:.ssh/id_dsa.pub
.file.19692 | 0 KB | 0.3 kB/s | ETA: 00:00:00 | 100%
[edit system]
user@host# show
root-authentication {
  ssh-rsa "1024 35 9727638204084251055468226757249864241630322
20740496252839038203869014158453496417001961060835872296
15634757491827360336127644187426594689320773910834481012
68312595772262546166799927831612350043866091586628382248
97467326056611921489539813965561563786211940327687806538
16960202749164163735913269396344008443 boojum@juniper.net"; #
SECRET-DATA
}

```

<b>Options</b>	<b>encrypted-password "password"</b> — MD5 or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password.
----------------	---

You cannot configure a blank password for **encrypted-password** using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

**plain-text-password**—Plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password.

**ssh-ecdsa "public/private-key"**—SSH ECDSA (variant of DSA that uses elliptic curve cryptography) public key. You can specify one or more public keys.

**ssh-dsa "public-key"**—SSH version 2 authentication. Specify the DSA (SSH version 2) public key. You can specify one or more public keys.

**ssh-rsa "public-key"**—SSH version 1 authentication. Specify the RSA (SSH version 1 and SSH version 2) public key. You can specify one or more public keys.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- *Understanding User Accounts*
- *Configuring the Root Password*
- *Recovering the Root Password*
- [authentication on page 538](#)

---

## root-login

---

**Syntax** root-login (allow | deny | deny-password);

**Hierarchy Level** [edit system services ssh]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Control user access through SSH.

**Default** Allow user access through SSH.

**Options** **allow**—Allow users to log in to the router or switch as root through SSH.

**deny**—Disable users from logging in to the router or switch as root through SSH.

**deny-password**—Allow users to log in to the router or switch as root through SSH when the authentication method (for example, RSA authentication) does not require a password.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring SSH Service for Remote Access to the Router or Switch*



## routing-engine-profile

---

<b>Syntax</b>	<pre> routing-engine-profile <i>profile-name</i> {     fields {         <i>field-name</i>;     }     file <i>filename</i>;     interval <i>minutes</i>; } </pre>
<b>Hierarchy Level</b>	[edit accounting-options]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Create a Routing Engine profile to collect selected Routing Engine statistics and write them to a file in the <code>/var/log</code> directory.
<b>Options</b>	<p><b><i>profile-name</i></b>—Name of the Routing Engine statistics profile.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Routing Engine Profile</i></li> </ul>

## size

---

<b>Syntax</b>	<code>size bytes;</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">file</a> <i>filename</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify attributes of an accounting-data log file.
<b>Options</b>	<b>bytes</b> —Maximum size of each log file, in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). When a log file (for example, <b>profilelog</b> ) reaches its maximum size, it is renamed <b>profilelog.0</b> , then <b>profilelog.1</b> , and so on, until the maximum number of log files is reached. Then the oldest log file is overwritten. If you do not specify a size, the file is closed, archived, and renamed when the time specified for the transfer interval is exceeded.  <b>Syntax:</b> <i>x</i> to specify bytes, <i>xk</i> to specify KB, <i>xm</i> to specify MB, <i>xg</i> to specify GB <b>Range:</b> 256 KB through 1 GB
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Maximum Size of the File</i></li></ul>

## source-classes

---

<b>Syntax</b>	<pre>source-classes {     source-class-name; }</pre>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">class-usage-profile</a> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the source classes for which statistics are collected.
<b>Options</b>	<b>source-class-name</b> —Name of the source class to include in the class usage profile.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring a Class Usage Profile</i></li></ul>

---

## start-time (Log File Transfer)

---

<b>Syntax</b>	<code>start-time <i>time</i>;</code>
<b>Hierarchy Level</b>	[edit accounting-options <a href="#">file</a> <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the start time for transfer of an accounting-data log file.
<b>Options</b>	<i>time</i> —Start time for file transfer. <b>Syntax:</b> <i>YYYY-MM-DD.hh:mm</i>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Start Time for File Transfer</i></li></ul>

## tacplus-options

---

<b>Syntax</b>	<pre>tacplus-options {   (exclude-cmd-attribute   no-cmd-attribute-value);   enhanced-accounting;   service-name <i>service-name</i>;   timestamp-and-timezone; }</pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>no-cmd-attribute-value</b> and <b>exclude-cmd-attribute</b> options introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 11.1 for QFX Series.</p> <p><b>timestamp-and-timezone</b> option introduced in Junos OS Release 12.2.</p> <p><b>enhanced-accounting</b> option introduced in Junos OS Release 14.1.</p>
<b>Description</b>	Configure TACACS+ options for authentication and accounting.
<b>Options</b>	<p><b>enhanced-accounting</b>—View the attribute values of a logged in user.</p> <p><b>exclude-cmd-attribute</b>—Exclude the <b>cmd</b> attribute value completely from start and stop accounting records to enable logging of accounting records in the correct log file on a TACACS+ server.</p> <p><b>no-cmd-attribute-value</b>—Set the <b>cmd</b> attribute value to an empty string in the TACACS+ accounting start and stop requests to enable logging of accounting records in the correct log file on a TACACS+ server.</p> <p><b>service-name <i>service-name</i></b>—Name of the authentication service used when you configure multiple TACACS+ servers to use the same authentication service.</p> <p><b>Default:</b> junos-exec</p> <p><b>timestamp-and-timezone</b>—Include this statement if you want start time, stop time, and timezone attributes included in start/stop accounting records.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring TACACS+ Authentication</i></li><li>• <i>Configuring TACACS+ System Accounting</i></li><li>• <i>Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication</i></li><li>• <i>enhanced-accounting</i></li></ul>

---

## tacplus-server

---

<b>Syntax</b>	<pre>tacplus-server server-address {     secret password;     single-connection;     source-address source-address;     timeout seconds; }</pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the IPv4 or IPv6 TACACS+ server.
<b>Options</b>	<b>server-address</b> —Address of the IPv4 or IPv6 TACACS+ authentication server.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring TACACS+ Authentication</i></li></ul>

## traceoptions (Address-Assignment Pool)

---

**Syntax**    `traceoptions {  
              file filename {  
                  files number;  
                  size maximum-file-size;  
                  match regex;  
                  (world-readable | no-world-readable);  
              }  
              flag address-assignment;  
              flag all;  
              flag configuration;  
              flag framework;  
              flag ldap;  
              flag local-authentication;  
              flag radius;  
          }`

**Hierarchy Level**    [edit system processes general-authentication-service]

**Release Information**    Flag for tracing address-assignment pool operations introduced in Junos OS Release 9.0.  
                              **option-name** option introduced in Junos OS Release 8.3.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description**    Configure tracing options.

**Options**    **file *filename***—Name of the file that receives the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory **/var/log**.

**files *number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

              If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

**Range:** 2 through 1000

**Default:** 3 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **address-assignment**—All address-assignment events
- **all**—All tracing operations
- **configuration**—Configuration events
- **framework**—Authentication framework events
- **ldap**—LDAP authentication events
- **local-authentication**—Local authentication events

- **radius**—RADIUS authentication events

**match *regex***—(Optional) Refine the output to include lines that contain the regular expression.

**no-world-readable**—(Optional) Restrict access to the originator of the trace operation only.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring Address-Assignment Pools*

## transfer-interval

**Syntax** transfer-interval *minutes*;

**Hierarchy Level** [edit accounting-options **file** *filename*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Specify the length of time the file remains open and receives new statistics before it is closed and transferred to an archive site.

**Options** ***minutes***—Time the file remains open and receives new statistics before it is closed and transferred to an archive site.

**Range:** 5 through 2880 minutes

**Default:** 30 minutes

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring the Transfer Interval of the File*

## uid

---

<b>Syntax</b>	<code>uid <i>uid-value</i>;</code>
<b>Hierarchy Level</b>	[edit system login <a href="#">user</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Numeric identifier associated with the user account name, either assigned by an administrator or assigned automatically when you commit the user configuration. It is used by applications that request numeric identifiers, such as some RADIUS queries or secure applications such as flow-tap monitoring.
<b>Options</b>	<b><i>uid-value</i></b> —Number associated with the login account. This value must be unique on the router or switch. <b>Range:</b> 100 through 64000
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS User Accounts</i></li></ul>



## user (Access)

---

<b>Syntax</b>	<pre> user username {   authentication {     class class-name;     (encrypted-password "password"   plain-text-password);     full-name complete-name;     load-key-file URL filename;     ssh-dsa "public-key" &lt;from hostname&gt;;     ssh-rsa "public-key" &lt;from hostname&gt;;     uid uid-value;   } } </pre>
<b>Hierarchy Level</b>	[edit system <a href="#">login</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Configure access permission for individual users.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Junos OS User Accounts</i></li> <li>• <a href="#">class on page 540</a></li> </ul>



# Administration

- [Routine Monitoring on page 579](#)
- [Operational Commands on page 581](#)

## Routine Monitoring

---

- [Managing Users \(J-Web Procedure\) on page 579](#)

### Managing Users (J-Web Procedure)

You can use the Users Configuration page for user information to add new users to an EX Series switch. For each account, you define a login name and password for the user and specify a login class for access privileges.

To configure users:

1. Select **Configure > System Properties > User Management**.

The User Management page displays details of users, the authentication order, the RADIUS servers and TACACS servers present.

2. Click **Edit**.
3. Click any of the following options on the **Users** tab:
  - **Add**—Select this option to add a user. Enter details as described in [Table 48 on page 580](#).
  - **Edit**—Select this option to edit an existing user's details. Enter details as described in [Table 48 on page 580](#).
  - **Delete**—Select this option to delete a user.
4. Click an option on the **Authentication Methods and Order** tab:
  - **Authentication Order**—Drag and drop the authentication type from the Available Methods section to the Selected Methods. Click the up or down buttons to modify the authentication order.
  - **RADIUS server**—Click one of the following options:
    - **Add**—Select this option to add an authentication server. Enter details as described in [Table 49 on page 581](#).

- **Edit**—Select this option to modify the authentication server details. Enter details as described in [Table 49 on page 581](#).
- **Delete**—Select this option to delete an authentication server from the list.
- TACACS server—Click one of the following options:
  - **Add**—Select this option to add an authentication server. Enter details as described in [Table 49 on page 581](#).
  - **Edit**—Select this option to modify the authentication server details. Enter details as described in [Table 49 on page 581](#).
  - **Delete**—Select this option to delete an authentication server from the list.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

**Table 48: User Management Configuration Page Summary**

Field	Function	Your Action
<b>User Information</b>		
Username (required)	Specifies the name that identifies the user.	Type the username. It must be unique within the switching platform. Do not include spaces, colons, or commas in the username.
User Id	Specifies the user identification.	Type the user's ID.
Full Name	Specifies the user's full name.	Type the user's full name. If the full name contains spaces, enclose it in quotation marks. Do not include colons or commas.
Login Class (required)	Defines the user's access privilege.	Select the user's login class from the list: <ul style="list-style-type: none"> <li>• <b>operator</b></li> <li>• <b>read-only</b></li> <li>• <b>super-user/superuser</b></li> <li>• <b>unauthorized</b></li> </ul> This list also includes any user-defined login classes.
Password	Specifies the login password for this user.	Type the login password for this user. The login password must meet these criteria: <ul style="list-style-type: none"> <li>• The password must be at least 6 characters long.</li> <li>• It can include alphabetic, numeric, and special characters, but not control characters.</li> <li>• It must contain at least one change of case or character class.</li> </ul>

Table 48: User Management Configuration Page Summary (*continued*)

Field	Function	Your Action
Confirm Password	Verifies the login password for this user.	Retype the login password for this user.

Table 49: Add an Authentication Server

Field	Function	Your Action
IP Address	Specifies the IP address of the server.	Type the server's 32-bit IP address, in dotted decimal notation.
Password	Specifies the password of the server.	Type the password of the server.
Confirm Password	Verifies that the password of the server is entered correctly.	Retype the password of the server.
Server Port	Specifies the port with which the server is associated.	Type the port number.
Source Address	Specifies the source address of the server.	Type the server's 32-bit IP address, in dotted decimal notation.
Retry Attempts	Specifies the number of login retries allowed after a login failure.	Type the number.  <b>NOTE:</b> Only 1 retry is permitted for a TACACS server.
Time out	Specifies the time interval to wait before the connection to the server is closed.	Type the interval in seconds.

**Related Documentation**

- [Configuring Management Access for the EX Series Switch \(J-Web Procedure\) on page 529](#)

## Operational Commands

- [request message](#)
- [show subscribers](#)

## request message

---

<b>Syntax</b>	<code>request message all message "text"</code> <code>request message message "text" (terminal <i>terminal-name</i>   user <i>user-name</i>)</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display a message on the screens of all users who are logged in to the router or switch or on specific screens.
<b>Options</b>	<b>all</b> —Display a message on the terminal of all users who are currently logged in.  <b>message "text"</b> —Message to display.  <b>terminal <i>terminal-name</i></b> —Name of the terminal on which to display the message.  <b>user <i>user-name</i></b> —Name of the user to whom to direct the message.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">request message message on page 582</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request message message

```
user@host> request message message "Maintenance window in 10 minutes" user maria
Message from user@host on tty0 at 20:27 ...
Maintenance window in 10 minutes
EOF
```

## show subscribers

**Syntax** show subscribers  
 <detail | extensive | terse>  
 <aci-interface-set-name *aci-interface-set-name*>  
 <address *address*>  
 <agent-circuit-identifier *agent-circuit-identifier-substring*>  
 <client-type *client-type*>  
 <count>  
 <id>  
 <interface *interface*>  
 <logical-system *logical-system*>  
 <mac-address *mac-address*>  
 <physical-interface *physical-interface-name*>  
 <profile-name *profile-name*>  
 <routing-instance *routing-instance*>  
 <stacked-vlan-id *stacked-vlan-id*>  
 <subscriber-state *subscriber-state*>  
 <user-name *user-name*>  
 <vci *vci-identifier*>  
 <vpi *vpi-identifier*>  
 <vlan-id *vlan-id*>

**Release Information** Command introduced in Junos OS Release 9.3.  
 Command introduced in Junos OS Release 9.3 for EX Series switches.  
**client-type**, **mac-address**, **subscriber-state**, and **extensive** options introduced in Junos OS Release 10.2.  
**count** option usage with other options introduced in Junos OS Release 10.2.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.  
 Options **aci-interface-set-name** and **agent-circuit-identifier** introduced in Junos OS Release 12.2.  
 The **physical-interface** and **user-name** options introduced in Junos OS Release 12.3.  
 Options **vci** and **vpi** introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.  
 Options **vci** and **vpi** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

**Description** Display information for active subscribers.

**Options** **detail | extensive | terse**—(Optional) Display the specified level of output.

**aci-interface-set-name**—(Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the DHCP or PPPoE control packets.

**address**—(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.168.17.1). If you specify the IP address as a prefix with a netmask (for example, 192.168.17.1/32), the router displays a message that the IP address is invalid, and rejects the command.

***agent-circuit-identifier-substring***—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified substring.

***client-type***—(Optional) Display subscribers whose client type matches the specified client type (DHCP, L2TP, PPP, PPPOE, VLAN, or static).

***count***—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the ***count*** option alone or with the ***address***, ***client-type***, ***interface***, ***logical-system***, ***mac-address***, ***profile-name***, ***routing-instance***, ***stacked-vlan-id***, ***subscriber-state***, or ***vlan-id*** options.

***id***—(Optional) Display a specific subscriber session whose session id matches the specified subscriber ID. You can display subscriber IDs by using the ***show subscribers extensive*** or the ***show subscribers interface extensive*** commands.

***interface***—(Optional) Display subscribers whose interface matches the specified interface.

***logical-system***—(Optional) Display subscribers whose logical system matches the specified logical system.

***mac-address***—(Optional) Display subscribers whose MAC address matches the specified MAC address.

***physical-interface-name***—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

***profile-name***—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

***routing-instance***—(Optional) Display subscribers whose routing instance matches the specified routing instance.

***stacked-vlan-id***—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

***subscriber-state***—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

***user-name***—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

***vci-identifier***—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is **0** through **255**.

***vpi-identifier***—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is **0** through **65535**.

***vlan-id***—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID.





**NOTE:** Due to display limitations, logical system and routing instance output values are truncated when necessary.

**Required Privilege Level**

view

**Related Documentation**

- *show subscribers summary*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

**List of Sample Output**

[show subscribers \(IPv4\) on page 589](#)  
[show subscribers \(IPv6\) on page 589](#)  
[show subscribers \(IPv4 and IPv6 Dual Stack\) on page 589](#)  
[show subscribers \(LNS on MX Series Routers\) on page 590](#)  
[show subscribers \(L2TP Switched Tunnels\) on page 590](#)  
[show subscribers client-type dhcp detail on page 590](#)  
[show subscribers count on page 590](#)  
[show subscribers address detail \(IPv6\) on page 590](#)  
[show subscribers detail \(IPv4\) on page 591](#)  
[show subscribers detail \(IPv6\) on page 591](#)  
[show subscribers detail \(IPv6 Static Demux Interface\) on page 592](#)  
[show subscribers detail \(L2TP LNS Subscribers on MX Series Routers\) on page 592](#)  
[show subscribers detail \(L2TP Switched Tunnels\) on page 592](#)  
[show subscribers detail \(Tunneled Subscriber\) on page 593](#)  
[show subscribers detail \(IPv4 and IPv6 Dual Stack\) on page 593](#)  
[show subscribers detail \(ACI Interface Set Session\) on page 594](#)  
[show subscribers detail \(PPPoE Subscriber Session with ACI Interface Set\) on page 594](#)  
[show subscribers extensive on page 594](#)  
[show subscribers extensive \(RPF Check Fail Filter\) on page 595](#)  
[show subscribers extensive \(L2TP LNS Subscribers on MX Series Routers\) on page 595](#)  
[show subscribers extensive \(IPv4 and IPv6 Dual Stack\) on page 595](#)  
[show subscribers extensive \(Effective Shaping-Rate\) on page 596](#)  
[show subscribers aci-interface-set-name detail \(Subscriber Sessions Using Specified ACI Interface Set\) on page 597](#)  
[show subscribers agent-circuit-identifier detail \(Subscriber Sessions Using Specified ACI Substring\) on page 597](#)  
[show subscribers interface extensive on page 598](#)  
[show subscribers logical-system terse on page 598](#)  
[show subscribers physical-interface count on page 599](#)  
[show subscribers routing-instance inst1 count on page 599](#)  
[show subscribers stacked-vlan-id detail on page 599](#)  
[show subscribers stacked-vlan-id vlan-id detail \(Combined Output\) on page 599](#)  
[show subscribers stacked-vlan-id vlan-id interface detail \(Combined Output for a Specific Interface\) on page 599](#)  
[show subscribers user-name detail on page 599](#)  
[show subscribers vlan-id on page 600](#)

[show subscribers vlan-id detail on page 600](#)

[show subscribers vpi vci extensive \(PPPoE-over-ATM Subscriber Session\) on page 600](#)

**Output Fields** [Table 50 on page 586](#) lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

**Table 50: show subscribers Output Fields**

Field Name	Field Description
<b>Interface</b>	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface.  The * character indicates a continuation of addresses for the same session.
<b>IP Address/VLAN ID</b>	Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i>  No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions the value is <b>Tunnel-switched</b> .
<b>User Name</b>	Name of subscriber.
<b>LS:RI</b>	Logical system and routing instance associated with the subscriber.
<b>Type</b>	Subscriber client type (DHCP, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
<b>IP Address</b>	Subscriber IPv4 address.
<b>IP Netmask</b>	Subscriber IP netmask.
<b>Primary DNS Address</b>	IP address of primary DNS server.
<b>Secondary DNS Address</b>	IP address of secondary DNS server.
<b>Primary WINS Address</b>	IP address of primary WINS server.
<b>Secondary WINS Address</b>	IP address of secondary WINS server.
<b>IPv6 Address</b>	Subscriber IPv6 address, or multiple addresses.
<b>IPv6 Prefix</b>	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
<b>IPv6 User Prefix</b>	IPv6 prefix obtained through ND/RA.
<b>IPv6 Address Pool</b>	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
<b>IPv6 Network Prefix Length</b>	Length of the network portion of the IPv6 address.
<b>IPv6 Prefix Length</b>	Length of the subscriber IPv6 prefix.

Table 50: show subscribers Output Fields (*continued*)

Field Name	Field Description
<b>Logical System</b>	Logical system associated with the subscriber.
<b>Routing Instance</b>	Routing instance associated with the subscriber.
<b>Interface Type</b>	Whether the subscriber interface is <b>Static</b> or <b>Dynamic</b> .
<b>Interface Set</b>	Internally generated name of the dynamic ACI interface set used by the subscriber session.
<b>Interface Set Type</b>	Interface type of the ACI interface set: <b>Dynamic</b> . This is the only ACI interface set type currently supported.
<b>Interface Set Session ID</b>	Identifier of the dynamic ACI interface set entry in the session database.
<b>Underlying Interface</b>	Name of the underlying interface for the subscriber session.
<b>Dynamic Profile Name</b>	Dynamic profile used for the subscriber.
<b>Dynamic Profile Version</b>	Version number of the dynamic profile used for the subscriber.
<b>MAC Address</b>	MAC address associated with the subscriber.
<b>State</b>	Current state of the subscriber session ( <b>Init</b> , <b>Configured</b> , <b>Active</b> , <b>Terminating</b> , <b>Tunneled</b> ).
<b>L2TP State</b>	Current state of the L2TP session, <b>Tunneled</b> or <b>Tunnel-switched</b> . When the value is <b>Tunnel-switched</b> , two entries are displayed for the subscriber; the first entry is at the LNS interface on the LTS and the second entry is at the LAC interface on the LTS.
<b>Tunnel switch Profile Name</b>	Name of the L2TP tunnel switch profile that initiates tunnel switching.
<b>Local IP Address</b>	IP address of the local gateway (LAC).
<b>Remote IP Address</b>	IP address of the remote peer (LNS).
<b>VLAN Id</b>	VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
<b>Stacked VLAN Id</b>	Stacked VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
<b>RADIUS Accounting ID</b>	RADIUS accounting ID associated with the subscriber.
<b>Agent Circuit ID</b>	Option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.
<b>Agent Remote ID</b>	Option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.
<b>DHCP Relay IP Address</b>	IP address used by the DHCP relay agent.

Table 50: show subscribers Output Fields (*continued*)

Field Name	Field Description
ATM VPI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
ATM VCI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.
Login Time	Date and time at which the subscriber logged in.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.
IPv4 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
IPv6 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
Session ID	ID number for a subscriber service session.
Underlying Session ID	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
Service Session Name	Service session profile name.
Session Timeout (seconds)	Number of seconds of access provided to the subscriber before the session is automatically terminated.
Idle Timeout (seconds)	Number of seconds subscriber can be idle before the session is automatically terminated.
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.

Table 50: show subscribers Output Fields (*continued*)

Field Name	Field Description
<b>ADF IPv4 Output Filter Name</b>	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
<b>ADF IPv6 Input Filter Name</b>	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
<b>ADF IPv6 Output Filter Name</b>	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
<b>IPv4 Input Filter Name</b>	Name assigned to the IPv4 input filter (client or service session).
<b>IPv4 Output Filter Name</b>	Name assigned to the IPv4 output filter (client or service session).
<b>IPv6 Input Filter Name</b>	Name assigned to the IPv6 input filter (client or service session).
<b>IPv6 Output Filter Name</b>	Name assigned to the IPv6 output filter (client or service session).
<b>IFL Input Filter Name</b>	Name assigned to the logical interface input filter (client or service session).
<b>IFL Output Filter Name</b>	Name assigned to the logical interface output filter (client or service session).

## Sample Output

### show subscribers (IPv4)

```

user@host> show subscribers
Interface          IP Address/VLAN ID  User Name          LS:RI
ge-1/3/0.1073741824 100                 WHOLESALE-CLIENT  default:default
demux0.1073741824   100.0.0.10         RETAILER1-CLIENT  test1:retailer1
demux0.1073741825   101.0.0.3          RETAILER2-CLIENT  test1:retailer2
demux0.1073741826   102.0.0.3          RETAILER2-CLIENT  test1:retailer2

```

### show subscribers (IPv6)

```

user@host> show subscribers
Interface          IP Address/VLAN ID  User Name          LS:RI
ge-1/0/0.0         2001::c0:0:0:0/74  WHOLESALE-CLIENT  default:default
*                  2002::1/128        subscriber-25      default:default

```

### show subscribers (IPv4 and IPv6 Dual Stack)

```

user@host> show subscribers
Interface          IP Address/VLAN ID  User Name
LS:RI
demux0.1073741834  0x8100.1002 0x8100.1
default:default
demux0.1073741835  0x8100.1001 0x8100.1
default:default
pp0.1073741836     61.1.1.1        dualstackuser1@ISP1.com

```

```
default:ASP-1
*                2041:1:1::/48
*                2061:1:1:1::/64
pp0.1073741837   23.1.1.3                dualstackuser2@ISP1.com
default:ASP-1
*                2001:1:2:5::/64
```

#### show subscribers (LNS on MX Series Routers)

```
user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
si-4/0/0.1     192.168.4.1        xyz@example.com default:default
```

#### show subscribers (L2TP Switched Tunnels)

```
user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
si-2/1/0.1073741842 Tunnel-switched    ap@lts.com     default:default
si-2/1/0.1073741843 Tunnel-switched    ap@lts.com     default:default
```

#### show subscribers client-type dhcp detail

```
user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 100.20.10.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:94:00:01:f3
State: Active
Radius Accounting ID: jnpr :2560
Login Time: 2009-08-25 14:43:56 PDT
```

#### show subscribers count

```
user@host> show subscribers count
Total Subscribers: 188, Active Subscribers: 188
```

#### show subscribers address detail (IPv6)

```
user@host> show subscribers address 100.16.12.137 detail
```

```

Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 100.16.12.137
IP Netmask: 255.0.0.0
IPv6 User Prefix: 1016:0:0:c88::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:0d:02:01:00:01
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: jnpr demux0.8201:6544
Session ID: 6544
Agent Circuit ID: if13720
Agent Remote ID: if13720
Login Time: 2012-05-21 13:37:27 PDT
Service Sessions: 1

```

#### show subscribers detail (IPv4)

```

user@host> show subscribers detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Primary DNS Address: 192.168.17.1
Secondary DNS Address: 192.168.17.2
Primary WINS Address: 192.168.22.1
Secondary WINS Address: 192.168.22.2
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Idle Timeout (seconds): 600
Login Time: 2009-08-25 14:43:52 PDT
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
33 2d 37 2d 30 37 05 01 06 0f 21 2c
Service Sessions: 2

```

#### show subscribers detail (IPv6)

```

user@host> show subscribers detail
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03
State: Active
Radius Accounting ID: 1

```

```
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
```

#### show subscribers detail (IPv6 Static Demux Interface)

```
user@host> show subscribers detail
Type: STATIC-INTERFACE
User Name: demux0.1@jnpr.net
IPv6 Prefix: 1:2:3:4:5:6:7:aa/128
Logical System: default
Routing Instance: default
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT
```

#### show subscribers detail (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers detail
Type: L2TP
User Name: user1@jnpr.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
```

#### show subscribers detail (L2TP Switched Tunnels)

```
user@host> show subscribers detail
Type: L2TP
User Name: ap@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741842
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.50.1.1
Remote IP Address: 192.168.20.3
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST

Type: L2TP
User Name: ap@example.com
Logical System: default
```



```

Routing Instance: default
Interface: si-2/1/0.1073741843
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.30.1.1
Remote IP Address: 172.20.1.10
Session ID: 22
Login Time: 2013-01-18 03:01:14 PST

```

#### show subscribers detail (Tunneled Subscriber)

```

user@host> show subscribers detail
Type: PPPoE
User Name: user1@example.com
Logical System: default
Routing Instance: default
Interface: pp0.1
State: Active, Tunneled
Radius Accounting ID: 512

```

#### show subscribers detail (IPv4 and IPv6 Dual Stack)

```

user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@ISP1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02

```

```
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
```

#### show subscribers detail (ACI Interface Set Session)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0
Interface Set: aci-1001-ge-1/0/0.2800
Interface Set Session ID: 0
Underlying Interface: ge-1/0/0.2800
Dynamic Profile Name: aci-vlan-set-profile-2
Dynamic Profile Version: 1
State: Active
Session ID: 1
Agent Circuit ID: aci-ppp-dhcp-20
Login Time: 2012-05-26 01:54:08 PDT
```

#### show subscribers detail (PPPoE Subscriber Session with ACI Interface Set)

```
user@host> show subscribers detail
Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.5
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Interface Set: aci-1001-demux0.1073741824
Interface Set Type: Dynamic
Interface Set Session ID: 2
Underlying Interface: demux0.1073741824
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:64:39:01:02
State: Active
Radius Accounting ID: 3
Session ID: 3
Agent Circuit ID: aci-ppp-dhcp-dvlan-50
Login Time: 2012-03-07 13:46:53 PST
```

#### show subscribers extensive

```
user@host> show subscribers extensive
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03
```

```

State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48

```

#### show subscribers extensive (RPF Check Fail Filter)

```

user@host> show subscribers extensive
...
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ae0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof
State: Active
Session ID: 9
VLAN Id: 100
Login Time: 2011-08-26 08:17:00 PDT
IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...

```

#### show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

```

user@host> show subscribers extensive
Type: L2TP
User Name: user1@jnpr.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out

```

#### show subscribers extensive (IPv4 and IPv6 Dual Stack)

```

user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1

```

```
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@ISP1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2061:1:1:1::1/64
IPv6 Framed Interface Id: 1:1:2:2
IPv4 Input Filter Name: FILTER-IN-pp0.1073741825-in
IPv4 Output Filter Name: FILTER-OUT-pp0.1073741825-out
IPv6 Input Filter Name: FILTER-IN6-pp0.1073741825-in
IPv6 Output Filter Name: FILTER-OUT6-pp0.1073741825-out

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Network Prefix Length: 48
```

#### show subscribers extensive (Effective Shaping-Rate)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741837
Interface type: Dynamic
Interface Set: ifset-1
Underlying Interface: ae1
Dynamic Profile Name: svlan-dhcp-test
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.201
VLAN Id: 0x8100.201
Login Time: 2011-11-30 00:18:04 PST
```

Effective shaping-rate: 31000000k

...

#### show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set)

```
user@host> show subscribers aci-interface-set-name aci-1003-ge-1/0/0.4001 detail
```

```
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT
```

```
Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT
```

#### show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring)

```
user@host> show subscribers agent-circuit-identifier aci-ppp-vlan detail
```

```
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT
```

```
Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
```

**Interface Set Type: Dynamic**  
**Interface Set Session ID: 13**  
Underlying Interface: ge-1/0/0.4001  
Dynamic Profile Name: aci-vlan-pppoe-profile  
Dynamic Profile Version: 1  
MAC Address: 00:00:65:26:01:02  
State: Active  
Radius Accounting ID: 14  
Session ID: 14  
**Agent Circuit ID: aci-ppp-vlan-10**  
Login Time: 2012-03-12 10:41:57 PDT

#### show subscribers interface extensive

```
user@host> show subscribers interface demux0.1073741826 extensive
Type: VLAN
User Name: test1@test.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 12
Session ID: 12
Stacked VLAN Id: 0x8100.1500
VLAN Id: 0x8100.2902
Login Time: 2011-10-20 16:21:59 EST

Type: DHCP
User Name: test1@test.com
IP Address: 172.16.200.6
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 21
Session ID: 21
Login Time: 2011-10-20 16:24:33 EST
Service Sessions: 2

Service Session ID: 25
Service Session Name: SUB-QOS
State: Active

Service Session ID: 26
Service Session Name: service-cb-content
State: Active
IPv4 Input Filter Name: content-cb-in-demux0.1073741826-in
IPv4 Output Filter Name: content-cb-out-demux0.1073741826-out
```

#### show subscribers logical-system terse

```
user@host> show subscribers logical-system test1 terse
Interface      IP Address/VLAN ID  User Name      LS:RI
demux0.1073741825  101.0.0.3          RETAILER1-CLIENT test1:retailer1
demux0.1073741826  102.0.0.3          RETAILER2-CLIENT test1:retailer2
```

### show subscribers physical-interface count

```
user@host> show subscribers physical-interface ge-1/0/0 count
Total subscribers: 3998, Active Subscribers: 3998
```

### show subscribers routing-instance inst1 count

```
user@host> show subscribers routing-instance inst1 count
Total Subscribers: 188, Active Subscribers: 183
```

### show subscribers stacked-vlan-id detail

```
user@host> show subscribers stacked-vlan-id 101 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

### show subscribers stacked-vlan-id vlan-id detail (Combined Output)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

### show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

### show subscribers user-name detail

```
user@host> show subscribers user-name larry1 detail
Type: DHCP
User Name: larry1
IP Address: 100.0.0.37
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.1
Interface type: Static
Dynamic Profile Name: foo
MAC Address: 00:10:94:00:00:01
State: Active
Radius Accounting ID: 1
Session ID: 1
```

```
Login Time: 2011-11-07 08:25:59 PST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
37 2d 30 2d 30 37 05 01 06 0f 21 2c
```

#### show subscribers vlan-id

```
user@host> show subscribers vlan-id 100
Interface          IP Address          User Name
ge-1/0/0.1073741824
ge-1/2/0.1073741825
```

#### show subscribers vlan-id detail

```
user@host> show subscribers vlan-id 100 detail
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT
```

#### show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session)

```
user@host> show subscribers vpi 40 vci 50 extensive
Type: PPPoE
User Name: testuser
IP Address: 100.0.0.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:65:23:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
ATM VPI: 40
ATM VCI: 50
Login Time: 2012-12-03 07:49:26 PST
IP Address Pool: pool_1
IPv6 Framed Interface Id: 200:65ff:fe23:102
```



# Troubleshooting Procedures

- Troubleshooting Loss of the Root Password on page 601

## Troubleshooting Loss of the Root Password

---

**Problem**    **Description:** If you forget the root password for a switch, use the password recovery procedure to reset the root password.



**NOTE:** You need physical access to the switch to recover the root password.




**TIP:** For a video on recovering the root password for routers, see *Recovering the Root Password*. The procedure is similar for switches.

**Solution**    To recover the root password:

1. Power off your switch by unplugging the power cord or turning off the power at the wall switch.
2. Insert one end of the Ethernet cable into the serial port on the management device and connect the other end to the console port on the back of the switch. See [Figure 1 on page 306](#).

Figure 2: Connecting to the Console Port on the EX Series Switch



3. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate COM port to use (for example, COM1).
4. Configure the port settings as follows:
  - Bits per second: 9600
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: None
5. Power on your switch by plugging in the power cord or turning on the power at the wall switch.
6. When the following prompt appears, press the Spacebar to access the switch's bootstrap loader command prompt:
 

Hit [Enter] to boot immediately, or space bar for command prompt.  
Booting [kernel] in 1 second...



**NOTE:** If the switch is in unattended mode for U-Boot, access to the bootstrap loader command prompt is blocked. If the root password is lost, you must reset the switch to the factory default configuration using the LCD panel. For more information, see [“Reverting to the Default Factory Configuration for the EX Series Switch”](#) on page 620.

7. At the following prompt, type **boot -s** to start up the system in single-user mode:  
**loader> boot -s**
8. At the following prompt, type **recovery** to start the root password recovery procedure:  
**Enter full path name of shell or 'recovery' for root password recovery or RETURN for /bin/sh: recovery**

A series of messages describe consistency checks, mounting of filesystems, and initialization and checkout of management services. Then the CLI prompt appears.

9. Enter configuration mode in the CLI:

```
user@switch> configure
```

10. Set the root password. For example:

```
user@switch# set system root-authentication plain-text-password
```

11. At the following prompt, enter the new root password. For example, juniper1:

```
user@switch# juniper1
```

```
Retype new password:
```

12. At the second prompt, reenter the new root password.

13. If you are finished configuring the network, commit the configuration.

```
root@switch# commit
```

```
commit complete
```

14. Exit configuration mode in the CLI.

```
root@switch# exit
```

15. Exit operational mode in the CLI.

```
root@switch> exit
```

16. At the prompt, enter **y** to reboot the switch.

```
Reboot the system? [y/n] y
```

#### **Related Documentation**

- *Connecting and Configuring an EX Series Switch (CLI Procedure)*
- *Connecting and Configuring an EX Series Switch (J-Web Procedure)*
- For information about configuring an encrypted root password, configuring SSH keys to authenticate root logins, and configuring special requirements for plain-text passwords, see *Configuring the Root Password*.



## PART 5

# Configuration and File Management

- [Overview on page 607](#)
- [Configuration on page 617](#)
- [Administration on page 641](#)
- [Troubleshooting Procedures on page 685](#)



## CHAPTER 13

# Overview

- [Software Overview on page 607](#)
- [Configuration Files Overview on page 609](#)

## Software Overview

---

- [Understanding Software Infrastructure and Processes on page 607](#)

### Understanding Software Infrastructure and Processes

Each switch runs the Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches on its general-purpose processors. Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the chassis.

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

With the J-Web interface and the command-line interface (CLI) to the Junos OS, you configure switching features and routing protocols and set the properties of network interfaces on your switch. After activating a software configuration, use either the J-Web or CLI user interface to monitor the switch, manage operations, and diagnose protocol and network connectivity problems.

- [Routing Engine and Packet Forwarding Engine on page 607](#)
- [Junos OS Processes on page 608](#)

### Routing Engine and Packet Forwarding Engine

---

A switch has two primary software processing components:

- **Packet Forwarding Engine**—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.
- **Routing Engine**—Provides three main functions:
  - Creates the packet forwarding switch fabric for the switch, providing route lookup, filtering, and switching on incoming data packets, then directing outbound packets to the appropriate interface for transmission to the network

- Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
- Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

### Junos OS Processes

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

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

Table 6 on page 30 describes the primary Junos OS processes.

**Table 51: Junos OS Processes**

Process	Name	Description
Chassis process	chassisd	<p>Detects hardware on the system that is used to configure network interfaces.</p> <p>Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered.</p> <p>Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.</p>
Ethernet switching process	eswd	<p>Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree protocol and access port security. The process is also responsible for managing Ethernet switching interfaces, VLANs, and VLAN interfaces.</p> <p>Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.</p>
Forwarding process	pfem	<p>Defines how routing protocols operate on the switch. The overall performance of the switch is largely determined by the effectiveness of the forwarding process.</p>
Interface process	dcd	<p>Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.</p>
Management process	mgd	<p>Provides communication between the other processes and an interface to the configuration database.</p> <p>Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured.</p> <p>Interacts with the other processes when commands are issued through one of the user interfaces on the switch.</p> <p>If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.</p>
Routing protocol process	rpd	<p>Defines how routing protocols such as RIP, OSPF, and BGP operate on the device, including selecting routes and maintaining forwarding tables.</p>



- Related Documentation**
- For more information about processes, see *Junos OS Network Operations Guide*
  - For more information about basic system parameters, supported protocols, and software processes, see *Junos OS System Basics Configuration Guide*

## Configuration Files Overview

- Understanding Configuration Files for EX Series Switches on page 609
- Using the CLI Viewer in the J-Web Interface to View Configuration Text on page 610
- Using the CLI Editor in the J-Web Interface to Edit Configuration Text on page 610
- Using the Point and Click CLI Tool in the J-Web Interface to Edit Configuration Text on page 611
- Understanding Automatic Refreshing of Scripts on EX Series Switches on page 613
- Understanding Autoinstallation of Configuration Files on page 613
- Configuration Files Terms on page 615

### Understanding Configuration Files for EX Series Switches

A configuration file stores the complete configuration of a switch. The current configuration of a switch is called the active configuration. You can alter this current configuration and you can also return to a previous configuration or to a rescue configuration. For more information, see “[Configuration Files Terms](#)” on page 615.

Juniper Networks Junos operating system (Junos OS) saves the 50 most recently committed configuration files on the switch so that you can return to a previous configuration. The configuration files are named:

- **juniper.conf.gz**—The current active configuration.
- **juniper.conf.1.gz** to **juniper.conf.49.gz**—Rollback configurations.

The configuration files are available in the **/config** folder on the switch.

To make changes to the configuration file, you have to work in the configuration mode in the CLI or use the configuration tools in the J-Web interface. When making changes to a configuration file, you are viewing and changing the candidate configuration file. The candidate configuration allows you to make configuration changes without causing operational changes to the active configuration or causing potential damage to your current network operations. Once you commit the changes made to the candidate configuration, the system updates the active configuration.

- Related Documentation**
- [Managing Configuration Files Through the Configuration History \(J-Web Procedure\) on page 641](#)
  - [Uploading a Configuration File \(CLI Procedure\) on page 617](#)
  - [Uploading a Configuration File \(J-Web Procedure\) on page 619](#)
  - [Loading a Previous Configuration File \(CLI Procedure\) on page 619](#)
  - [Reverting to the Rescue Configuration for the EX Series Switch on page 624](#)

- [Configuration Files Terms on page 615](#)

## Using the CLI Viewer in the J-Web Interface to View Configuration Text

---



**NOTE:** This topic applies only to the J-Web Application package.

---

To view the entire configuration file contents in text format, select **Configure > CLI Tools > CLI Viewer**. The main pane displays the configuration in text format.

Each level in the hierarchy is indented to indicate each statement's relative position in the hierarchy. Each level is generally set off with braces, with an open brace ( { ) at the beginning of each hierarchy level and a closing brace ( } ) at the end. If the statement at a hierarchy level is empty, the braces are not displayed. Each leaf statement ends with a semicolon ( ; ), as does the last statement in the hierarchy.

This indented representation is used when the configuration is displayed or saved as an ASCII file. However, when you load an ASCII configuration file, the format of the file is not so strict. The braces and semicolons are required, but the indentation and use of new lines are not required in ASCII configuration files.

### Related Documentation

- [Understanding J-Web Configuration Tools](#)

## Using the CLI Editor in the J-Web Interface to Edit Configuration Text

---



**NOTE:** This topic applies only to the J-Web Application package.

---

Use the CLI Editor to edit configuration if you know the Junos OS CLI or prefer a command interface.

To edit the entire configuration in text format:

---



**CAUTION:** We recommend that you use this method to edit and commit the configuration only if you have experience editing configurations through the CLI.

---

1. Select **Configure > CLI Tools > CLI Editor**. The work area displays the configuration in a text editor.
2. Navigate to the hierarchy level you want to edit.

You can edit the candidate configuration using standard text editor operations—insert lines (by using the Enter key), delete lines, and modify, copy, and paste text.

3. Click **Commit** to load and commit the configuration.

The switching platform checks the configuration for the correct syntax before committing it.

- Related Documentation**
- [CLI User Interface Overview on page 313](#)
  - [Understanding J-Web Configuration Tools](#)

## Using the Point and Click CLI Tool in the J-Web Interface to Edit Configuration Text



**NOTE:** This topic applies only to the J-Web Application package.

To edit the configuration on a series of pages of clickable options that steps you through the hierarchy, select **Configure > CLI Tools > Point&Click CLI**. The side pane displays the top level of the configured hierarchy, and the work area displays configured hierarchy options and the Icon Legend.

To expand or hide the hierarchy of all the statements in the side pane, click **Expand all** or **Hide all**. To expand or hide an individual statement in the hierarchy, click the expand (+) or collapse (–) icon to the left of the statement.



**TIP:** Only those statements included in the committed configuration are displayed in the hierarchy.

The configuration information in the work area consists of configuration options that correspond to configuration statements. Configuration options that contain subordinate statements are identified by the term *Nested*.

To include, edit, or delete statements in the candidate configuration, click one of the links described in [Table 52 on page 611](#). Then specify configuration information by typing in a field, selecting a value from a list, or selecting a check box (toggle).

**Table 52: J-Web Edit Point & Click Configuration Links**

Link	Function
Add new entry	Displays fields and lists for a statement identifier, allowing you to add a new identifier to a statement.
Configure	Displays information for a configuration option that has not been configured, allowing you to include a statement.
Delete	Deletes the corresponding statement or identifier from the configuration. All subordinate statements and identifiers contained within a deleted statement are also discarded.
Edit	Displays information for a configuration option that has already been configured, allowing you to edit a statement.
Identifier	Displays fields and lists for an existing statement identifier, allowing you to edit the identifier.

As you navigate through the configuration, the hierarchy level is displayed at the top of the work area. You can click a statement or identifier in the hierarchy to display the corresponding configuration options in the work area.

The work area includes icons that display information about statements and identifiers when you place your cursor over them. [Table 53 on page 612](#) describes these icons.

**Table 53: J-Web Edit Point & Click Configuration Icons**

Icon	Function
C	Displays a comment about a statement.
I	Indicates that a statement is inactive.
M	Indicates that a statement has been added or modified but has not been committed.
*	Indicates that the statement or identifier is required in the configuration.
?	Provides online help information.

After typing or selecting your configuration edits, click a button in the work area (described in [Table 54 on page 612](#)) to apply your changes or cancel them, refresh the display, or discard parts of the candidate configuration. An updated configuration does not take effect until you commit it.

**Table 54: J-Web Edit Point & Click Configuration Buttons**

Button	Function
Refresh	Updates the display with any changes to the configuration made by other users.
Commit	Verifies edits and applies them to the current configuration file running on the switch.
Discard	Removes edits applied to or deletes existing statements or identifiers from the candidate configuration.

- Related Documentation**
- [CLI User Interface Overview on page 313](#)
  - [Understanding J-Web Configuration Tools](#)

## Understanding Automatic Refreshing of Scripts on EX Series Switches

You can automatically refresh **commit**, **event**, and **op** scripts using operational mode commands on EX Series switches. The commands are:

- [request system scripts refresh-from commit](#)
- [request system scripts refresh-from event](#)
- [request system scripts refresh-from op](#)

The existing Junos operating system (Junos OS) command-line interface (CLI) **refresh** and **refresh-from** configuration mode statements have been extended to work with Junos XML management protocol and NETCONF XML management protocol sessions.

### Related Documentation

- [Understanding Autoinstallation of Configuration Files on page 613](#)
- [CLI User Interface Overview on page 313](#)
- [Junos OS Junos XML Management Protocol Guide](#)
- [Junos OS NETCONF XML Management Protocol Guide](#)

## Understanding Autoinstallation of Configuration Files

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

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

This topic describes:

- [Typical Uses for Autoinstallation on page 613](#)
- [Autoinstallation Configuration Files and IP Addresses on page 614](#)
- [Typical Autoinstallation Process on a New Switch on page 614](#)

### Typical Uses for Autoinstallation

Typical uses for autoinstallation of the software include:

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

## Autoinstallation Configuration Files and IP Addresses

---

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

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

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

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

## Typical Autoinstallation Process on a New Switch

---

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

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

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

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

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

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

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

**Related  
Documentation**

- [Configuring Autoinstallation of Configuration Files \(CLI Procedure\) on page 626](#)
- [Connecting and Configuring an EX Series Switch \(CLI Procedure\)](#)
- [Connecting and Configuring an EX Series Switch \(J-Web Procedure\)](#)
- [Configuration Files Terms on page 615](#)

## Configuration Files Terms

[Table 55 on page 615](#) lists the various configuration file terms used for EX Series switches and their definitions.

**Table 55: Configuration File Terms**

Term	Definition
active configuration	The current committed configuration of a switch.
candidate configuration	A working copy of the configuration that allows users to make configurational changes without causing any operational changes until this copy is committed.
configuration group	Group of configuration statements that can be inherited by the rest of the configuration.

Table 55: Configuration File Terms (*continued*)

Term	Definition
commit a configuration	Have the candidate configuration checked for proper syntax, activated, and marked as the current configuration file running on the switching platform.
configuration hierarchy	The Junos OS configuration consists of a hierarchy of statements. There are two types of statements: container statements, which contain other statements, and leaf statements, which do not contain other statements. All the container and leaf statements together form the configuration hierarchy.
default configuration	The default configuration contains the initial values set for each configuration parameter when a switch is shipped.
rescue configuration	Well-known configuration that recovers a switch from a configuration that denies management access. You set a current committed configuration to be the rescue configuration through the J-Web interface or CLI.
roll back a configuration	Return to a previously committed configuration.

**Related  
Documentation**

- [EX2200 Switch Default Configuration](#)
- [EX3200 Default Configuration](#)
- [EX4200 Default Configuration](#)
- [EX4500 Default Configuration](#)
- [EX8200 Switch Default Configuration](#)
- [Loading a Previous Configuration File \(CLI Procedure\) on page 619](#)
- [Managing Configuration Files Through the Configuration History \(J-Web Procedure\) on page 641](#)
- [Reverting to the Rescue Configuration for the EX Series Switch on page 624](#)
- [Understanding Configuration Files for EX Series Switches on page 609](#)



## CHAPTER 14

# Configuration

- Configuration Tasks on page 617
- Configuration Statements on page 629

### Configuration Tasks

---

- Uploading a Configuration File (CLI Procedure) on page 617
- Uploading a Configuration File (J-Web Procedure) on page 619
- Loading a Previous Configuration File (CLI Procedure) on page 619
- Reverting to the Default Factory Configuration for the EX Series Switch on page 620
- Reverting to the Rescue Configuration for the EX Series Switch on page 624
- Setting or Deleting the Rescue Configuration (CLI Procedure) on page 625
- Setting or Deleting the Rescue Configuration (J-Web Procedure) on page 625
- Configuring Autoinstallation of Configuration Files (CLI Procedure) on page 626
- Using the Commit Options to Commit Configuration Changes (J-Web Procedure) on page 628

### Uploading a Configuration File (CLI Procedure)

You can create a configuration file on your local system, copy the file to the EX Series switch and then load the file into the CLI. After you have loaded the configuration file, you can commit it to activate the configuration on the switch. You can also edit the configuration interactively using the CLI and commit it at a later time.

To upload a configuration file from your local system:

1. Create the configuration file using a text editor such as Notepad, making sure that the syntax of the configuration file is correct. For more information about testing the syntax of a configuration file see the *Junos OS System Basics and Services Command Reference*.
2. In the configuration text file, use an option to perform the required action when the file is loaded. [Table 56 on page 618](#) lists and describes some options for the **load** command.

Table 56: Options for the load command

Options	Description
<b>merge</b>	Combines the current active configuration and the configuration in <i>filename</i> or the one that you type at the terminal. A <b>merge</b> operation is useful when you are adding a new section to an existing configuration. If the active configuration and the incoming configuration contain conflicting statements, the statements in the incoming configuration override those in the active configuration.
<b>override</b>	Discards the current candidate configuration and loads the configuration in <i>filename</i> or the one that you type at the terminal. When you use the <b>override</b> option and commit the configuration, all system processes reparse the configuration. You can use the <b>override</b> option at any level of the hierarchy.
<b>replace</b>	Searches for the <b>replace</b> tags, deletes the existing statements of the same name, if any, and replaces them with the incoming configuration. If there is no existing statement of the same name, the <b>replace</b> operation adds the statements marked with the <b>replace</b> tag to the active configuration.  <b>NOTE:</b> For this operation to work, you must include <b>replace</b> tags in the text file or in the configuration you type at the terminal.

- Press Ctrl+A to select all the text in the configuration file.
- Press Ctrl+C to copy the contents of the configuration text file to the Clipboard.
- Log in to the switch using your username and password.
- To enter configuration mode:  

```
user@switch> configure
```

You will see this output, with the hash or pound mark indicating configuration mode.

**Entering configuration mode**

```
[edit]
user@switch#
```
- Load the configuration file:  

```
[edit]
user@switch# load merge terminal
```
- At the cursor, paste the contents of the Clipboard using the mouse and the Paste icon:  

```
[edit]
user@switch# load merge terminal
[Type ^D at a new line to end input]
>Cursor is here. Paste the contents of the clipboard here<
```
- Press Enter.
- Press Ctrl+D to set the end-of-file marker.

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt. You can also edit the configuration interactively using the CLI and commit it at a later time.

- Related Documentation**
- [Uploading a Configuration File \(J-Web Procedure\) on page 619](#)
  - [Understanding Configuration Files for EX Series Switches on page 609](#)

## Uploading a Configuration File (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You can create a configuration file on your local system, copy the file to the EX Series switch and then load the file into the CLI. After you have loaded the configuration file, you can commit it to activate the configuration on the switch. You can also edit the configuration interactively using the CLI and commit it at a later time.

To upload a configuration file from your local system:

1. Select **Maintain > Config Management > Upload**.  
The work area displays the File to Upload box.
2. Specify the name of the file to upload using one of the following methods:
  - Type the absolute path and filename in the File to Upload box.
  - Click **Browse** to navigate to the file.
3. Click **Upload and Commit** to upload and commit the configuration.  
The switch checks the configuration for the correct syntax before committing it.

- Related Documentation**
- [Uploading a Configuration File \(CLI Procedure\) on page 617](#)
  - [Understanding J-Web Configuration Tools](#)
  - [Understanding Configuration Files for EX Series Switches on page 609](#)

## Loading a Previous Configuration File (CLI Procedure)

You can return to a previously committed configuration file if you need to revert to a previous configuration. The EX Series switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

### Syntax

**rollback** <number>

### Options

- **none**—Return to the most recently saved configuration.
- **number**—Configuration to return to.
  - **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.

- **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
user@switch# commit
```

#### Related Documentation

- [Managing Configuration Files Through the Configuration History \(J-Web Procedure\) on page 641](#)
- [Configuration Files Terms on page 615](#)
- [For more information on rollback, see \*Junos OS CLI User Guide\*.](#)

## Reverting to the Default Factory Configuration for the EX Series Switch

If for any reason the current active configuration fails, you can revert to the factory-default configuration.

You can also roll back to a previous configuration, as described in “[Loading a Previous Configuration File \(CLI Procedure\)](#)” on page 619, or revert to the rescue configuration, as described in “[Reverting to the Rescue Configuration for the EX Series Switch](#)” on page 624.



**TIP:** If you have lost the root password, it is not necessary to revert to the factory-default configuration to reset it. See “[Troubleshooting Loss of the Root Password](#)” on page 305.

The factory-default configuration contains the basic configuration settings for the switch. This is the first configuration of the switch and it is loaded when the switch is first powered on. For the factory-default configuration file for your switch, see the complete list under the Configuration tab of *Configuration File Management on EX Series Switches*.



**TIP:** You can run the EZsetup script to complete the initial configuration of the switch *after* reverting to the factory-default configuration. (The EZsetup script is available only on fixed configuration switches, it is not available on modular switches.) For information on completing the initial configuration using either the CLI or the J-Web interface, see *Connecting and Configuring an EX Series Switch (CLI Procedure)* or *Connecting and Configuring an EX Series Switch (J-Web Procedure)*.

You can revert to the factory-default configuration by using the **Menu** button to the right of the LCD panel or by using the **request system zeroize** operational command or the **load factory-default** configuration command. (If your switch model does not have an LCD panel, use these commands.) You can also use the **load factory-default** command to revert to the factory-default configuration file that contains all default settings *except* the root password setting, which is retained.


These procedures are described in the following sections:

- [Reverting to the Factory-Default Configuration by Using the LCD Panel on page 621](#)
- [Reverting to the Factory-Default Configuration by Using the request system zeroize Command on page 622](#)
- [Reverting to the Factory-Default Configuration by Using the load factory-default Command on page 622](#)

### Reverting to the Factory-Default Configuration by Using the LCD Panel

To set the switch to the factory-default configuration, use the LCD panel and buttons on the front panel of the switch shown in [Figure 3 on page 621](#). If the switch model does not have an LCD panel, use one of the CLI commands described in the following sections.

**Figure 3: EX Series Switch LCD Panel**



**NOTE:** To revert a member switch of a Virtual Chassis to the factory-default configuration, first disconnect the cables connected to the Virtual Chassis ports (VCPs) to avoid affecting Virtual Chassis configuration parameters (member ID, mastership priority, and setting of VCP uplinks) on other members. See *Disconnecting a Fiber-Optic Cable from an EX Series Switch*, *Disconnecting a Virtual Chassis Cable from an EX4200 Switch*, or *Disconnecting a Virtual Chassis Cable from an EX4500 Switch*.

To revert to the factory-default configuration by using the LCD panel:

1. Press the **Menu** button until you see MAINTENANCE MENU on the panel.
2. Press the **Enter** button.
3. Press **Menu** until you see FACTORY DEFAULT.
4. Press **Enter**. The display says RESTORE DEFAULT?

5. Press **Enter**. The screen flashes **FACTORY DEFAULT IN PROGRESS** and returns to the idle menu.
6. Complete the initial configuration of the switch. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* or *Connecting and Configuring an EX Series Switch (J-Web Procedure)*

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

---

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

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



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

---

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

1. `user@switch> request system zeroize`  
warning: System will be rebooted and may not boot without configuration  
Erase all data, including configuration and log files? [yes,no] (yes)
  2. Type **yes** to remove configuration and log files and revert to the factory default configuration.
- 



**NOTE:** The **auto-image-upgrade** statement is added under the **[edit chassis]** hierarchy level when you use this procedure, and thus the automatic image upgrade feature is made available on the switch.

---

3. Complete the initial configuration of the switch. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* or *Connecting and Configuring an EX Series Switch (J-Web Procedure)*

### Reverting to the Factory-Default Configuration by Using the `load factory-default` Command

---

The **load factory-default** command is a standard Junos OS configuration command that replaces the current active configuration with the factory-default configuration (except the root password setting, which by default is not set but which you must set in order to commit the new configuration in this procedure).

If you want to run the EZsetup script to complete the initial configuration of the switch after you revert to the factory-default configuration, do not use the **load factory-default**

command. Instead do the reversion using either the LCD panel or the **request system zeroize** command. If you use the **load factory-default** command to revert to the factory-default configuration, the configuration for the root password is retained and the EZsetup script will not run. (The EZsetup script is available only on fixed configuration switches, it is not available on modular switches.)



**NOTE:** The **load factory-default** command by itself is not supported on EX3300, EX4200, EX4500, and EX4550 switches configured in a Virtual Chassis.

To revert to the factory-default configuration by using the **load factory-default** command:



**NOTE:** If you use this procedure, you must delete the system commit factory settings, set the root password, and commit the configuration. These steps are not required when you revert to the factory-default configuration by using **request system zeroize**. Also, the **auto-image-upgrade** statement is not added to the configuration when you use this procedure; it *is* added to the configuration when you use **request system zeroize**.

1. [edit]  
user@switch# **load factory-default**
2. [edit]  
user@switch# **delete system commit factory-settings**
3. [edit]  
user@switch# **set system root-authentication plain-text-password**
4. [edit]  
user@switch# **commit**
5. Check the member ID and mastership priority with the **show virtual-chassis** command and check to see whether there are remaining settings for uplink VCPs by using the **show virtual-chassis vc-port** command.

#### Related Documentation

- [Connecting and Configuring an EX Series Switch \(CLI Procedure\)](#)
- [Connecting and Configuring an EX Series Switch \(J-Web Procedure\)](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Understanding Configuration Files for EX Series Switches on page 609](#)
- [For more information about the load factory-default command, see Junos OS CLI User Guide.](#)

## Reverting to the Rescue Configuration for the EX Series Switch

If someone inadvertently commits a configuration that denies management access to an EX Series switch and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration by using the LCD panel on the switch. The rescue configuration is a previously committed, valid configuration.

You can also revert to the default factory configuration, as described in [“Reverting to the Default Factory Configuration for the EX Series Switch” on page 620](#).

Before you begin to revert to the rescue configuration:

- Ensure that you have physical access to the switch.
- A rescue configuration for the switch must have been previously set.

To revert the switch to the rescue configuration:

1. At the LCD panel on the switch, press **Menu** until you see **MAINTENANCE MENU**.
2. Press **Enter**.
3. Press **Menu** until you see **Load Rescue**.
4. Press **Enter**.
5. When **Commit Rescue** is displayed, press **Enter**.

The LCD panel displays the message **Commit Rescue in Progress**. When the reversion is complete, it displays the idle menu.



.....  
**NOTE:** If there is no rescue configuration saved on the switch, the message **Commit rescue failed** is displayed.  
.....

### Related Documentation

- [Setting or Deleting the Rescue Configuration \(CLI Procedure\) on page 625](#)
- [Setting or Deleting the Rescue Configuration \(J-Web Procedure\) on page 625](#)
- [LCD Panel in EX3200 Switches](#)
- [LCD Panel in EX4200 Switches](#)
- [LCD Panel in EX4500 Switches](#)
- [LCD Panel in an EX8200 Switch](#)
- [Configuration Files Terms on page 615](#)



## Setting or Deleting the Rescue Configuration (CLI Procedure)

A rescue configuration is a well-known configuration that recovers a switch from a configuration that denies management access. You set a current committed configuration to be the rescue configuration through the J-Web interface or CLI.

If someone inadvertently commits a configuration that denies management access to an EX Series switch and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration by using the LCD panel on the switch. The rescue configuration is a previously committed, valid configuration. We recommend that the rescue configuration include the IP address (accessible from the network) for the management port.

To set the current active configuration as the rescue configuration:

```
user@switch> request system configuration rescue save
```

To delete an existing rescue configuration:

```
user@switch> request system configuration rescue delete
```

### Related Documentation

- [Setting or Deleting the Rescue Configuration \(J-Web Procedure\) on page 625](#)
- [Reverting to the Rescue Configuration for the EX Series Switch on page 624](#)
- [Loading a Previous Configuration File \(CLI Procedure\) on page 619](#)
- [Configuration Files Terms on page 615](#)
- [For information on show system configuration rescue, see \*Junos OS System Basics and Services Command Reference\*.](#)

## Setting or Deleting the Rescue Configuration (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

A rescue configuration is a well-known configuration that recovers a switch from a configuration that denies management access. You set a current committed configuration to be the rescue configuration through the J-Web interface or CLI.

If someone inadvertently commits a configuration that denies management access to an EX Series switch and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration by using the LCD panel on the switch. The rescue configuration is a previously committed, valid configuration. We recommend that the rescue configuration include the IP address (accessible from the network) for the management port.

To view, set, or delete the rescue configuration using the J-Web interface, select **Maintain > Config Management > Rescue**. On the Rescue page, you can perform the following tasks:

- View the current rescue configuration—Click **View rescue configuration**.

- Set the current running configuration as the rescue configuration—Click **Set rescue configuration**.
- Delete the current rescue configuration—Click **Delete rescue configuration**.

**Related  
Documentation**

- [Setting or Deleting the Rescue Configuration \(CLI Procedure\) on page 625](#)
- [Reverting to the Rescue Configuration for the EX Series Switch on page 624](#)
- [Configuration Files Terms on page 615](#)

## Configuring Autoinstallation of Configuration Files (CLI Procedure)

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

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

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

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

the IP addresses of the hosts providing TFTP and DNS services. Connect this interface to the switch.

- If you are using **hostname.conf** files for autoinstallation, you must also complete the following tasks:
  - Configure the DHCP server to provide a **hostname.conf** filename to each switch. Each switch uses its **hostname.conf** filename to request a configuration file from the TFTP server. Copy the necessary **hostname.conf** configuration files to the TFTP server.
  - Create a default configuration file named **network.conf**, and copy it to the TFTP server. This file contains IP-address-to-hostname mapping entries. If the DHCP server does not send a **hostname.conf** filename to a new switch, the switch uses **network.conf** to resolve its hostname based on its IP address.

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

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

To configure autoinstallation:

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

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



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

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

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

#### Related Documentation

- [Verifying Autoinstallation Status on page 644](#)
- [Understanding Autoinstallation of Configuration Files on page 613](#)
- [Understanding DHCP Services for Switches](#)

## Using the Commit Options to Commit Configuration Changes (J-Web Procedure)

You can use the single-commit feature to commit all outstanding configuration changes in the J-Web interface on EX Series switches simultaneously. This helps in reducing the time J-Web takes for committing configurations because when changes are committed at every step, rollback configurations pile up.

For example, suppose you want to delete a firewall filter and add a new one. With immediate commits, you would need to commit your changes twice for this action. Using single commit, you can decrease the number of commits to one, thus saving time for working on other configurations.

When you edit a configuration, you work on a copy of the current configuration, which is your candidate configuration. The changes you make to the candidate configuration are visible through the user interface immediately, allowing other users to edit those configurations, but they do not take effect on the switch until you commit the changes. When you commit the configuration, the candidate file is checked for proper syntax, activated, and marked as the current, operational software configuration file. If multiple users are editing the configuration when you commit the candidate configuration, changes made by all users take effect.

You can configure the commit options to either commit all configuration changes together or commit each configuration change immediately using the J-Web Commit Preference page.



**NOTE:** There are some pages on which configuration changes must be committed immediately. For such pages, if you configure the commit options for a single commit, the system displays warning notifications that remind you to commit your changes immediately. An example of such a page is the Interface Page (Configure > Interface).

To configure the commit options on an EX Series switch using the J-Web interface:

1. Select **Commit Options**.



**NOTE:** All action links except **Preference** are disabled unless you edit, add, or delete a configuration.

2. Choose an action. See [Table 57 on page 629](#) for details on the actions.
3. Configure the commit options by selecting **Preference**. See [Table 58 on page 629](#) for details on preference options.

Table 57: Commit Options

Menu Item	Function	Your Action
Commit	Commits the candidate configuration of the current user session, along with changes from other user sessions.	<ol style="list-style-type: none"> <li>1. Select <b>Commit Options &gt; Commit</b>.  Changes are committed after the system validates your configuration. A window displays that the configuration was successfully committed or that the commit failed.</li> <li>2. Click <b>OK</b>.  Click <b>Details</b> to view the commit log.</li> </ol>
Compare	Displays the XML log of pending uncommitted configurations on the device.	<ol style="list-style-type: none"> <li>1. Select <b>Commit Options &gt; Compare</b>.  The XML log of pending configurations on the devices are displayed similar to the CLI interface, in a “human-readable” form.</li> <li>2. Click <b>Close</b>.</li> </ol>
Discard	Discards the candidate configuration of your current session, along with changes from other user sessions.	<ol style="list-style-type: none"> <li>1. Select <b>Commit Options &gt; Discard</b>.</li> <li>2. Click <b>OK</b> to confirm the discard action.  Your changes are discarded after the system validates your configuration.</li> </ol>
Preference	Indicates your choice of committing all global configurations together or committing each configuration change immediately.	<ol style="list-style-type: none"> <li>1. Select <b>Commit Options &gt; Preference</b>. The Commit Preference page is displayed.</li> <li>2. Configure the commit options by selecting your preference. See <a href="#">Table 58 on page 629</a> for details on preference options.</li> </ol>

Table 58: Commit Preference Options

Option	Function
Validate and commit configuration changes	Sets the system to validate and force an immediate commit on every screen after every configuration change.
Validate configuration changes	<p>Loads all the configuration changes for an accumulated single commit. If there are errors in loading the configuration, the errors are logged. This is the default mode.</p> <p>Once you select this option, you need to select <b>Commit Options &gt; Commit</b> to commit your changes.</p>

- Related Documentation**
- *J-Web User Interface for EX Series Switches Overview*
  - *EX Series Switch Software Features Overview*

## Configuration Statements

- [archival on page 630](#)
- [archive-sites \(Configuration File\) on page 631](#)
- [autoinstallation on page 633](#)

- [synchronize \(Commit configuration\) on page 634](#)
- [configuration on page 635](#)
- [configuration-servers on page 636](#)
- [interfaces on page 637](#)
- [transfer-interval \(Configuration\) on page 638](#)
- [transfer-on-commit on page 639](#)

---

## archival

---

**Syntax**

```
archival {  
  configuration {  
    archive-sites {  
      file://<path>/<filename>;  
      ftp://username@host:<port>url-path password password;  
      http://username@host:<port>url-path password password;  
      pasvftp://username@host:<port>url-path password password;  
      scp://username@host:<port>url-path password password;  
    }  
    transfer-interval interval;  
    transfer-on-commit;  
  }  
}
```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure copying of the currently active configuration to an archive site. An archive site can be a file, or an FTP or SCP location.



**NOTE:** The edit system archival hierarchy is not available on QFabric systems.

---

**Options** The remaining statements are explained separately.



**NOTE:** The [edit system archival] hierarchy is not available on QFabric systems.



---

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- [Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site](#)

## archive-sites (Configuration File)

<b>Syntax</b>	<pre>archive-sites {     file://&lt;path&gt;/&lt;filename&gt;;     ftp://username@host:&lt;port&gt;url-path password password;     http://username@host:&lt;port&gt;url-path password password;     pasvftp://username@host:&lt;port&gt;url-path password password;     scp://username@host:&lt;port&gt;url-path password password; }</pre>
<b>Hierarchy Level</b>	[edit system archival configuration]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Specify where to transfer the current configuration files. When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks ( " ") and enclose the IPv6 host address in brackets ( [ ] ). For example, "scp://username&lt;:password&gt;@[ipv6-host-address]&lt;:port&gt;/url-path"</p> <p>If you specify more than one archive site, the router or switch attempts to transfer the configuration files to the first archive site in the list, moving to the next only if the transfer fails.</p> <p>The destination filename is saved in the following format, where <i>n</i> corresponds to the number of the compressed configuration rollback file that has been archived:</p> <p><b><i>router-name_juniper.conf.n.gz_YYYYMMDD_HHMMSS.</i></b></p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router or switch is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.</p> </div> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> The [edit system archival] hierarchy is not available on QFabric systems.</p> </div>
<b>Options</b>	<p>The prefix used in the configuration statement determines the form of transfer:</p> <p><b>file://</b> —transfer on a path to a named file</p> <p><b>ftp://</b> —transfer using active FTP server</p> <p><b>pasvftp://</b> —transfer to a device that only accepts passive FTP services</p>

**scp://** —transfer to a known host using background SCP file transfers

<b>Required Privilege</b>	system—To view this statement in the configuration.
<b>Level</b>	system-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site</i></li><li>• <i>Junos OS Commit Model for Router or Switch Configuration</i></li><li>• <a href="#">configuration on page 635</a></li><li>• <a href="#">transfer-on-commit on page 639</a></li></ul>
------------------------------	---



## autoinstallation

<b>Syntax</b>	<pre> autoinstallation {   configuration-servers {     url;   }   interfaces {     interface-name {       bootp;       rarp;     }   } } </pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p>
<b>Description</b>	<p>For ACX Series routers, J Series Services routers, and EX Series switches only. Download a configuration file automatically from an FTP, Hypertext Transfer Protocol (HTTP), or Trivial FTP (TFTP) server. When you power on a router or switch configured for autoinstallation, it requests an IP address from a Dynamic Host Configuration Protocol (DHCP) server. Once the router or switch has an address, it sends a request to a configuration server and downloads and installs a configuration.</p>
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>ACX Series Autoinstallation Overview</i></li> <li>• <i>Before You Begin Autoinstallation on an ACX Series Universal Access Router</i></li> <li>• <i>Autoinstallation Configuration of ACX Series Universal Access Routers</i></li> <li>• <i>USB Autoinstallation on ACX Series Routers</i></li> <li>• <i>Verifying Autoinstallation on ACX Series Universal Access Routers</i></li> <li>• <i>show system autoinstallation status</i></li> <li>• <i>Upgrading Software by Using Automatic Software Download</i></li> <li>• <i>J Series Services Router Basic LAN and WAN Access Configuration Guide</i></li> <li>• <a href="#">configuration-servers on page 636</a></li> <li>• <a href="#">idle-timeout on page 552</a></li> </ul>

## synchronize (Commit configuration)

---

<b>Syntax</b>	synchronize;
<b>Hierarchy Level</b>	[edit system commit]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	<p>For devices with multiple Routing Engines only. Configure a <b>commit</b> command to automatically result in a <b>commit synchronize</b> action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the <b>commit</b> command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engines. All Routing Engines then perform a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on all Routing Engines.</p> <p>Starting with Junos OS Release 9.3, accounting of events and operations on a backup Routing Engine is not supported on accounting servers such as TACACS+ or RADIUS. Logging of accounting events is supported only for events and operations on a master Routing Engine.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically</i></li></ul>

## configuration

**Syntax**

```
configuration {
  transfer-interval interval;
  transfer-on-commit;
  archive-sites {
    file://<path>/<filename>;
    ftp://username@host:<port>url-path password password;
    http://username@host:<port>url-path password password;
    pasvftp://username@host:<port>url-path password password;
    scp://username@host:<port>url-path password password;
  }
}
```

**Hierarchy Level** [edit system archival]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the router or switch to periodically transfer its currently active configuration (or after each commit).



**NOTE:** The [edit system archival] hierarchy is not available on QFabric systems.

**Options** The remaining statements are explained separately.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- *Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site*
- *archive*
- [archive-sites on page 631](#)
- [transfer-interval on page 638](#)
- [transfer-on-commit on page 639](#)

## configuration-servers

---

<b>Syntax</b>	<pre>configuration-servers {     url; }</pre>
<b>Hierarchy Level</b>	[edit system <a href="#">autoinstallation</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	For J Series Services Routers and EX Series switches only, configure the URL address of a server from which to obtain configuration files. Examples of URLs:  <b><i>tftp://hostname/path/filename</i></b>  <b><i>ftp://username:prompt@ftp.hostname.net/filename /</i></b>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Upgrading Software by Using Automatic Software Download</i></li><li>• Getting Started Guide for your router model</li><li>• <a href="#">autoinstallation on page 633</a></li><li>• <a href="#">idle-timeout on page 552</a></li></ul>



## interfaces

---

<b>Syntax</b>	<pre> interfaces {   interface-name {     bootp;     rarp;     slarp;   } } </pre>
<b>Hierarchy Level</b>	[edit system <a href="#">autoinstallation</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	For J Series Services Routers and EX Series switches only. Configure the interface on which to perform autoinstallation. A request for an IP address is sent from the interface. Specify the IP address procurement protocol.
<b>Options</b>	<p><b>rarpbootp</b>—Send requests over serial interfaces with Frame Relay.</p> <p><b>rarp</b>—Send requests over Ethernet interfaces.</p> <p><b>slarp</b>—(On J Series Services Routers only) Send requests over serial interfaces with HDLC.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Upgrading Software by Using Automatic Software Download</i></li> <li>• <i>J Series Services Router Basic LAN and WAN Access Configuration Guide</i></li> <li>• <a href="#">autoinstallation on page 633</a></li> </ul>

## transfer-interval (Configuration)

---

<b>Syntax</b>	<code>transfer-interval <i>interval</i>;</code>
<b>Hierarchy Level</b>	[edit system archival configuration]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the router or switch to periodically transfer its currently active configuration to an archive site.
<hr/>	
<div> <b>NOTE:</b> The <code>edit system archival</code> hierarchy is not available on QFabric systems.</div> <hr/>	
<b>Options</b>	<b><i>interval</i></b> —Interval at which to transfer the current configuration to an archive site. <b>Range:</b> 15 through 2880 minutes
<hr/>	
<div> <b>NOTE:</b> The <code>[edit system archival]</code> hierarchy is not available on QFabric systems.</div> <hr/>	
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site</i></li><li>• <i>archive</i></li><li>• <a href="#">configuration on page 635</a></li><li>• <a href="#">transfer-on-commit on page 639</a></li></ul>

## transfer-on-commit

<b>Syntax</b>	transfer-on-commit;
<b>Hierarchy Level</b>	[edit system archival configuration]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the router or switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration.



**NOTE:** When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks ( " ") and enclose the IPv6 host address in brackets ( [ ] ). For example, "ftp://username<:password>@[ipv6-host-address]<:port>/url-path" .



**NOTE:** The [edit system archival] hierarchy is not available on QFabric systems.

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site</i></li> <li>• <i>archive</i></li> <li>• <a href="#">configuration on page 635</a></li> <li>• <a href="#">transfer-interval on page 638</a></li> </ul>





# Administration

- [Routine Monitoring on page 641](#)
- [Operational Commands on page 645](#)

## Routine Monitoring

---

- [Managing Configuration Files Through the Configuration History \(J-Web Procedure\) on page 641](#)
- [Verifying Autoinstallation Status on page 644](#)

## Managing Configuration Files Through the Configuration History (J-Web Procedure)

---



NOTE: This topic applies only to the J-Web Application package.

---

Use the Configuration History function to manage configuration files.

1. [Displaying Configuration History on page 641](#)
2. [Displaying Users Editing the Configuration on page 642](#)
3. [Comparing Configuration Files with the J-Web Interface on page 643](#)
4. [Downloading a Configuration File with the J-Web Interface on page 643](#)
5. [Loading a Previous Configuration File with the J-Web Interface on page 643](#)

### Displaying Configuration History

---

To manage configuration files with the J-Web interface, select **Maintain > Config Management > History**. The main pane displays History — Database Information page.

[Table 59 on page 642](#) summarizes the contents of the display.

The configuration history display allows you to:

- View a configuration.
- Compare two configurations.
- Download a configuration file to your local system.
- Roll back the configuration to any of the previous versions stored on the switch.

Table 59: J-Web Configuration History Summary

Field	Description
Number	Version of the configuration file.
Date/Time	Date and time the configuration was committed.
User	Name of the user who committed the configuration.
Client	Method by which the configuration was committed: <ul style="list-style-type: none"> <li><b>cli</b>—A user entered a Junos OS CLI command.</li> <li><b>junoscript</b>—A Junos XML protocol client performed the operation. Commit operations performed by users through the J-Web interface are identified in this way.</li> <li><b>snmp</b>—An SNMP <b>set</b> request started the operation.</li> <li><b>other</b>—Another method was used to commit the configuration.</li> </ul>
Comment	Comment.
Log Message	Method used to edit the configuration: <ul style="list-style-type: none"> <li>Imported via paste— Configuration was edited and loaded with the <b>Configure &gt; CLI Tools &gt; Edit Configuration Text</b> option.</li> <li>Imported upload [<i>filename</i>]<b>—</b>Configuration was uploaded with the <b>Configure &gt; CLI Tools &gt; Point Click Editor</b> option.</li> <li>Modified via J-Web Configure — Configuration was modified with the J-Web Configure menu.</li> <li>Rolled back via <i>user-interface</i><b>—</b> Configuration was rolled back to a previous version through the user interface specified by <i>user-interface</i>, which can be Web Interface or CLI.</li> </ul>
Action	Action to perform with the configuration file. The action can be <b>Download</b> or <b>Rollback</b> .

### Displaying Users Editing the Configuration

To display a list of users editing the switching platform configuration, select **Config Management > History**. The list is displayed as Database Information in the main pane. [Table 60 on page 642](#) summarizes the Database Information display.

Table 60: J-Web Configuration Database Information Summary

Field	Description
User Name	Name of user editing the configuration.
Start Time	Time of day the user logged in to the switch.
Idle Time	Elapsed time since the user issued a configuration command from the CLI.
Terminal	Terminal on which the user is logged in.
PID	Process identifier assigned to the user by the switching platform.
Edit Flags	Designates a private or exclusive edit.

**Table 60: J-Web Configuration Database Information Summary** *(continued)*

Field	Description
Edit Path	Level of the configuration hierarchy that the user is editing.

### Comparing Configuration Files with the J-Web Interface

To compare any two of the past 50 committed configuration files:

1. Select **Config Management > History**. A list of the current and the previous 49 configurations is displayed as Configuration History in the main pane.
2. Select the check boxes to the left of the two configuration versions you want to compare.
3. Click **Compare**.

The main pane displays the differences between the two configuration files at each hierarchy level as follows:

- Lines that have changed are highlighted side by side in green.
- Lines that exist only in the more recent configuration file are displayed in red on the left.
- Lines that exist only in the older configuration file are displayed in blue on the right.

### Downloading a Configuration File with the J-Web Interface

To download a configuration file from the switch to your local system:

1. Select **Config Management > History**. A list of current and previous 49 configurations is displayed as Configuration History in the main pane.
2. In the Action column, click **Download** for the version of the configuration you want to download.
3. Select the options your Web browser provides that allow you to save the configuration file to a target directory on your local system.

The file is saved as an ASCII file.

### Loading a Previous Configuration File with the J-Web Interface

To load (roll back) and commit a previous configuration file stored on the switching platform:

1. Select **Config Management > History**. A list of current and previous 49 configurations is displayed as Configuration History in the main pane.
2. In the Action column, click **Rollback** for the version of the configuration you want to load.

The main pane displays the results of the rollback operation.



**NOTE:** When you click **Rollback**, the switch loads and commits the selected configuration. This behavior is different from the switch's behavior that occurs after you enter the **rollback** configuration mode command from the CLI. In the latter case, the configuration is loaded but not committed.

- Related Documentation**
- [Loading a Previous Configuration File \(CLI Procedure\) on page 619](#)
  - [Understanding Configuration Files for EX Series Switches on page 609](#)
  - [Understanding J-Web Configuration Tools](#)

## Verifying Autoinstallation Status

**Purpose** Display the status of the autoinstallation feature.

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

## Sample Output

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

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

- Related Documentation**
- [Configuring Autoinstallation of Configuration Files \(CLI Procedure\) on page 626](#)

## Operational Commands

---

- clear log
- clear system commit
- file archive
- file checksum md5
- file checksum sha1
- file checksum sha-256
- file compare
- file copy
- file delete
- file list
- file rename
- file show
- request system configuration rescue delete
- request system configuration rescue save
- request system scripts refresh-from commit
- request system scripts refresh-from event
- request system scripts refresh-from op
- request system zeroize
- show system commit
- show system configuration archival
- show system configuration rescue
- show system rollback
- test configuration

## clear log

---

<b>Syntax</b>	<code>clear log <i>filename</i></code> <code>&lt;all&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Remove contents of a log file.
<b>Options</b>	<i>filename</i> —Name of the specific log file to delete.  <code>all</code> —(Optional) Delete the specified log file and all archived versions of it.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show log on page 1249</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear log on page 646</a>
<b>Output Fields</b>	See <a href="#">file list</a> for an explanation of output fields.

## Sample Output

### clear log

The following sample commands list log file information, clear the contents of a log file, and then display the updated log file information:

```
user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel          26450 Jun 23 18:47 /var/log/sampled
total 1

user@host> clear log lcc0-re0:sampled
lcc0-re0:
-----

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel           57 Sep 15 03:44 /var/log/sampled
total 1
```

## clear system commit

---

<b>Syntax</b>	clear system commit
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Clear any pending commit operation.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	maintenance (or the actual user who scheduled the commit)
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show system commit on page 675</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear system commit on page 647</a> <a href="#">clear system commit (None Pending) on page 647</a> <a href="#">clear system commit (User Does Not Have Required Privilege Level) on page 647</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### clear system commit

```
user@host> clear system commit
Pending commit cleared.
```

#### clear system commit (None Pending)

```
user@host> clear system commit
No commit scheduled.
```

#### clear system commit (User Does Not Have Required Privilege Level)

```
user@host> clear system commit
error: Permission denied
```

## file archive

---

<b>Syntax</b>	<code>file archive destination <i>destination</i> source <i>source</i> &lt;compress&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.
<b>Options</b>	<p><b>destination <i>destination</i></b>—Destination of the archived file or files. Specify the destination as a URL or filename. The Junos OS adds one of the following suffixes if the destination filename does not already have it:</p> <ul style="list-style-type: none"><li>• For archived files—The suffix <b>.tar</b></li><li>• For archived and compressed files—The suffix <b>.tgz</b></li></ul> <p><b>source <i>source</i></b>—Source of the original file or files. Specify the source as a URL or filename.</p> <p><b>compress</b>—(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix <b>.tgz</b>.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Format for Specifying Filenames and URLs in Junos OS CLI Commands</i></li></ul>
<b>List of Sample Output</b>	<a href="#">file archive (Multiple Files) on page 648</a> <a href="#">file archive (Single File) on page 648</a> <a href="#">file archive (with Compression) on page 649</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file archive (Multiple Files)

The following sample command archives all message files in the local directory `/var/log/messages` as the single file **messages-archive.tar**.

```
user@host> file archive source /var/log/messages* destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

### file archive (Single File)

The following sample command archives one message file in the local directory `/var/log/messages` as the single file **messages-archive.tar**.



```
user@host> file archive source /var/log/messages destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host
```

### file archive (with Compression)

The following sample command archives and compresses all message files in the local directory **/var/log/messages** as the single file **messages-archive.tgz**.

```
user@host> file archive compress source /var/log/messages* destination
/var/log/messages-archive.tgz
/usr/bin/tar: Removing leading / from absolute path names in the archive.
```

## file checksum md5

---

<b>Syntax</b>	<code>file checksum md5 &lt;pathname&gt; filename</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Calculate the Message Digest 5 (MD5) checksum of a file.
<b>Options</b>	<b>pathname</b> —(Optional) Path to a filename. <b>filename</b> —Name of a local file for which to calculate the MD5 checksum.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <a href="#">file checksum sha-256 on page 652</a></li><li>• <a href="#">file checksum sha1 on page 651</a></li><li>• <a href="#">op on page 90</a></li></ul>
<b>List of Sample Output</b>	<a href="#">file checksum md5 on page 650</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = 2a3b69e43f9bd4893729cc16f505a0f5
```

## file checksum sha1

<b>Syntax</b>	<code>file checksum sha1 &lt;pathname&gt; filename</code>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.
<b>Options</b>	<p><b>pathname</b>—(Optional) Path to a filename.</p> <p><b>filename</b>—Name of a local file for which to calculate the SHA-1 checksum.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <a href="#">file checksum md5 on page 650</a></li> <li>• <a href="#">file checksum sha-256 on page 652</a></li> <li>• <a href="#">op on page 90</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">file checksum sha1 on page 651</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax
```

```
SHA1 (/var/db/scripts/commitscript.slax) = ba9e47120c7ce55cff29afd73eacd370e162c676
```

## file checksum sha-256

---

<b>Syntax</b>	<code>file checksum sha-256 &lt;pathname&gt; filename</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.
<b>Options</b>	<b>pathname</b> —(Optional) Path to a filename.  <b>filename</b> —Name of a local file for which to calculate the SHA-256 checksum.
<b>Required Privilege Level</b>	maintenance view view-configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <a href="#">file checksum md5 on page 650</a></li><li>• <a href="#">file checksum sha1 on page 651</a></li><li>• <a href="#">op on page 90</a></li></ul>
<b>List of Sample Output</b>	<a href="#">file checksum sha-256 on page 652</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### file checksum sha-256

```
user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =
94c2b061fb55399e15babd2529453815601a602b5c98e5c12ed929c9d343dd71
```

## file compare

<b>Syntax</b>	<pre>file compare (files <i>filename filename</i>) &lt;context   unified&gt; &lt;ignore-white-space&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Compare two local files and describe the differences between them in default, context, or unified output styles:</p> <ul style="list-style-type: none"> <li>• <b>Default</b>—In the first line of output, <b>c</b> means lines were changed between the two files, <b>d</b> means lines were deleted between the two files, and <b>a</b> means lines were added between the two files. The numbers preceding this alphabetical marker represent the first file, and the lines after the alphabetical marker represent the second file. A left angle bracket (&lt;) in front of output lines refers to the first file. A right angle bracket (&gt;) in front of output lines refers to the second file.</li> <li>• <b>Context</b>—The display is divided into two parts. The first part is the first file; the second part is the second file. Output lines preceded by an exclamation point (!) have changed. Additions are marked with a plus sign (+), and deletions are marked with a minus sign (-).</li> <li>• <b>Unified</b>—The display is preceded by the line number from the first and the second file (xx,xxx,x). Before the line number, additions to the file are marked with a plus sign (+), and deletions to the file are marked with a minus sign (-). The body of the output contains the affected lines. Changes are viewed as additions plus deletions.</li> </ul>
<b>Options</b>	<p><b>files <i>filename</i></b>—Names of two local files to compare.</p> <p><b>context</b>—(Optional) Display output in context format.</p> <p><b>ignore-white-space</b>—(Optional) Ignore changes in the amount of white space.</p> <p><b>unified</b>—(Optional) Display output in unified format.</p>
<b>Required Privilege Level</b>	none
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Format for Specifying Filenames and URLs in Junos OS CLI Commands</i></li> <li>• <i>Viewing Core Files from Junos OS Processes</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">file compare files on page 654</a></p> <p><a href="#">file compare files context on page 654</a></p> <p><a href="#">file compare files unified on page 654</a></p> <p><a href="#">file compare files unified ignore-white-space on page 654</a></p>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file compare files

```
user@host> file compare files /tmp/one /tmp/two
100c100
<          full-name "File 1";
---
>          full-name "File 2";
102c102
<          class foo; # 'foo' is not defined
---
>          class super-user;
```

### file compare files context

```
user@host> file compare files /tmp/one /tmp/two context
*** /tmp/one   Wed Dec  3 17:12:50 2003
--- /tmp/two   Wed Dec  3 09:13:14 2003
*****
*** 97,104 ****
        }
    }
    user bill {
!         full-name "Bill Smith";
!         class foo; # 'foo' is not defined
        authentication {
            encrypted-password SECRET;
        }
--- 97,105 ----
    }
    user bill {
!         full-name "Bill Smith";
!         uid 1089;
!         class super-user;
        authentication {
            encrypted-password SECRET;
        }
    }
```

### file compare files unified

```
user@host> file compare files /tmp/one /tmp/two unified
--- /tmp/one   Wed Dec  3 17:12:50 2003
+++ /tmp/two   Wed Dec  3 09:13:14 2003
@@ -97,8 +97,9 @@
    }
}
user bill {
-     full-name "Bill Smith";
-     class foo; # 'foo' is not defined
+     full-name "Bill Smith";
+     uid 1089;
+     class super-user;
    authentication {
        encrypted-passwordSECRET;
    }
}
```

### file compare files unified ignore-white-space

```
user@host> file compare files /tmp/one /tmp/two unified ignore-white-space
```

```
--- /tmp/one    Wed Dec  3 09:13:10 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -99,7 +99,7 @@
     user bill {
         full-name "Bill Smith";
         uid 1089;
-        class foo; # 'foo' is not defined
+        class super-user;
         authentication {
             encrypted-password <SECRET>; # SECRET-DATA
         }
     }
```

## file copy

---

<b>Syntax</b>	<code>file copy <i>source destination</i></code> <code>&lt;source-address <i>address</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <b>source-address</b> option added in Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series switches.
<b>Description</b>	Copy files from one location to another location on the local device or to a location on a remote device reachable by the local device.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Format for Specifying Filenames and URLs in Junos OS CLI Commands</i></li><li>• <i>Default Directories for Junos OS File Storage on the Router or Switch</i></li><li>• <i>Copying a Configuration File from One Routing Engine to the Other</i></li></ul>
<b>List of Sample Output</b>	<a href="#">Copy a File from the Local Device to a Personal Computer on page 656</a> <a href="#">Copy a Configuration File between Routing Engines on page 656</a> <a href="#">Copy a Log File between Routing Engines on page 656</a> <a href="#">Copy a File from a TX Matrix Plus Router to a T1600 Router Connected to the TX Matrix Plus on page 657</a> <a href="#">Copy a File Using File Transfer Protocol on page 657</a> <a href="#">Copy a File Using File Transfer Protocol and Requiring a Password on page 657</a> <a href="#">Copy a File Using Secure Copy Protocol (scp) on page 657</a>

## Sample Output

The following are examples of a variety of file copy scenarios.

### Copy a File from the Local Device to a Personal Computer

```
user@host> file copy /var/tmp/rpd.core.4 mypc:/c/junipero/tmp
...transferring.file..... |           0 KB |    0.3 kB/s | ETA: 00:00:00 | 100%
```

### Copy a Configuration File between Routing Engines

The following sample command copies a configuration file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy /config/juniper.conf re1:/var/tmp/copied-juniper.conf
```

### Copy a Log File between Routing Engines

The following sample command copies a log file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy lcc0-re0:/var/log/chassisd lcc0-re1:/var/tmp
```



### Copy a File from a TX Matrix Plus Router to a T1600 Router Connected to the TX Matrix Plus

The following sample command copies a text file from Routing Engine 1 on the switch-fabric chassis sfc0 to Routing Engine 1 on the line-card chassis lcc0:

```
user@host> file copy sfc0-re1:/tmp/sample.txt lcc0-re1:/var/tmp
```

### Copy a File Using File Transfer Protocol

To use anonymous FTP to copy a local file to a remote system, enter the following command:

```
user@host> file copy filename ftp://hostname/filename
```

In the following example, `/config/juniper.conf` is the local file and `hostname` is the FTP server:

```
user@host> file copy /config/juniper.conf ftp://hostname/juniper.conf
Receiving ftp: //hostname/juniper.conf (2198 bytes): 100%
2198 bytes transferred in 0.0 seconds (2.69 MBps)
```

### Copy a File Using File Transfer Protocol and Requiring a Password

To use FTP where you require more privacy and are prompted for a password, enter the following command:

```
root@host> file copy filename ftp://user@hostname/filename
```

In the following example, `/config/juniper.conf` is the local file and `hostname` is the FTP server:

```
root@host> file copy /config/juniper.conf ftp://user@hostname/juniper.conf
Password for user@hostname: *****
Receiving ftp: //user@hostname/juniper.conf (2198 bytes): 100%
2198 bytes transferred in 0.0 seconds (2.69 MBps)
```

### Copy a File Using Secure Copy Protocol (scp)

To use scp to copy a local file to a remote system, enter the following command:

```
root@host> file copy filename scp://user@hostname/path/filename
```

In the following example, `/config/juniper.conf` is the local file, `user` is the username, and `ssh-host` is the scp server:

```
root@host> file copy /config/juniper.conf scp://user@ssh-host/tmp/juniper.conf
user@ssh-host's password: *****
juniper.conf          100%
|*****|
2198          00:00
```

## file delete

---

<b>Syntax</b>	<code>file delete <i>filename</i></code> <code>&lt;purge&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Delete a file on the local router or switch.
<b>Options</b>	<b><i>filename</i></b> —Name of the file to delete. For a routing matrix, include chassis information in the filename if the file to be deleted is not local to the Routing Engine from which the command is issued.  <b><i>purge</i></b> —(Optional) Overwrite regular files before deleting them.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file delete on page 658</a> <a href="#">file delete (Routing Matrix) on page 658</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file delete

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

### file delete (Routing Matrix)

```
user@host> file list lcc0-re0:/var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete lcc0-re0:/var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

## file list

<b>Syntax</b>	file list <detail   recursive> <filename>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display a list of files on the local router or switch.
<b>Options</b>	<p><b>none</b>—Display a list of all files for the current directory.</p> <p><b>detail   recursive</b>—(Optional) Display detailed output or descend recursively through the directory hierarchy, respectively.</p> <p><b>filename</b>—(Optional) Display a list of files. For a routing matrix, the filename must include the chassis information.</p>
<b>Additional Information</b>	The default directory is the home directory of the user logged in to the router or switch. To view available directories, enter a space and then a backslash (/) after the <b>file list</b> command. To view files within a specific directory, include a backslash followed by the directory and, optionally, subdirectory name after the <b>file list</b> command.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file list on page 659</a> <a href="#">file list (Routing Matrix) on page 659</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file list

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
```

### file list (Routing Matrix)

```
user@host> file list lcc0-re0:var/tmp
lcc0-re0:
-----
/var/tmp/:
.gdbinit
.pccardd
Test/
chassisd*
chassisd.nathan*
check_time*
```

```
cores/  
diagTestPrep*  
diagtest*  
diagtest.regress*  
do_switchovers*  
dump_test*  
err.manoj.log  
esw_clearstats*  
esw_counter*  
esw_debug*  
esw_debug_ge*  
esw_filt_test*  
esw_filter_tnp_addr*  
esw_getstats*  
esw_phy*  
esw_stats*
```

## file rename

<b>Syntax</b>	<code>file rename <i>source destination</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Rename a file on the local router or switch.
<b>Options</b>	<i>destination</i> —New name for the file.  <i>source</i> —Original name of the file. For a routing matrix, the filename must include the chassis information.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file rename on page 661</a> <a href="#">file rename (Routing Matrix) on page 661</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file rename

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

### file rename (Routing Matrix)

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
-----

/var/tmp:
.pccardd
sartre.conf
snmpd
syslogd.core-tarball.0.tgz
```

```
user@host> file rename lcc0-re0:/var/tmp/snmpd /var/tmp/snmpd.rr
```

```
user@host> file list lcc0-re1:/var/tmp
```

```
lcc0-re1:
```

```
-----
```

```
/var/tmp:
```

```
.pccardd
```

```
sartre.conf
```

```
snmpd.rr
```

```
syslogd.core-tarball.0.tgz
```

## file show

<b>Syntax</b>	<code>file show filename</code> <encoding (base64   raw)>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the contents of a file.
<b>Options</b>	<b>filename</b> —Name of a file. For a routing matrix, the filename must include the chassis information.  <b>encoding (base64   raw)</b> —(Optional) Encode file contents with base64 encoding or show raw text.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file show on page 663</a> <a href="#">file show (Routing Matrix) on page 663</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file show

```
user@host> file show /var/log/messages
Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
...
```

### file show (Routing Matrix)

```
user@host> file show lcc0-re0:/var/tmp/.gdbinit
lcc0-re0:
-----
#####
# Settings
#####

set print pretty


#####
# Basic stuff
#####

define msgbuf
    printf "%s", msgbufp->msg_ptr
end
```

```
# hex dump of a block of memory
# usage: dump address length
define dump
  p $arg0, $arg1
  set $ch = $arg0
  set $j = 0
  set $n = $arg1
  while ($j < $n)
    #printf "%x %x ",&$ch[$j],$ch[$j]
    printf "%x ",$ch[$j]
    set $j = $j + 1
    if (!($j % 16))
      printf "\n"
    end
  end
end
end
```



## request system configuration rescue delete

<b>Syntax</b>	request system configuration rescue delete
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Delete an existing rescue configuration.
<div>  <b>NOTE:</b> The [edit system configuration] hierarchy is not available on QFabric systems. </div>	
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request system configuration rescue save on page 666</a></li> <li>• <a href="#">request system software rollback on page 756</a></li> <li>• <a href="#">show system commit on page 675</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system configuration rescue delete on page 665</a>
<b>Output Fields</b>	This command produces no output.


### Sample Output

request system configuration rescue delete

```
user@host> request system configuration rescue delete
```

## request system configuration rescue save

---

<b>Syntax</b>	request system configuration rescue save
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Save the most recently committed configuration as the rescue configuration so that you can return to it at any time by using the <b>rollback</b> command.
<div> <b>NOTE:</b> The [edit system configuration] hierarchy is not available on QFabric systems.</div>	
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">request system software delete on page 752</a></li><li>• <a href="#">request system software rollback on page 756</a></li><li>• <a href="#">show system commit on page 675</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request system configuration rescue save on page 666</a>
<b>Output Fields</b>	This command produces no output.

### Sample Output

#### request system configuration rescue save

```
user@host> request system configuration rescue save
```

## request system scripts refresh-from commit

<b>Syntax</b>	<code>request system scripts refresh-from commit file <i>file-name</i> url <i>url-path</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.1 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Automatically download the initial Junos OS configuration and a set of standard commit scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.</p> <p>The Junos XML management protocol equivalent for this operational mode command is:</p> <pre>&lt;request-script-refresh-from&gt;   &lt;type&gt;commit&lt;/type&gt;   &lt;file&gt;file-name&lt;/file&gt;   &lt;URL&gt;URL&lt;/URL&gt; &lt;/request-script-refresh-from&gt;</pre>
<b>Options</b>	<p><b>file <i>file-name</i></b>—Name of the file to be downloaded.</p> <p><b>url <i>url-path</i></b>—URL of the file to be downloaded.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Automatic Refreshing of Scripts on EX Series Switches on page 613</a></li> <li>• <a href="#">Junos OS Junos XML Management Protocol Guide</a></li> <li>• <a href="#">Junos OS NETCONF XML Management Protocol Guide</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system scripts refresh-from commit file config.txt url http://host1.juniper.net on page 667</a>

### Sample Output

`request system scripts refresh-from commit file config.txt url http://host1.juniper.net`

```
user@switch> request system scripts refresh-from commit file config.txt url
http://host1.juniper.net
user@switch>
```

## request system scripts refresh-from event

---

<b>Syntax</b>	<code>request system scripts refresh-from event file <i>file-name</i> url <i>url-path</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.1 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Automatically download the initial Junos OS configuration and a set of standard event scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.</p> <p>The Junos XML management protocol equivalent for this operational mode command is:</p> <pre>&lt;request-script-refresh-from&gt;   &lt;type&gt;event&lt;/type&gt;   &lt;file&gt;file-name&lt;/file&gt;   &lt;URL&gt;URL&lt;/URL&gt; &lt;/request-script-refresh-from&gt;</pre>
<b>Options</b>	<p><code>file <i>file-name</i></code>—Name of the file to be downloaded.</p> <p><code>url <i>url-path</i></code>—URL of the file to be downloaded.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Automatic Refreshing of Scripts on EX Series Switches on page 613</a></li><li>• <a href="#">Junos OS Junos XML Management Protocol Guide</a></li><li>• <a href="#">Junos OS NETCONF XML Management Protocol Guide</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request system scripts refresh-from event file config.txt url http://host1.juniper.net on page 668</a>

### Sample Output

`request system scripts refresh-from event file config.txt url http://host1.juniper.net`

```
user@switch> request system scripts refresh-from event file config.txt url http://host1.juniper.net
user@switch>
```

## request system scripts refresh-from op

<b>Syntax</b>	<code>request system scripts refresh-from op file <i>file-name</i> url <i>url-path</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.1 for EX Series switches.
<b>Description</b>	<p>Automatically download the initial Junos OS configuration and a set of standard op scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.</p> <p>The Junos XML management protocol equivalent for this operational mode command is:</p> <pre>&lt;request-script-refresh-from&gt;   &lt;type&gt;op&lt;/type&gt;   &lt;file&gt;file-name&lt;/file&gt;   &lt;URL&gt;URL&lt;/URL&gt; &lt;/request-script-refresh-from&gt;</pre>
<b>Options</b>	<p><b>file <i>file-name</i></b>—Name of the file to be downloaded.</p> <p><b>url <i>url-path</i></b>—URL of the file to be downloaded.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Automatic Refreshing of Scripts on EX Series Switches on page 613</a></li> <li>• <a href="#">Junos OS Junos XML Management Protocol Guide</a></li> <li>• <a href="#">Junos OS NETCONF XML Management Protocol Guide</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system scripts refresh-from op file config.txt url http://host1.juniper.net on page 669</a>


### Sample Output

`request system scripts refresh-from op file config.txt url http://host1.juniper.net`

```
user@switch> request system scripts refresh-from op file config.txt url http://host1.juniper.net
user@switch>
```

## request system zeroize

---

Syntax	request system zeroize <media> <local>
Release Information	Command introduced before Junos OS Release 9.0. Command introduced in Junos OS Release 11.2 for EX Series switches. Option <b>media</b> added in Junos OS Release 11.4 for EX Series switches. Command introduced in Junos OS Release 12.2 for MX Series routers. Command introduced in Junos OS Release 12.3 for the QFX Series. Option <b>local</b> added in Junos OS Release 14.1.
Description	<div> <b>NOTE:</b> The <b>media</b> option is not available on the QFX Series.</div> <p>Remove all configuration information on the Routing Engines and reset all key values. If the device has dual Routing Engines, the command is broadcast to all Routing Engines on the device. The command removes all data files, including customized configuration and log files, by unlinking the files from their directories. The command removes all user-created files from the system including all plain-text passwords, secrets, and private keys for SSH, local encryption, local authentication, IPsec, RADIUS, TACACS+, and SNMP.</p> <p>This command reboots the device and sets it to the factory default configuration. After the reboot, you cannot access the device through the management Ethernet interface. Log in through the console as <b>root</b> and start the Junos OS CLI by typing <b>cli</b> at the prompt.</p> <p>To completely erase user-created data so that it is unrecoverable, use the <b>media</b> option.</p>
Options	<p><b>media</b>—(Optional) In addition to removing all configuration and log files, causes memory and the media to be scrubbed, removing all traces of any user-created files. Every storage device attached to the system is scrubbed, including disks, flash drives, removable USBs, and so on. The duration of the scrubbing process is dependent on the size of the media being erased. As a result, the <b>request system zeroize media</b> operation can take considerably more time than the <b>request system zeroize</b> operation. However, the critical security parameters are all removed at the beginning of the process.</p> <p><b>local</b>—(Optional) Remove all the configuration information and restore all the key values on the active Routing Engine.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">request system snapshot on page 737</a></li><li>• <a href="#">request system snapshot</a></li><li>• <a href="#">Reverting to the Default Factory Configuration for the EX Series Switch on page 620</a></li></ul>

- [Reverting to the Rescue Configuration for the EX Series Switch on page 624](#)
- *Reverting to the Default Factory Configuration*
- *Reverting to the Rescue Configuration*
- *Reverting to the Default Factory Configuration by Using the request system zeroize Command*

List of Sample Output [request system zeroize on page 671](#)  
[request system zeroize media on page 672](#)

## Sample Output

### request system zeroize

```
user@host> request system zeroize
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (no) yes

0 1 1 0 0 0 done

syncing disks... All buffers synced.
Uptime: 5d19h20m26s
recorded reboot as normal shutdown
Rebooting...

U-Boot 1.1.6 (Mar 11 2011 - 04:39:06)

Board: EX4200-24T 2.11
EPLD: Version 6.0 (0x85)
DRAM: Initializing (1024 MB)
FLASH: 8 MB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.4
(user@juniper.net, Fri Mar 11 03:03:36 UTC 2011)
Memory: 1024MB
bootsequencing is enabled
bootsuccess is set
new boot device = disk0s1:
Loading /boot/defaults/loader.conf
/kernel data=0x915c84+0xa1260 syms=[0x4+0x7cbd0+0x4+0xb1c19]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
```

```
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 11.1R1.8 #0: 2011-03-09 20:14:25 UTC
```

```
user@juniper.net:/volume/build/junos/11.1/release/11.1R1.8/obj-powerpc/bsd/kernels/
    JUNIPER-EX/kernel
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080
...
```

#### request system zeroize media

```
user@host> request system zeroize media
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (no) yes

warning: ipsec-key-management subsystem not running - not needed by configuration.
warning: zeroizing fpc0

{master:0}
root> Waiting (max 60 seconds) for system process `vnlr' to stop...done
...
Syncing disks, vnodes remaining...2 4 2 4 3 2 1 1 0 0 0 done

syncing disks... All buffers synced.
Uptime: 14m50s
recorded reboot as normal shutdown
Rebooting...

U-Boot 1.1.6 (Apr 21 2011 - 13:58:42)

Board: EX4200-48PX 1.1
EPLD: Version 8.0 (0x82)
DRAM: Initializing (512 MB)
FLASH: 8 MB
NAND: No NAND device found!!!
0 MiB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.2
(vtseng@svl-junos-pool27.juniper.net, Fri Feb 26 17:48:51 PST 2010)
Memory: 512MB
Loading /boot/defaults/loader.conf
/kernel data=0x9abfdc+0xb06e4 syms=[0x4+0x83b30+0x4+0xbd7c6]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel] in 1 second... Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
```



```

Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.4R1.2 #0: 2011-10-27 18:05:39 UTC
user@juniper.net:/volume/build/junos/11.4/release/11.4R1.2/obj-powerpc/
bsd/kernels/JUNIPER-EX/kernel
can't re-use a leaf (all_slot_serialid)!
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080<EMCP,TBEN,EN_MAS7_UPDATE>
real memory = 511705088 (488 MB)
avail memory = 500260864 (477 MB)
ETHERNET SOCKET BRIDGE initialising
Initializing EXSERIES platform properties ...
. . .
Automatic reboot in progress...
Media check on da0 on ex platforms
** /dev/da0s2a
FILE SYSTEM CLEAN; SKIPPING CHECKS
clean, 20055 free (31 frags, 2503 blocks, 0.0% fragmentation)
zeroizing /dev/da0s1a ...
. . .
zeroizing /dev/da0s3d ...
. . .
zeroizing /dev/da0s3e ...
. . .
zeroizing /dev/da0s4d ...
. . .
zeroizing /dev/da0s4e ...
. . .

syncing disks... All buffers synced.
Uptime: 3m40s
Rebooting...

U-Boot 1.1.6 (Apr 21 2011 - 13:58:42)

Board: EX4200-48PX 1.1
EPLD: Version 8.0 (0x82)
DRAM: Initializing (512 MB)
FLASH: 8 MB
NAND: No NAND device found!!!
0 MiB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.2
(vtseng@svl-junos-pool27.juniper.net, Fri Feb 26 17:48:51 PST 2010)
Memory: 512MB
Loading /boot/defaults/loader.conf
/kernel data=0x9abfdc+0xb06e4 syms=[0x4+0x83b30+0x4+0xbd7c6]


Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel] in 1 second... Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb

```

```
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.4R1.2 #0: 2011-10-27 18:05:39 UTC
user@juniper.net:/volume/build/junos/11.4/release/11.4R1.2/obj-powerpc/
bsd/kernels/JUNIPER-EX/kernel
can't re-use a leaf (all_slot_serialid!)
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HIDO 80004080 <EMCP,TBEN,EN_MAS7_UPDATE>
real memory = 511705088 (488 MB)
avail memory = 500260864 (477 MB)
ETHERNET SOCKET BRIDGE initialising
Initializing EXSERIES platform properties ...
. . .
Automatic reboot in progress...
Media check on da0 on ex platforms
** /dev/da0s1a
FILE SYSTEM CLEAN; SKIPPING CHECKS
clean, 20064 free (48 frags, 2502 blocks, 0.1% fragmentation)
zeroizing /dev/da0s2a ...
. . .
Creating initial configuration...mgd: error: Cannot open configuration file:
/config/juniper.conf
mgd: warning: activating factory configuration
mgd: commit complete
mgd: -----
mgd: Please login as 'root'. No password is required.
mgd: To start Initial Setup, type 'ezsetup' at the JUNOS prompt.
mgd: To start JUNOS CLI, type 'cli' at the JUNOS prompt.
mgd: -----
Setting initial options: debugger_on_panic=NO debugger_on_break=NO.
Starting optional daemons: .
Doing initial network setup:
. . .

Amnesiac (ttyu0)
```

## show system commit

<b>Syntax</b>	<pre>show system commit &lt;revision&gt; &lt;server&gt;</pre>	
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <b>server</b> introduced in Junos OS Release 12.1 for the PTX Series router.</p> <p>Option <b>revision</b> introduced in Junos OS Release 14.1.</p>	
<b>Description</b>	Display the system commit history and any pending commit operation.	
<b>Options</b>	<p><b>none</b>—Display the last 50 commit operations listed, most recent to first.</p> <p><b>revision</b>—(Optional) Display the revision number of the active configuration of the Routing Engine(s).</p> <p><b>server</b>—(Optional) Display commit server status.</p>	
	<div>  <p><b>NOTE:</b> By default, the status of the commit server is “Not running”. The commit server starts running only when a commit job is added to the batch.</p> </div>	
<b>Required Privilege Level</b>	view	
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear system commit on page 647</a></li> <li><a href="#">show system commit revision</a></li> </ul>	
<b>List of Sample Output</b>	<a href="#">show system commit on page 677</a> <a href="#">show system commit (At a Particular Time) on page 677</a> <a href="#">show system commit (At the Next Reboot) on page 677</a> <a href="#">show system commit (Rollback Pending) on page 677</a> <a href="#">show system commit (QFX Series) on page 677</a>	
<b>Output Fields</b>	<p><a href="#">Table 61 on page 675</a> describes the output fields for the <b>show system commit</b> command. Output fields are listed in the approximate order in which they appear.</p>	

**Table 61: show system commit Output Fields**

Field Name	Field Description	Level of Output
<b>&lt;number&gt;</b>	Displays the last 50 commit operations listed, most recent to first. The identifier <b>&lt;number&gt;</b> designates a configuration created for recovery using the <b>request system configuration rescue save</b> command.	<b>none</b>

Table 61: show system commit Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>&lt;time-stamp&gt;</b>	Date and time of the commit operation.	<b>none</b>
<b>&lt;root&gt;/&lt;username&gt;</b>	User who executed the commit operation.	<b>none</b>
<b>&lt;method&gt;</b>	<p>Method used to execute the commit operation:</p> <ul style="list-style-type: none"> <li>• <b>CLI</b>—CLI interactive user performed the commit operation.</li> <li>• <b>Junos XML protocol</b>—Junos XML protocol client performed the commit operation.</li> <li>• <b>synchronize</b>—The <b>commit synchronize</b> command was performed on the other Routing Engine.</li> <li>• <b>snmp</b>—An SNMP <b>set</b> request caused the commit operation.</li> <li>• <b>button</b>—A button on the router or switch was pressed to commit a rescue configuration for recovery.</li> <li>• <b>autoinstall</b>—A configuration obtained through autoinstallation was committed.</li> <li>• <b>other</b>—When there is no login name associated with the session, the values for user and client default to root and other. For example, during a reboot after package installation, mgd commits the configuration as a system commit, and there is no login associated with the commit.</li> </ul>	<b>none</b>

## Sample Output

### show system commit

```
user@host> show system commit
0   2003-07-28 19:14:04 PDT by root via other
1   2003-07-25 22:01:36 PDT by regress via cli
2   2003-07-25 22:01:32 PDT by regress via cli
3   2003-07-25 21:30:13 PDT by root via button
4   2003-07-25 13:46:48 PDT by regress via cli
5   2003-07-25 05:33:21 PDT by root via autoinstall
...
rescue 2002-05-10 15:32:03 PDT by root via other
```

### show system commit (At a Particular Time)

```
user@host> show system commit
commit requested by root via cli at Tue May  7 15:59:00 2002
```

### show system commit (At the Next Reboot)

```
user@host> show system commit
commit requested by root via cli at reboot
```

### show system commit (Rollback Pending)

```
user@host> show system commit
0 2005-01-05 15:00:37 PST by root via cli commit confirmed, rollback in 3mins
```

### show system commit (QFX Series)

```
user@switch> show system commit
0 2011-11-25 19:17:49 PST by root via cli
```

## show system configuration archival

---

**Syntax**    show system configuration archival

**Release Information**    Introduced in Junos OS Release 7.6.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Display directory and number of files queued for archival transfer.



**NOTE:** The [edit system configuration] hierarchy is not available on QFabric systems.

---

**Options**    This command has no options.

**Required Privilege Level**    maintenance

**List of Sample Output**    [show system configuration archival on page 678](#)

### Sample Output

show system configuration archival

```
user@host> show system configuration archival

/var/transfer/config/:
total 8
```

## show system configuration rescue

<b>Syntax</b>	show system configuration rescue
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display a rescue configuration, if one exists.



**NOTE:** The [edit system configuration] hierarchy is not available on QFabric systems.

<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show system configuration archival on page 678</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show system configuration rescue on page 679</a>

## Sample Output

### show system configuration rescue

```


user@switch> show system configuration rescue
version "7.3"; groups {
  global {
    system {
      host-name router1;
      domain-name customer.net;
      domain-search [ customer.net ];
      backup-router 192.168.124.254;
      name-server {
        172.17.28.11;
        172.17.28.101;
        172.17.28.100;
        172.17.28.10;
      }
      login {
        user regress {
          uid 928;
          class ;
          shell csh;
          authentication {
            encrypted-password "$1$kPU..$w.4FGRAGanJ8U4Yq6sbj7."; ##
SECRET-DATA
          }
        }
      }
    }
  }
  services {

```

```
        ftp;  
        rlogin;  
        rsh;  
        telnet;  
    }  
}  
.....
```



## show system rollback

<b>Syntax</b>	<code>show system rollback <i>number</i></code> <code>&lt;compare <i>number</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the contents of a previously committed configuration, or the differences between two previously committed configurations.
<div>  <b>NOTE:</b> The <code>show system rollback</code> command is a purely operational mode command and cannot be issued with <code>run</code> from the configuration mode.         </div>	
<b>Options</b>	<p><b><i>number</i></b>—Number of a configuration to view. The output displays the configuration. The range of values is 0 through 49.</p> <p><b><code>compare <i>number</i></code></b>—(Optional) Number of another previously committed (rollback) configuration to compare to rollback <b><i>number</i></b>. The output displays the differences between the two configurations. The range of values is 0 through 49.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show system rollback compare on page 681</a>

## Sample Output

### show system rollback compare

```

user@host> show system rollback 3 compare 1
[edit]
+ interfaces {
+   ge-1/1/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 14.1.1.1/30;
+       }
+     }
+   }
+   ge-1/2/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 13.1.1.1/30;
+       }
+     }
+   }
+ }

```

```
+      }
+    }
+    ge-1/3/0 {
+      unit 0 {
+        family inet {
+          filter {
+            input mf_plp;
+          }
+          address 12.1.1.1/30;
+        }
+      }
+    }
+  }
+}
```

## test configuration

<b>Syntax</b>	<code>test configuration <i>filename</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Verify that the syntax of a configuration file is correct. If the configuration contains any syntax or commit check errors, a message is displayed to indicate the line number and column number in which the error was found. This command only accepts text files.
<b>Options</b>	<i>filename</i> —Name of the configuration file.  <b>syntax-only</b> —Check the syntax of a partial configuration file, without checking for commit errors. This option introduced in Junos OS Release 12.1.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">test configuration on page 683</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### test configuration

```

user@host> test configuration terminal
[Type ^D to end input]
system {
host-name bluesky;
paris-23;
login;
}
terminal:3:(8) syntax error: paris
[edit system]
    'paris-23;'
    syntax error
terminal:4:(11) statement must contain additional statements: ;
[edit system login]
    'login ;'
    statement must contain additional statements
configuration syntax failed

```



# Troubleshooting Procedures

- Troubleshooting Loss of the Root Password on page 685

## Troubleshooting Loss of the Root Password

---

**Problem**    **Description:** If you forget the root password for a switch, use the password recovery procedure to reset the root password.



**NOTE:** You need physical access to the switch to recover the root password.




**TIP:** For a video on recovering the root password for routers, see *Recovering the Root Password*. The procedure is similar for switches.

**Solution**    To recover the root password:

1. Power off your switch by unplugging the power cord or turning off the power at the wall switch.
2. Insert one end of the Ethernet cable into the serial port on the management device and connect the other end to the console port on the back of the switch. See [Figure 1 on page 306](#).

Figure 4: Connecting to the Console Port on the EX Series Switch



3. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate COM port to use (for example, COM1).
4. Configure the port settings as follows:
  - Bits per second: 9600
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: None
5. Power on your switch by plugging in the power cord or turning on the power at the wall switch.
6. When the following prompt appears, press the Spacebar to access the switch's bootstrap loader command prompt:
 

Hit [Enter] to boot immediately, or space bar for command prompt.  
Booting [kernel] in 1 second...



**NOTE:** If the switch is in unattended mode for U-Boot, access to the bootstrap loader command prompt is blocked. If the root password is lost, you must reset the switch to the factory default configuration using the LCD panel. For more information, see [“Reverting to the Default Factory Configuration for the EX Series Switch”](#) on page 620.

7. At the following prompt, type **boot -s** to start up the system in single-user mode:  
**loader> boot -s**
8. At the following prompt, type **recovery** to start the root password recovery procedure:  
**Enter full path name of shell or 'recovery' for root password recovery or RETURN for /bin/sh: recovery**

A series of messages describe consistency checks, mounting of filesystems, and initialization and checkout of management services. Then the CLI prompt appears.

9. Enter configuration mode in the CLI:

```
user@switch> configure
```

10. Set the root password. For example:

```
user@switch# set system root-authentication plain-text-password
```

11. At the following prompt, enter the new root password. For example, juniper1:

```
user@switch# juniper1
```

```
Retype new password:
```

12. At the second prompt, reenter the new root password.

13. If you are finished configuring the network, commit the configuration.

```
root@switch# commit
```

```
commit complete
```

14. Exit configuration mode in the CLI.

```
root@switch# exit
```

15. Exit operational mode in the CLI.

```
root@switch> exit
```

16. At the prompt, enter **y** to reboot the switch.

```
Reboot the system? [y/n] y
```

#### Related Documentation

- *Connecting and Configuring an EX Series Switch (CLI Procedure)*
- *Connecting and Configuring an EX Series Switch (J-Web Procedure)*
- For information about configuring an encrypted root password, configuring SSH keys to authenticate root logins, and configuring special requirements for plain-text passwords, see *Configuring the Root Password*.





## PART 6

# Software Installation

- [Overview on page 691](#)
- [Configuration on page 711](#)
- [Administration on page 719](#)
- [Troubleshooting Procedures on page 785](#)



## CHAPTER 17

# Overview

- [Installation Overview on page 691](#)
- [Software Overview on page 700](#)
- [Licenses Overview on page 701](#)

## Installation Overview

---

- [Understanding Software Installation on EX Series Switches on page 691](#)
- [Understanding System Snapshot on EX Series Switches on page 694](#)
- [Understanding Resilient Dual-Root Partitions on Switches on page 695](#)
- [Junos OS Package Names on page 698](#)

## Understanding Software Installation on EX Series Switches

A Juniper Networks EX Series Ethernet Switch is delivered with the Juniper Networks Junos operating system (Junos OS) preinstalled. As new features and software fixes become available, you must upgrade your software to use them. You can also downgrade Junos OS to a previous release.

This topic covers:

- [Overview of the Software Installation Process on page 691](#)
- [Software Package Security on page 692](#)
- [Installing Software on a Virtual Chassis on page 692](#)
- [Installing Software on Switches with Redundant Routing Engines on page 692](#)
- [Installing Software Using Automatic Software Download on page 693](#)
- [Autoinstalling a Configuration File on an EX2200 or EX3300 Switch from a Disk-on-Key USB Memory Stick on page 693](#)
- [Troubleshooting Software Installation on page 693](#)

### Overview of the Software Installation Process

---

An EX Series switch is delivered with a domestic version of Junos OS preinstalled. When you connect power to the switch, it starts (boots) from the installed software.

You upgrade Junos OS on an EX Series switch by copying a software package to your switch or another system on your local network, then use either the J-Web interface or the command-line interface (CLI) to install the new software package on the switch. Finally, you reboot the switch; it boots from the upgraded software. After a successful upgrade, you should back up the new current configuration to a secondary device. You should follow this procedure regardless of whether you are installing a domestic or controlled Junos OS package.

During a successful upgrade, the upgrade package removes all files from `/var/tmp` and completely reinstalls the existing software. It retains configuration files, and similar information, such as secure shell and host keys, from the previous version. The previous software package is preserved in a separate disk partition, and you can manually revert back to it if necessary. If the software installation fails for any reason, such as loss of power during the installation process, the system returns to the originally active installation when you reboot.

---

### Software Package Security

All Junos OS releases are delivered in signed packages that contain digital signatures to ensure official Juniper Networks software. For more information about signed software packages, see the *Junos OS Installation and Upgrade Guide*.

---

### Installing Software on a Virtual Chassis

You can connect individual EX Series switches together to form one unit and manage the unit as a single device, called a Virtual Chassis. The Virtual Chassis operates as a single network entity composed of member switches. Each member switch in a Virtual Chassis must be running the same version of Junos OS. See *EX Series Virtual Chassis Software Features Overview* for a list of switches that can be used in a Virtual Chassis.

For ease of management, a Virtual Chassis provides flexible methods to upgrade software releases. You can deploy a new software release to all member switches of a Virtual Chassis or to only a particular member switch.

You can also upgrade the software on an EX4200, EX4500, mixed EX4200 and EX4500, and EX8200 Virtual Chassis using nonstop software upgrade (NSSU). NSSU takes advantage of graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) to ensure no disruption to the control plane during the upgrade. You can minimize disruption to network traffic by defining link aggregation groups (LAGs) such that the member links of each LAG reside on different line cards (on EX8200 Virtual Chassis) or on different members (on EX4200, EX4500, mixed EX4200 and EX4500 Virtual Chassis). During an NSSU, the line cards and Virtual Chassis members are upgraded one at a time, so that traffic continues to flow through the other line cards or members while that line card or member is being upgraded.

---

### Installing Software on Switches with Redundant Routing Engines

You can install software on a switch with redundant Routing Engines in one of two ways:

- Perform an NSSU—An NSSU upgrades both Routing Engines with a single command and with a minimum of network disruption. An NSSU takes advantage of GRES and NSR to ensure no disruption to the control plane. You can minimize disruption to

network traffic by defining LAGs such that the member links of each LAG reside on different line cards. The line cards are upgraded one at a time, so that traffic continues to flow through the other line cards while a line card is being upgraded.

You cannot use NSSU to downgrade the software running on a switch.

For more information about NSSU, see [“Understanding Nonstop Software Upgrade on EX Series Switches” on page 2488](#). See *EX Series Switch Software Features Overview* for a list of switches that support NSSU.

- Upgrade each Routing Engine manually—You can perform a Junos OS installation on each Routing Engine separately, starting with the backup Routing Engine. You can use this procedure to downgrade the software running on a switch. See *Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)*.

### Installing Software Using Automatic Software Download

The automatic software download feature uses the DHCP message exchange process to download and install software packages. Users can define a path to a software package on the DHCP server and then the DHCP server communicates this path to EX Series switches acting as DHCP clients as part of the DHCP message exchange process. The DHCP clients that have been configured for automatic software download receive these messages and, when the software package name in the DHCP server message is different from that of the software package that booted the DHCP client switch, download and install the software package. See *Upgrading Software by Using Automatic Software Download*.

### Autoinstalling a Configuration File on an EX2200 or EX3300 Switch from a Disk-on-Key USB Memory Stick

You can use an autoinstallation process to configure the software on an EX2200 or EX3300 switch. You can use a configuration file that is in either text format or XML format. If you want to use an XML-formatted file, you use a Junos Space platform to create the configuration file. You place the configuration file on a Disk-on-Key USB memory stick. See *Autoinstalling a Configuration File from a Disk-on-Key USB Memory Stick onto an EX2200 or EX3300 Switch*.

### Troubleshooting Software Installation

If Junos OS loads but the CLI is not working for any reason, or if the switch has no software installed, you can use the recovery installation procedure to install the software on the switch. See [“Troubleshooting Software Installation” on page 785](#).



**NOTE:** You can also use this procedure to load two versions of Junos OS in separate partitions on the switch.

#### Related Documentation

- *Downloading Software Packages from Juniper Networks*
- *Installing Software on EX Series Switches (J-Web Procedure)*
- *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*

- [Installing Software on an EX Series Switch with Redundant Routing Engines \(CLI Procedure\)](#)
- [Understanding Nonstop Software Upgrade on EX Series Switches on page 2488](#)

## Understanding System Snapshot on EX Series Switches

You can create copies of the software running a Juniper Networks EX Series Ethernet Switch using the system snapshot feature. The system snapshot feature takes a “snapshot” of the files currently used to run the switch and copies them to an alternate storage location. You can then use this snapshot to boot the switch at the next bootup or as a backup boot option.

The switch can boot from either internal flash media or external (USB) flash media. The contents of the snapshot vary depending on whether you create the snapshot on the media that the switch booted from or on the media that it did not boot from:

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

- If you create the snapshot on the media that the switch did not boot from, the following partitions on the boot media are included in the snapshot: **root**, **altroot**, **var**, **var/tmp**, **config**.

The **root** partition is the primary boot partition, and the **altroot** partition is the backup boot partition.

- If you create the snapshot on the media that the switch booted from, the root partition that the switch booted from is copied to the alternate root partition. The **var**, **var/tmp**, and **config** partitions are not copied as part of the snapshot because they already exist on the boot media.

The system snapshot feature has the following limitations:

- You cannot use snapshots to move files to any destination outside the switch other than an installed external USB flash drive or switches that are members of the same Virtual Chassis as the switch on which you created the snapshot..
- Snapshot commands, like all commands executed on a Virtual Chassis, are executed on the local member switch. If different member switches request the snapshot, the snapshot command is pushed to the Virtual Chassis member creating the snapshot and is executed on that member, and the output is then returned to the switch that initiated the process. For instance, if the command to create an external snapshot on member 3 is entered on member 1, the snapshot of internal memory on member 3 is taken on external memory on member 3. The output of the process is seen on member 1. No files move between the switches.

### Related Documentation

- [Understanding Software Installation on EX Series Switches on page 691](#)
- [Creating a Snapshot and Using It to Boot an EX Series Switch on page 713](#)

Understanding Resilient Dual-Root Partitions on Switches

Resilient dual-root partitioning, introduced on Juniper Networks EX Series Ethernet Switches in Juniper Networks Junos operating system (Junos OS) Release 10.4R3, provides additional resiliency to switches in the following ways:

- Allows the switch to boot transparently from the second (alternate) root partition if the system fails to boot from the primary root partition.
- Provides separation of the root Junos OS file system from the `/var` file system. If corruption occurs in the `/var` file system (a higher probability than in the root file system because of the greater frequency of reads and writes in `/var`), the root file system is insulated from the corruption.



**NOTE:** For instructions on upgrading to a release that supports resilient dual-root partitions from a release that does not, see the release notes. The procedure for upgrading to a resilient dual-root partition release is different from the normal upgrade procedure.

This topic covers:

- [Resilient Dual-Root Partition Scheme \(Junos OS Release 10.4R3 and Later\) on page 695](#)
- [Automatic Fixing of Corrupted Primary Root Partition with the Automatic Snapshot Feature on page 696](#)
- [Earlier Partition Scheme \(Junos OS Release 10.4R2 and Earlier\) on page 697](#)
- [Understanding Upgrading or Downgrading Between Resilient Dual-Root Partition Releases and Earlier Releases on page 697](#)

Resilient Dual-Root Partition Scheme (Junos OS Release 10.4R3 and Later)

EX Series switches that ship with Junos OS Release 10.4R3 or later are configured with a root partition scheme that is optimized for resiliency, as shown in [Table 62 on page 695](#).

Table 62: Resilient Dual-Root Partition Scheme

Slice 1	Slice 2	Slice 3		Slice 4
s1a	s2a	s3e	s3d	s4d
/	/	/var	/var/tmp	/config
(root Junos OS )	(root Junos OS )			

In the resilient dual-root partition scheme, the `/var` file system is contained in a separate slice (Slice 3) from the root file systems; the `/config` directory is contained in its own slice (Slice 4); and switches ship from the factory with identical Junos OS images in Slice 1 and Slice 2. The `/var` file system, which has a greater frequency of reads and writes than the root file systems and is therefore more likely to have corruption issues, is isolated from the root directories and the `/config` directory. If the switch fails to boot from the

active partition, the switch automatically boots from the alternate root partition and triggers an alarm.

### Automatic Fixing of Corrupted Primary Root Partition with the Automatic Snapshot Feature

---

Resilient dual-root partitioning also provides the *automatic snapshot* feature, which allows the switch to automatically fix a corrupt Junos OS file in the primary root partition. If the automatic snapshot feature is enabled, the switch automatically takes a snapshot of the Junos OS root file system in the alternate root partition and copies it onto the primary root partition, thereby repairing the corrupt file in the primary root partition. The automatic snapshot procedure takes place whenever the system reboots from the alternate root partition, regardless of whether the reboot is due to a command or due to corruption of the primary root partition.



---

#### NOTE:

- EX9200 switches do not support the automatic snapshot feature.
  - The automatic snapshot feature is enabled by default on the following EX Series switches:
    - EX4550 switches
    - EX Series switches that ship with Junos OS Release 12.3R1 or later
  - The automatic snapshot feature is disabled by default on EX Series switches (except the EX4550 switches) running Junos OS Release 12.2 or earlier.
  - If the automatic snapshot feature was disabled by default before the switch was upgraded to Junos OS Release 12.3R1 or later, the feature remains disabled (for backward compatibility) by default after the upgrade.
  - If the automatic snapshot feature is enabled in a Virtual Chassis configuration, the automatic snapshot procedure takes place whenever any member of the Virtual Chassis reboots from its alternate root partition.
  - You can enable the automatic snapshot feature by configuring the **auto-snapshot** statement at the **[edit system]** hierarchy level.
- 

The automatic snapshot feature provides an additional layer of fault protection if you maintain the same version of Junos OS in both partitions of resilient dual-root partitions. When **auto-snapshot** is enabled, repair happens automatically. Therefore, the switch does not issue an alarm to indicate that the system has rebooted from the alternate partition. However, it does log the event. You cannot execute a manual snapshot when an automatic snapshot procedure is in process. The login banner indicates that an automatic snapshot operation is in progress and that banner is removed only after the snapshot operation is complete. The next reboot happens from the primary partition.





**NOTE:** EX Series switches that ship with Junos OS Release 10.4R3 or later are configured with identical Junos OS images in the primary root partition (Slice 1) and the alternate root partition (Slice 2).

However, if you do *not* maintain the same version of Junos OS in both partitions, you might want to disable the automatic snapshot feature. If you have an earlier version of Junos OS in the alternate partition and the system reboots from the alternate root partition, the automatic snapshot feature causes the later Junos OS version to be replaced with the earlier version.

When automatic snapshot is disabled and the system reboots from the alternate root partition, it triggers an alarm indicating that the system has rebooted from its alternate partition.

### Earlier Partition Scheme (Junos OS Release 10.4R2 and Earlier)

The partition scheme used in Junos OS 10.4R2 and earlier is shown in [Table 63 on page 697](#).

**Table 63: Earlier Partition Scheme**

Slice 1		Slice 2		Slice 3	
s1a	s1f	s2a	s2f	s3d	s3e
/	/var	(empty until initial software upgrade)	(empty until initial software upgrade)	/var/tmp	/config
(root Junos OS)					

This is the partitioning scheme for a switch shipped with Release 10.4R2 or earlier (or after you reformat the disk during a downgrade from Release 10.4R3 or later to Release 10.4R2 or earlier). In this partitioning scheme, the switch comes from the factory with only one Junos OS image installed in the root Junos OS partition of Slice 1. The first time that you perform a software upgrade, the new Junos OS image is installed in Slice 2. If the switch fails to boot, you must manually trigger it to boot from the alternate partition (rebooting from the alternate partition does not occur automatically).

### Understanding Upgrading or Downgrading Between Resilient Dual-Root Partition Releases and Earlier Releases

Upgrading from Release 10.4R2 or earlier to Release 10.4R3 or later differs from other upgrades in two important ways:

- You must install a new loader software package in addition to installing the new Junos OS image.
- Rebooting after the upgrade reformats the disk from three partitions to four partitions. See [Table 62 on page 695](#).

You can perform all operations for this special software upgrade from the CLI.



**CAUTION:** Back up any important log files because the `/var/log` files are not saved or restored during an upgrade from Release 10.4R2 or earlier to a release that supports resilient dual-root partitions (Release 10.4R3 or later).

We recommend that you also save your `/config` files and any important log files to an external medium because if there is a power interruption during the upgrade process, they might be lost.

**Related  
Documentation**

- *Resilient Dual-Root Partitions Frequently Asked Questions*
- *EX Series Virtual Chassis Overview*
- *EX8200 Virtual Chassis Overview*

## Junos OS Package Names

You upgrade the Juniper Networks Junos operating system (Junos OS) on a Juniper Networks EX Series Ethernet Switch by copying a software package to your switch or another system on your local network, then install the new software package on the switch.

Two versions of a Junos OS image—a controlled version that supports Media Access Control Security (MACsec) and a domestic version that does not support MACsec—are available for EX Series switches. A domestic version of Junos OS is available for all EX Series switches; a controlled version of Junos OS is only available for EX Series switches on Junos OS releases that support MACsec. The domestic version of Junos OS on EX Series switches can be used on any switch in any geography. The controlled version of Junos OS contains encryption and is not available to customers in all geographies.



**NOTE:** The controlled version of Junos OS contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of the controlled version of Junos OS is strictly controlled under United States export laws. The export, import, and use of the controlled version of Junos OS is also subject to controls imposed under the laws of other countries.

If you have questions about acquiring the controlled version of Junos OS in your country, contact the Juniper Networks Trade Compliance group at [compliance\\_helpdesk@juniper.net](mailto:compliance_helpdesk@juniper.net).



**NOTE:** The domestic version of Junos OS on EX Series switches is intended for use on any switch in any worldwide location.

For most Junos packages on other Juniper Networks products, the domestic package is used for products installed in the United States and Canada only while an export package is used for products installed in any worldwide location.

**domestic-signed** indicates the domestic software package.

A domestic software package name is in the following format:

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

A controlled software package name is in the following format:

***package-name-m.nZx.y-controlled-signed.tgz***

where:

- ***package-name*** is the name of the package—for example, ***jinstall-ex-4200***.
- ***m.n*** is the software release, with ***m*** representing the major release number and ***n*** representing the minor release number—for example, ***9.5***.
- ***Z*** indicates the type of software release, where ***R*** indicates released software and ***B*** indicates beta-level software.
- ***x.y*** represents the version of the major software release (***x***) and an internal tracking number (***y***)—for example, ***1.6***.
- ***domestic-signed*** indicates the domestic software package.
- ***controlled-signed*** indicates the controlled software package.

A sample EX Series software domestic package name is:

***jinstall-ex-4200-9.5R1.6-domestic-signed.tgz***

A sample EX Series controlled package name is:

***jinstall-ex-4200-13.2X50-D15.3-controlled-signed.tgz***

#### Related Documentation

- *Installing Software on EX Series Switches (J-Web Procedure)*
- *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*
- *Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)*
- *Downloading Software Packages from Juniper Networks*
- [Understanding Software Installation on EX Series Switches on page 691](#)

## Software Overview

---

- [Understanding Software Infrastructure and Processes on page 700](#)

### Understanding Software Infrastructure and Processes

Each switch runs the Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches on its general-purpose processors. Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the chassis.

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

With the J-Web interface and the command-line interface (CLI) to the Junos OS, you configure switching features and routing protocols and set the properties of network interfaces on your switch. After activating a software configuration, use either the J-Web or CLI user interface to monitor the switch, manage operations, and diagnose protocol and network connectivity problems.

- [Routing Engine and Packet Forwarding Engine on page 700](#)
- [Junos OS Processes on page 700](#)

### Routing Engine and Packet Forwarding Engine

---

A switch has two primary software processing components:

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

### Junos OS Processes

---

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

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

Table 6 on page 30 describes the primary Junos OS processes.

**Table 64: Junos OS Processes**

Process	Name	Description
Chassis process	chassisd	<p>Detects hardware on the system that is used to configure network interfaces.</p> <p>Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered.</p> <p>Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.</p>
Ethernet switching process	eswd	<p>Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree protocol and access port security. The process is also responsible for managing Ethernet switching interfaces, VLANs, and VLAN interfaces.</p> <p>Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.</p>
Forwarding process	pfem	<p>Defines how routing protocols operate on the switch. The overall performance of the switch is largely determined by the effectiveness of the forwarding process.</p>
Interface process	dcd	<p>Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.</p>
Management process	mgd	<p>Provides communication between the other processes and an interface to the configuration database.</p> <p>Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured.</p> <p>Interacts with the other processes when commands are issued through one of the user interfaces on the switch.</p> <p>If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.</p>
Routing protocol process	rpd	<p>Defines how routing protocols such as RIP, OSPF, and BGP operate on the device, including selecting routes and maintaining forwarding tables.</p>

- Related Documentation**
- [For more information about processes, see \*Junos OS Network Operations Guide\*](#)
  - [For more information about basic system parameters, supported protocols, and software processes, see \*Junos OS System Basics Configuration Guide\*](#)

## Licenses Overview

- [Understanding Software Licenses for EX Series Switches on page 702](#)

## Understanding Software Licenses for EX Series Switches

To enable and use some of the Juniper Networks operating system (Junos OS) features, you must purchase, install, and manage separate software licenses. If the switch has the appropriate software license, you can configure and use these features.

The Junos OS feature license (that is, the purchased authorization code) is universal. However, to conform to Junos OS feature licensing requirements, you must install a unique license key (a combination of the authorization code and the switch's serial number) on each switch.

For a Virtual Chassis deployment, two license keys are recommended for redundancy—one for the device in the master role and the other for the device in the backup role:

- In an EX8200 Virtual Chassis, the devices in the master and backup roles are always XRE200 External Routing Engines.
- In all other Virtual Chassis, the devices in the master and backup roles are switches.

You do not need additional license keys for Virtual Chassis member switches that are in the linecard role or for the redundant Routing Engine (RE) modules or the redundant Switch Fabric and Routing Engine (SRE) modules in an EX8200 member switch.

This topic describes:

- [Purchasing a Software Feature License on page 702](#)
- [Features Requiring a License on EX2200 Switches on page 703](#)
- [Features Requiring a License on EX3300 Switches on page 704](#)
- [Features Requiring a License on EX4300 Switches on page 705](#)
- [Features Requiring a License on EX4600 Switches on page 706](#)
- [Features Requiring a License on EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 Switches on page 707](#)
- [License Warning Messages on page 708](#)

---

### Purchasing a Software Feature License

The following sections list features that require separate licenses. To purchase a software license, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show chassis hardware** command.



**NOTE:** You are required to provide the 12-digit serial number when purchasing a license for an XRE200 External Routing Engine in an EX8200 Virtual Chassis.

The serial number listed on the XRE200 External Routing Engine serial ID label is 16 digits long. Use the last 12 digits of the 16-digit serial number to purchase the license.

You can use the `show chassis hardware` command output to display the 12-digit serial number of the XRE200 External Routing Engine.

### Features Requiring a License on EX2200 Switches

For EX2200 switches, the following features can be added to basic Junos OS by installing an enhanced feature license (EFL):

- Bidirectional Forwarding Detection (BFD)
- Connectivity fault management (IEEE 802.1ag)
- IGMP (Internet Group Management Protocol) version 1 (IGMPv1), IGMPv2, and IGMPv3
- OSPFv1/v2 (with four active interfaces)
- Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode
- Q-in-Q tunneling (IEEE 802.1ad)
- Real-time performance monitoring (RPM)
- Virtual Router
- Virtual Router Redundancy Protocol (VRRP)

Table 65 on page 703 lists the EFLs that you can purchase for EX2200 switch models. If you have the license, you can run all of the enhanced software features mentioned above on your EX2200 switch.

**Table 65: Junos OS EFL Part Number on EX2200 Switches**

Switch Model	EFL Part Number
EX2200-C-12P-2G EX2200-C-12T-2G	EX-12-EFL
EX2200-24T-4G EX2200-24P-4G EX2200-24T-DC-4G	EX-24-EFL
EX2200-48T-4G EX2200-48P-4G	EX-48-EFL

### Features Requiring a License on EX3300 Switches

Two types of licenses are available on EX3300 switches: enhanced feature licenses (EFLs) and advanced feature licenses (AFLs).

To use the following features on the EX3300 switches, you must install an EFL:

- Bidirectional Forwarding Detection (BFD)
- IGMP (Internet Group Management Protocol) version 1 (IGMPv1), IGMPv2, and IGMPv3
- IPv6 routing protocols: Multicast Listener Discovery version 1 and 2 (MLD v1/v2), OSPFv3, PIM multicast, VRRPv6, virtual router support for unicast and filter-based forwarding (FBF)
- OSPFv1/v2
- Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode
- Q-in-Q tunneling (IEEE 802.1ad)
- Virtual Router
- Virtual Router Redundancy Protocol (VRRP)

[Table 66 on page 704](#) lists the EFLs that you can purchase for EX3300 switch models. If you have the license, you can run all of the enhanced software features mentioned above on your EX3300 switch.

**Table 66: Junos OS EFL Part Number on EX3300 Switches**

Switch Model	EFL Part Number
EX3300-24T EX3300-24P EX3300-24T-DC	EX-24-EFL
EX3300-48T EX3300-48T-BF EX3300-48P	EX-48-EFL

To use the following feature on EX3300 switches, you must install an AFL:

- Border Gateway Protocol (BGP) and multiprotocol BGP (MBGP)
- IPv6 routing protocols: IPv6 BGP and IPv6 for MBGP
- Virtual routing and forwarding (VRF) BGP

[Table 67 on page 705](#) lists the AFLs that you can purchase for EX3300 switch models. For EX3300 switches, you must purchase and install a corresponding EFL along with the AFL to enable the advanced license features. If you have both these licenses, you can run all of the advanced software features mentioned above on your EX3300 switch.



**Table 67: Junos OS AFL Part Number on EX3300 Switches**

Switch Model	AFL Part Number
EX3300-24T EX3300-24P EX3300-24T-DC	EX-24-AFL
EX3300-48T EX3300-48T-BF EX3300-48P	EX-48-AFL

### Features Requiring a License on EX4300 Switches

Two types of licenses are available on EX4300 switches: enhanced feature licenses (EFLs) and advanced feature licenses (AFLs).

To use the following features on the EX4300 switches, you must install an EFL:

- Bidirectional Forwarding Detection (BFD)
- Connectivity fault management (IEEE 802.1ag)
- IGMP (Internet Group Management Protocol) version 1 (IGMPv1), IGMPv2, and IGMPv3
- Multicast Source Discovery Protocol (MSDP)
- OSPFv2/v3
- Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode
- Real-time performance monitoring (RPM)
- RIPng (RIP next generation)
- Unicast reverse-path forwarding (RPF)
- Virtual Router
- Virtual Router Redundancy Protocol (VRRP)

[Table 68 on page 705](#) lists the EFLs that you can purchase for EX4300 switch models. If you have the license, you can run all of the enhanced software features mentioned above on your EX4300 switch.

**Table 68: Junos OS EFL Part Number on EX4300 Switches**

Switch Model	EFL Part Number
EX4300-24T EX4300-24P	EX4300-24-EFL
EX4300-48P EX4300-48T EX4300-48T-AFI EX4300-48T-DC EX4300-48T-DC-AFI	EX4300-48-EFL

**Table 68: Junos OS EFL Part Number on EX4300 Switches (*continued*)**

Switch Model	EFL Part Number
EX4300-32F EX4300-32F-DC	EX4300-32F-EFL

To use the following features on EX4300 switches, you must install an AFL:

- Border Gateway Protocol (BGP) and multiprotocol BGP (MBGP)
- Intermediate System-to-Intermediate System (IS-IS)

Table 69 on page 706 lists the AFLs that you can purchase for EX4300 switch models. For EX4300 switches, you must purchase and install a corresponding EFL along with the AFL to enable the advanced license features. If you have both these licenses, you can run all of the advanced software features mentioned above on your EX4300 switch.

**Table 69: Junos OS AFL Part Number on EX4300 Switches**

Switch Model	AFL Part Number
EX4300-24T EX4300-24P	EX4300-24-AFL
EX4300-48P EX4300-48T EX4300-48T-AFI EX4300-48T-DC EX4300-48T-DC-AFI	EX4300-48-AFL
EX4300-32F EX4300-32F-DC	EX4300-32F-AFL

You must download a MACsec feature license to enable MACsec. The MACsec feature license is an independent feature license; the enhanced feature licenses (EFLs) or advanced feature licenses (AFLs) that must be purchased to enable some features on EX Series switches cannot be purchased to enable MACsec.

To purchase a feature license for MACsec, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with a feature license file and a license key.

MACsec is supported on EX4300 switches.

### Features Requiring a License on EX4600 Switches

To use the following features on EX4300 switches, you must install an advanced feature license:

- Border Gateway Protocol (BGP) and multiprotocol BGP (MBGP)
- Intermediate System-to-Intermediate System (IS-IS)
- Multiprotocol Label Switching (MPLS)

[Table 70 on page 707](#) lists the AFLs that you can purchase for EX4300 switch models.

**Table 70: Junos OS AFL Part Number on EX4600 Switches**

Switch Model	AFL Part Number
EX4600-40F	EX4600-AFL

#### Features Requiring a License on EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 Switches

To use the following features on EX3200, EX4200, EX4500, EX4550, EX8200, and EX9200 switches, you must install an advanced feature license (AFL):

- Border Gateway Protocol (BGP) and multiprotocol BGP (MBGP)
- Intermediate System-to-Intermediate System (IS-IS)
- IPv6 routing protocols: IS-IS for IPv6, IPv6 BGP, IPv6 for MBGP
- Logical systems (available only on EX9200 switches)
- MPLS with RSVP-based label-switched paths (LSPs) and MPLS-based circuit cross-connects (CCCs) (Not supported on EX9200 switches)

To use the following features on Juniper Networks EX6200 Ethernet Switches, you must install an advanced feature license (AFL):

- Border Gateway Protocol (BGP)
- Intermediate System-to-Intermediate System (IS-IS)
- IPv6 routing protocols: IS-IS for IPv6, IPv6 BGP

[Table 71 on page 707](#) lists the AFLs that you can purchase for EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 switches. If you have the license, you can run all of the advanced software features mentioned above on your EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, or EX9200 switch.

**Table 71: Junos OS AFL Part Number on EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 Switches**

Switch Model	AFL Part Number
EX3200-24P EX3200-24T EX4200-24F EX4200-24P EX4200-24PX EX4200-24T	EX-24-AFL
EX3200-48P EX3200-48T EX4200-48F EX4200-48P EX4200-48PX EX4200-48T	EX-48-AFL

**Table 71: Junos OS AFL Part Number on EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 Switches (*continued*)**

Switch Model	AFL Part Number
EX4500-40F-BF EX4500-40F-BF-C EX4500-40F-FB EX4500-40F-FB-C	EX-48-AFL
EX4550	EX4550-AFL
EX6210	EX6210-AFL
EX8208	EX8208-AFL
EX8216	EX8216-AFL
EX-XRE200	EX-XRE200-AFL
EX9204	EX9204-AFL
EX9208	EX9208-AFL
EX9214	EX9214-AFL

You must download a MACsec feature license to enable MACsec. The MACsec feature license is an independent feature license; the enhanced feature licenses (EFLs) or advanced feature licenses (AFLs) that must be purchased to enable some features on EX Series switches cannot be purchased to enable MACsec.

To purchase a feature license for MACsec, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with a feature license file and a license key.

MACsec is currently supported on EX4200 and EX4550 switches.

### License Warning Messages

For using features that require a license, you must install and configure a license key. To obtain a license key, use the contact information provided in your certificate.

If you have not purchased the AFL or EFL and installed the license key, you receive warnings when you try to commit the configuration:

```
[edit protocols]
  'bgp'
    warning: requires 'bgp' license
error: commit failed: (statements constraint check failed)
```

The system generates system log (**syslog**) alarm messages notifying you that the feature requires a license—for example:

```
Sep 3 05:59:11 craftd[806]: Minor alarm set, BGP Routing Protocol usage
requires a license
Sep 3 05:59:11 alarmd[805]: Alarm set: License color=YELLOW, class=CHASSIS,
reason=BGP Routing Protocol usage requires a license
Sep 3 05:59:11 alarmd[805]: LICENSE_EXPIRED: License for feature bgp(47) expired
```

Output of the **show system alarms** command displays the active alarms:

```
user@switch> show system alarms
1 alarm currently active
Alarm time          Class  Description
2009-09-03 06:00:11 UTC  Minor  BGP Routing Protocol usage requires a license
```

#### Related Documentation

- [Managing Licenses for the EX Series Switch \(CLI Procedure\) on page 714](#)
- [Managing Licenses for the EX Series Switch \(J-Web Procedure\) on page 716](#)
- [Monitoring Licenses for the EX Series Switch on page 724](#)
- [License Key Components for the EX Series Switch on page 723](#)
- [EX Series Switch Software Features Overview](#)



## CHAPTER 18

# Configuration

- [Registering the Switch on page 711](#)
- [Booting the Switch on page 711](#)
- [Managing Licenses on page 714](#)

## Registering the Switch

---

- [Registering the EX Series Switch with the J-Web Interface on page 711](#)

## Registering the EX Series Switch with the J-Web Interface

---



**NOTE:** This topic applies only to the J-Web Application package.

---

You can register your EX Series switch with the J-Web interface so that you can request technical assistance as and when required. To register an EX Series switch:

1. In the J-Web interface, select **Maintain > Customer Support > Product Registration**. For an EX8200 Virtual Chassis configuration, select the member from the list.  
  
Note the serial number that is displayed.
2. Click **Register**. Enter the serial number in the page that is displayed.

### Related Documentation

- [EX Series Switch Software Features Overview](#)

## Booting the Switch

---

- [Booting an EX Series Switch Using a Software Package Stored on a USB Flash Drive on page 712](#)
- [Creating a Snapshot and Using It to Boot an EX Series Switch on page 713](#)

## Booting an EX Series Switch Using a Software Package Stored on a USB Flash Drive

There are two methods of getting Junos OS stored on a USB flash drive before using the software to boot the switch. You can pre-install the software onto the USB flash drive before inserting the USB flash drive into the USB port, or you can use the system snapshot feature to copy files from internal switch memory to the USB flash drive.

To move files into USB flash memory by using a system snapshot and use those files to boot the switch, see [“Creating a Snapshot and Using It to Boot an EX Series Switch” on page 713](#). We recommend that you use this method to boot the switch from a USB flash drive if your switch is running properly.

If you need to pre-install the software onto the USB flash drive, you can use the method described in this topic. Pre-installing Junos OS onto a USB flash drive to boot the switch can be done at any time and is particularly useful when the switch boots to the loader prompt because the switch cannot locate the Junos OS in internal flash memory.

Ensure that you have the following tools and parts available to boot the switch from a USB flash drive:

- A USB flash drive that meets the EX Series switch USB port specifications. See *USB Port Specifications for an EX Series Switch*.
- A computer or other device that you can use to download the software package from the Internet and copy it to the USB flash drive.

To download a Junos OS package onto a USB flash drive before inserting the USB flash drive:

1. Download the Junos OS package that you want to place onto the EX Series switch from the Internet onto the USB flash drive by using your computer or other device. See *Downloading Software Packages from Juniper Networks*.
2. Remove the USB flash drive from the computer or other device.
3. Insert the USB flash drive into the USB port on the switch.
4. This step can be performed only when the prompt for the loader script (**loader>**) is displayed. The loader script starts when the Junos OS loads but the CLI is not working for any reason or if the switch has no software installed.

Install the software package onto the switch:

```
loader> install source
```

where **source** represents the name and location of the Junos OS package on the USB flash drive. The Junos OS package on a flash drive is commonly stored in the root drive as the only file—for example, **file:///jinstall-ex-4200-9.4R1.5-domestic-signed.tgz**.

### Related Documentation

- *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*
- *Installing Software on EX Series Switches (J-Web Procedure)*
- [Understanding Software Installation on EX Series Switches on page 691](#)



- See *EX2200 Switches Hardware Overview* for USB port location.
- See *Rear Panel of an EX3200 Switch* for USB port location.
- See *Rear Panel of an EX3300 Switch* for USB port location.
- See *Rear Panel of an EX4200 Switch* for USB port location.
- See *EX4300 Switches Hardware Overview* for USB port location.
- See *Front Panel of an EX4500 Switch* for USB port location.
- See *EX4550 Switches Hardware Overview* for USB port location.
- See *Switch Fabric and Routing Engine (SRE) Module in an EX6200 Switch* for USB port location.
- See *Switch Fabric and Routing Engine (SRE) Module in an EX8208 Switch* for USB port location.
- See *Routing Engine (RE) Module in an EX8216 Switch* for USB port location.

## Creating a Snapshot and Using It to Boot an EX Series Switch

The system snapshot feature takes a “snapshot” of the files currently used to run the switch and copies them to an alternate storage location. You can then use this snapshot to boot the switch at the next bootup or as a backup boot option.

This topic includes the following tasks:

- [Creating a Snapshot on a USB Flash Drive and Using It to Boot the Switch on page 713](#)
- [Creating a Snapshot on an Internal Flash Drive and Using it to Boot the Switch on page 714](#)
- [Creating a Snapshot on the Alternate Slice of the Boot Media on page 714](#)

### Creating a Snapshot on a USB Flash Drive and Using It to Boot the Switch

You can create a snapshot on USB flash memory after a switch is booted by using files stored in internal memory.

Ensure that you have the following tools and parts available before creating a snapshot on a USB flash drive:

- A USB flash drive that meets the switch USB port specifications. See *USB Port Specifications for an EX Series Switch*.

To create a snapshot on USB flash memory and use it to boot the switch:

1. Place the snapshot into USB flash memory:  

```
user@switch> request system snapshot partition media external
```
2. (Optional) Perform this step if you want to boot the switch now using the snapshot stored on the USB flash drive. If you created the snapshot as a backup, do not perform this step.

- To reboot the switch using the most recently created snapshot:  
user@switch> **request system reboot media external**
- To reboot the switch using a snapshot in a specific partition on the USB flash drive:  
user@switch> **request system reboot media external slice alternate**

### Creating a Snapshot on an Internal Flash Drive and Using it to Boot the Switch

You can create a snapshot in internal memory after a switch is booted by using files stored in external memory.

To create a snapshot in internal memory and use it to boot the switch:

1. Place the snapshot files in internal memory:  
user@switch> **request system snapshot partition media internal**
2. (Optional) Perform this step if you want to boot the switch now using the newly created snapshot. If you created the snapshot as a backup, do not perform this step.
  - To reboot the switch using the most recently created snapshot:  
user@switch> **request system reboot media internal**
  - To reboot the switch using a snapshot in a specific partition in internal memory:  
user@switch> **request system reboot media internal slice alternate**

### Creating a Snapshot on the Alternate Slice of the Boot Media

The alternate slice of the boot media contains a backup software image that the switch can boot from if it is unable to boot from the primary slice. When you upgrade software, the new software image gets copied only to the primary slice of the boot media.

To create a snapshot of the currently booted software image on the backup slice of the boot media:

```
user@switch> request system reboot slice alternate
```

#### **Related Documentation**

- [Verifying That a System Snapshot Was Created on an EX Series Switch on page 719](#)
- [Understanding System Snapshot on EX Series Switches on page 694](#)

## Managing Licenses

- [Managing Licenses for the EX Series Switch \(CLI Procedure\) on page 714](#)
- [Managing Licenses for the EX Series Switch \(J-Web Procedure\) on page 716](#)

### **Managing Licenses for the EX Series Switch (CLI Procedure)**

To enable and use some Junos OS features on an EX Series switch, you must purchase, install, and manage separate software licenses. Each switch requires one license. For a Virtual Chassis deployment, two licenses are recommended for redundancy. After you

have configured the features, you see a warning message if the switch does not have a license for the feature.

Before you begin managing licenses, be sure that you have:

- Obtained the needed licenses. For information about how to purchase software licenses, contact your Juniper Networks sales representative.
- Understand what makes up a license key. For more information, see [“License Key Components for the EX Series Switch” on page 723](#).

This topic includes the following tasks:

- [Adding New Licenses on page 715](#)
- [Deleting Licenses on page 715](#)
- [Saving License Keys on page 715](#)

### Adding New Licenses

To add one or more new license keys on the switch, with the CLI:

1. Add the license key or keys:
  - To add one or more license keys from a file or URL, specify the filename of the file or the URL where the key is located:  

```
user@switch> request system license add filename | url
```
  - To add a license key from the terminal:  

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

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

### Deleting Licenses

To delete one or more license keys from the switch with the CLI, specify the license ID:

```
user@switch> request system license delete license-id
```

You can delete only one license at a time.

### Saving License Keys

To save the installed license keys to a file (which can be a URL) or to the terminal:

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

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

```
user@switch> request system license save ftp://user@switch/license.conf
```

#### Related Documentation

- [Managing Licenses for the EX Series Switch \(J-Web Procedure\) on page 716](#)
- [Monitoring Licenses for the EX Series Switch on page 724](#)

- [Understanding Software Licenses for EX Series Switches on page 702](#)

## Managing Licenses for the EX Series Switch (J-Web Procedure)

---



**NOTE:** This topic applies only to the J-Web Application package.

---

To enable and use some Junos OS features on an EX Series switch, you must purchase, install, and manage separate software licenses. Each switch requires one license. For a Virtual Chassis deployment, two licenses are recommended for redundancy. After you have configured the features, you see a warning message if the switch does not have a license for the feature.

Before you begin managing licenses, be sure that you have:

- Obtained the needed licenses. For information about how to purchase software licenses, contact your Juniper Networks sales representative.
- Understand what makes up a license key. For more information, see [“License Key Components for the EX Series Switch” on page 723](#).

This topic includes the following tasks:

- [Adding New Licenses on page 716](#)
- [Deleting Licenses on page 716](#)
- [Displaying License Keys on page 717](#)
- [Downloading Licenses on page 717](#)

### Adding New Licenses

---

To add one or more new license keys on the switch, with the J-Web license manager:

1. In the J-Web interface, select **Maintain > Licenses**.
2. Under Installed Licenses, click **Add** to add a new license key or keys.
3. Do *one* of the following, using a blank line to separate multiple license keys:
  - In the License File URL box, type the full URL to the destination file containing the license key or keys to be added.
  - In the License Key Text box, paste the license key text, in plain-text format, for the license to be added.
4. Click **OK** to add the license key or keys.

A list of features that use the license key is displayed. The table also lists the ID, state, and version of the license key.

### Deleting Licenses

---

To delete one or more license keys from a switch with the J-Web license manager:

1. In the J-Web interface, select **Maintain > Licenses**.
2. Select the check box of the license or licenses you want to delete.
3. Click **Delete**.

---

### Displaying License Keys

To display the license keys installed on a switch with the J-Web license manager:

1. In the J-Web interface, select **Maintain > Licenses**.
2. Under Installed Licenses, click **Display Keys** to display all the license keys installed on the switch.

A screen displaying the license keys in text format appears. Multiple licenses are separated by a blank line.

---

### Downloading Licenses

To download the license keys installed on the switch with the J-Web license manager:

1. In the J-Web interface, select **Maintain > Licenses**.
2. Under Installed Licenses, click **Download Keys** to download all the license keys installed on the switch to a single file.
3. Select **Save it to disk** and specify the file to which the license keys are to be written. You can also download the license file to your system.

#### Related Documentation

- [Managing Licenses for the EX Series Switch \(CLI Procedure\) on page 714](#)
- [Monitoring Licenses for the EX Series Switch on page 724](#)
- [Understanding Software Licenses for EX Series Switches on page 702](#)



## CHAPTER 19

# Administration

- [Routine Monitoring on page 719](#)
- [Monitoring Licenses on page 723](#)
- [Operational Commands on page 725](#)

## Routine Monitoring

---

- [Verifying That a System Snapshot Was Created on an EX Series Switch on page 719](#)
- [Verifying Junos OS and Boot Loader Software Versions on an EX Series Switch on page 720](#)

### Verifying That a System Snapshot Was Created on an EX Series Switch

**Purpose** Verify that a system snapshot was created with the proper files on an EX Series switch.

**Action** View the snapshot:

```
user@switch> show system snapshot media external
Information for snapshot on      external (/dev/dar1s1a) (backup)
Creation date: Mar 19 03:37:18 2012
JUNOS version on snapshot:
  jbase   : ex-12.1I20120111_0048_user
  jcrypto-ex: 12.1I20120111_0048_user
  jdocs-ex: 12.1I20120111_0048_user
  jroute-ex: 12.1I20120111_0048_user
  jswitch-ex: 12.1I20120111_0048_user
  jweb-ex: 12.1I20120111_0048_user
Information for snapshot on      external (/dev/dar1s2a) (primary)
Creation date: Mar 19 03:38:25 2012
JUNOS version on snapshot:
  jbase   : ex-12.2I20120305_2240_user
  jcrypto-ex: 12.2I20120305_2240_user
  jdocs-ex: 12.2I20120305_2240_user
  jroute-ex: 12.2I20120305_2240_user
  jswitch-ex: 12.2I20120305_2240_user
  jweb-ex: 12.2I20120305_2240_user
```

**Meaning** The output shows the date and time when the snapshot was created and the packages that are part of the snapshot. Check to see that the date and time match the time when you created the snapshot.

You can compare the output of this command to the output of the **show system software** command to ensure that the snapshot contains the same packages as the software currently running the switch.

#### Related Documentation

- [Creating a Snapshot and Using It to Boot an EX Series Switch on page 713](#)

## Verifying Junos OS and Boot Loader Software Versions on an EX Series Switch

Before or after upgrading or downgrading Junos OS, you might need to verify the Junos OS version. You might also need to verify the boot loader software version if you are upgrading to or downgrading from a release that supports resilient dual-root partitions (Junos OS Release 10.4R3 and later).

This topic includes:

- [Verifying the Number of Partitions and File System Mountings on page 720](#)
- [Verifying the Loader Software Version on page 721](#)
- [Verifying Which Root Partition Is Active on page 722](#)
- [Verifying the Junos OS Version in Each Root Partition on page 722](#)

### Verifying the Number of Partitions and File System Mountings

**Purpose** Between Junos OS Release 10.4R2 and Release 10.4R3, upgrades were made to further increase resiliency of root partitions, which required reformatting the disk from three partitions to four partitions. If your switch is running Release 10.4R2 or earlier, it has three partitions, and if it is running Release 10.4R3 or later, it has four partitions.

**Action** Verify how many partitions the disk has, as well as where each file system is mounted, by using the following command:

```
user@switch> show system storage
fpc0:
```

```
-----
Filesystem  Size  Used  Avail  Capacity  Mounted on
/dev/da0s1a 184M  124M   45M    73%      /
devfs       1.0K   1.0K    0B    100%    /dev
/dev/md0     37M   37M    0B    100%    /packages/mnt/jbase
/dev/md1     18M   18M    0B    100%
/packages/mnt/jcrypto-ex-10.4I20110121_0509_hbRPSRLI15184421081
/dev/md2     6.1M   6.1M    0B    100%
/packages/mnt/jdocs-ex-10.4I20110121_0509_hbRPSRLI15184421081
/dev/md3    154M  154M    0B    100%
/packages/mnt/jkernel-ex-10.4I20110121_0509_hbRPSRLI15184421081
/dev/md4     23M   23M    0B    100%
/packages/mnt/jpfe-ex42x-10.4I20110121_0509_hbRPSRLI15184421081
/dev/md5     46M   46M    0B    100%
/packages/mnt/jroute-ex-10.4I20110121_0509_hbRPSRLI15184421081
/dev/md6     28M   28M    0B    100%
/packages/mnt/jswitch-ex-10.4I20110121_0509_hbRPSRLI15184421081
/dev/md7     22M   22M    0B    100%
/packages/mnt/jweb-ex-10.4I20110121_0509_hbRPSRLI15184421081
/dev/md8    126M  10.0K  116M     0%    /tmp
/dev/da0s3e 123M   632K  112M     1%    /var
```



```

/dev/da0s3d 369M 20K 339M 0% /var/tmp
/dev/da0s4d 62M 62K 57M 0% /config
/dev/md9 118M 12M 96M 11% /var/rundb
procfs 4.0K 4.0K 0B 100% /proc
/var/jail/etc 123M 632K 112M 1%
/packages/mnt/jweb-ex-10.4I20110121_0509_hbRPSRLI15184421081/jail/var/etc
/var/jail/run 123M 632K 112M 1%
/packages/mnt/jweb-ex-10.4I20110121_0509_hbRPSRLI15184421081/jail/var/run
/var/jail/tmp 123M 632K 112M 1%
/packages/mnt/jweb-ex-10.4I20110121_0509_hbRPSRLI15184421081/jail/var/tmp
/var/tmp 369M 20K 339M 0%
/packages/mnt/jweb-ex-10.4I20110121_0509_hbRPSRLI15184421081/jail/var/tmp/uploads
devfs 1.0K 1.0K 0B 100%
/packages/mnt/jweb-ex-10.4I20110121_0509_hbRPSRLI15184421081/jail/dev

```

**Meaning** The presence of the partition name containing **s4d** indicates that there is a fourth slice. If this were a three-slice partition scheme, in place of **s1a**, **s3e**, **s3d**, and **s4d**, you would see **s1a**, **s1f**, **s2a**, **s2f**, **s3d**, and **s3e** and you would not see **s4d**.

### Verifying the Loader Software Version

**Purpose** For the special case of upgrading from Junos OS Release 10.4R2 or earlier to Release 10.4R3 or later, you must upgrade the loader software.

**Action** For EX Series switches except EX8200 switches:

```

user@switch> show chassis firmware
Part          Type      Version
FPC 0         uboot     U-Boot 1.1.6 (Jan  3 2011 - 16:14:58) 1.0.0

              loader   FreeBSD/PowerPC U-Boot bootstrap loader 2.4

```

For EX8200 switches:

```

user@switch> show chassis firmware
Part          Type      Version
FPC 0         uboot     U-Boot 1.1.6 (Jan  3 2011 - 16:14:58) 3.5.0

              loader   FreeBSD/PowerPC U-Boot bootstrap loader 2.4

```

**Meaning** For EX Series switches other than EX8200 switches, with Junos OS Release 10.4R3 or later installed:

- If there is version information following the timestamp for **U-Boot** (**1.0.0** in the preceding example), then the loader software does not require upgrading.
- If there is no version number following the timestamp for **U-boot**, then the loader software requires upgrading.



**NOTE:** If the software version is Release 10.4R2 or earlier, no version number is displayed following the timestamp for **U-boot**, regardless of the loader software version installed. If you do not know whether you have installed the new loader software, we recommend that you upgrade the loader software when you upgrade the software version.

For EX8200 switches, if the version number following the timestamp for **U-Boot** is earlier than **3.5.0**, you must upgrade the loader software when you upgrade the software version.

### Verifying Which Root Partition Is Active

---

**Purpose** Switches running Release 10.4R3 or later have resilient dual-root partition functionality, which includes the ability to boot transparently from the inactive partition if the system fails to boot from the primary root partition.

You can verify which root partition is active using the following command:

**Action** user@switch> **show system storage partitions**  
fpc0:

```
-----  
Boot Media: internal (da0)  
Active Partition: da0s1a  
Backup Partition: da0s2a  
Currently booted from: active (da0s1a)
```

```
Partitions information:  
Partition Size Mountpoint  
s1a      184M /  
s2a      184M altroot  
s3d      369M /var/tmp  
s3e      123M /var  
s4d       62M /config  
s4e             unused (backup config)
```

**Meaning** The **Currently booted from:** field shows which root partition is active.

### Verifying the Junos OS Version in Each Root Partition

---

**Purpose** Each switch contains two root partitions. We recommend that you copy the same Junos OS version in each partition when you upgrade. In Junos OS Release 10.4R2 and earlier, you might choose to have different Junos OS release versions in each partition. You might have different versions during a software upgrade and before you have finished verifying the new software installation. To enable a smooth reboot if corruption is found in the primary root file system, ensure that the identical Junos OS images are in each root partition. For Release 10.4R2 and earlier, you must manually reboot the switch from the backup root partition. However, for Release 10.4R3 and later, the switch reboots automatically from the backup root partition if it fails to reboot from the active root partition.

**Action** Verify whether both root partitions contain the same image by using the following command:

```
user@switch> show system snapshot media internal  
Information for snapshot on      internal (/dev/da0s1a) (backup)  
Creation date: Jan 11 03:02:59 2012  
JUNOS version on snapshot:  
jbase : ex-12.2I20120305_2240_user  
jcrypto-ex: 12.2I20120305_2240_user  
jdocs-ex: 12.2I20120305_2240_user
```

```

jroute-ex: 12.2I20120305_2240_user
jswitch-ex: 12.2I20120305_2240_user
jweb-ex: 12.2I20120305_2240_user
Information for snapshot on      internal (/dev/da0s2a) (primary)
Creation date: Mar 6 02:24:08 2012
JUNOS version on snapshot:
jbase   : ex-12.2I20120305_2240_user
jcrypto-ex: 12.2I20120305_2240_user
jdocs-ex: 12.2I20120305_2240_user
jroute-ex: 12.2I20120305_2240_user
jswitch-ex: 12.2I20120305_2240_user
jweb-ex: 12.2I20120305_2240_user

```

**Meaning** The command shows which Junos OS version is installed on each media partition. Verify that the same version is installed on both partitions.

- Related Documentation**
- [Troubleshooting Software Installation on page 785](#)
  - [Troubleshooting a Switch That Has Booted from the Backup Junos OS Image on page 788](#)
  - [Understanding Resilient Dual-Root Partitions on Switches on page 695](#)
  - [Resilient Dual-Root Partitions Frequently Asked Questions](#)

## Monitoring Licenses

- [License Key Components for the EX Series Switch on page 723](#)
- [Monitoring Licenses for the EX Series Switch on page 724](#)

### License Key Components for the EX Series Switch

When you purchase a license for a Junos OS feature that requires a separate license, you receive a license key.

A license key consists of two parts:

- License ID—Alphanumeric string that uniquely identifies the license key. When a license is generated, it is given a license ID.
- License data—Block of binary data that defines and stores all license key objects.

For example, in the following typical license key, the string **Junos204558** is the license ID, and the trailing block of data is the license data:

```

Junos204558 aeaqea qmijhd amrqha ztfmbu gqzama uqceds
ra32zr lsevik ftvjed o4jy5u fynzzj mgviyl
kgioyf ardb5g sj7wnt rsfkd wbjf5a sg

```

The license data defines the device ID for which the license is valid and the version of the license.

- Related Documentation**
- [Managing Licenses for the EX Series Switch \(CLI Procedure\) on page 714](#)
  - [Managing Licenses for the EX Series Switch \(J-Web Procedure\) on page 716](#)

- [Understanding Software Licenses for EX Series Switches on page 702](#)

## Monitoring Licenses for the EX Series Switch

To enable and use some Junos OS features on the EX Series switch, you must purchase, install, and manage the appropriate software licenses. Each switch requires one license. For a Virtual Chassis deployment, two licenses are recommended for redundancy.

To monitor your installed licenses, perform the following tasks:

- [Displaying Installed Licenses and License Usage Details on page 724](#)
- [Displaying Installed License Keys on page 725](#)

### Displaying Installed Licenses and License Usage Details

**Purpose** Verify that the expected license is installed and active on the switch and fully covers the switch configuration.

**Action** From the CLI, enter the **show system license** command. (To display only the **License usage** list, enter the **show system license usage** command. To display only the **Licenses installed** output, enter **show system license installed**.)

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

Feature name	Licenses	Licenses	Licenses	Expiry
	used	installed	needed	
bgp	1	1	0	permanent
isis	0	1	0	permanent
ospf3	0	1	0	permanent
ripng	0	1	0	permanent
mpls	0	1	0	permanent

Licenses installed:

```
License identifier: JUNOS204558
```

```
License version: 2
```

```
Valid for device: BN0208380000
```

```
Features:
```

```
ex-series - Licensed routing protocols in ex-series
```

```
permanent
```

**Meaning** The output shows the license or licenses (for Virtual Chassis deployments) installed on the switch and license usage. Verify the following information:

- If a feature that requires a license is configured (used), a license is installed on the switch. The **Licenses needed** column must show that no licenses are required.
- The appropriate number of licenses is installed. Each switch requires one license. For a Virtual Chassis deployment, two licenses are recommended for redundancy.
- The expected license is installed.

### Displaying Installed License Keys

**Purpose** Verify that the expected license keys are installed on the switch.

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

```
user@switch> show system license keys
JUNOS204558 abcdef qhijkl mnopqr stuvwxyz efghij
               klmnop qrstuv wxyzab cdefgh ijklmn opqrst
               uvwxyz 61abcd efgh21 31efgh yzabcd
```

**Meaning** The output shows the license key or keys (for Virtual Chassis deployments) installed on the switch. Verify that each expected license key is present.

- Related Documentation**
- [Managing Licenses for the EX Series Switch \(CLI Procedure\) on page 714](#)
  - [Managing Licenses for the EX Series Switch \(J-Web Procedure\) on page 716](#)
  - [Understanding Software Licenses for EX Series Switches on page 702](#)

## Operational Commands

- [request system license add](#)
- [request system license delete](#)
- [request system license save](#)
- [request system reboot](#)
- [request system reboot](#)
- [request system snapshot](#)
- [request system software add](#)
- [request system software delete](#)
- [request system software rollback](#)
- [request system software validate](#)
- [show system auto-snapshot](#)
- [show system boot-messages](#)
- [show system license](#)

- [show system snapshot](#)
- [show system storage partitions \(EX Series Switches Only\)](#)

## request system license add

<b>Syntax</b>	<code>request system license add (<i>filename</i>   terminal)</code>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 9.5 for SRX Series devices.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Add a license key.
<b>Options</b>	<p><b><i>filename</i></b>—License key from a file or URL. Specify the filename or the URL where the key is located.</p> <p><b><i>terminal</i></b>—License key from the terminal.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Adding New Licenses (CLI Procedure)</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system license add on page 727</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system license add

```

user@host> request system license add terminal
E408408918 aeaqib qcsbja okbuqe rcmxnq vjocwf uxfsta
           z5ufjb kdrmt6 57bimv 2f3ddp qttcdn 627q4a
           jx4s5x hiri
E408408918: successfully added
add license complete (no errors)

```

## request system license delete

---

<b>Syntax</b>	<code>request system license delete ( <i>license-identifier</i>   license-identifier-list [ <i>licenseid001</i> <i>licenseid002</i> <i>licenseid003</i> ]   all )</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Option <b>license-identifier-list</b> introduced in Junos OS Release 13.1.
<b>Description</b>	Delete a license key. You can choose to delete one license at a time, all licenses at once, or a list of license identifiers enclosed in brackets.
<b>Options</b>	<b><i>license-identifier</i></b> —Text string that uniquely identifies a license key.  <b>license-identifier-list [ <i>licenseid001</i> <i>licenseid002</i> <i>licenseid003</i>.... ]</b> —Delete multiple license identifiers as a list enclosed in brackets.  <b>all</b> —Delete all licenses on the device.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Deleting a License (CLI Procedure)</i></li></ul>



## request system license save

---

<b>Syntax</b>	<code>request system license save (<i>filename</i>   terminal)</code>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 9.5 for SRX Series devices.</p>
<b>Description</b>	Save installed license keys to a file or URL.
<b>Options</b>	<p><b><i>filename</i></b>—License key from a file or URL. Specify the filename or the URL where the key is located.</p> <p><b><i>terminal</i></b>—License key from the terminal.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Saving License Keys</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system license save on page 729</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system license save

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

## request system reboot

---

<b>List of Syntax</b>	<a href="#">Syntax on page 730</a> <a href="#">Syntax (EX Series Switches) on page 730</a> <a href="#">Syntax (TX Matrix Router) on page 730</a> <a href="#">Syntax (TX Matrix Plus Router) on page 730</a> <a href="#">Syntax (MX Series Router) on page 730</a>
<b>Syntax</b>	<code>request system reboot</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;media (compact-flash   disk   removable-compact-flash   usb)&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code> <code>&lt;other-routing-engine&gt;</code>
<b>Syntax (EX Series Switches)</b>	<code>request system reboot</code> <code>&lt;all-members&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;local&gt;</code> <code>&lt;media (external   internal)&gt;</code> <code>&lt;member <i>member-id</i>&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code> <code>&lt;other-routing-engine&gt;</code> <code>&lt;slice <i>slice</i>&gt;</code>
<b>Syntax (TX Matrix Router)</b>	<code>request system reboot</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;media (compact-flash   disk)&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code> <code>&lt;other-routing-engine&gt;</code>
<b>Syntax (TX Matrix Plus Router)</b>	<code>request system reboot</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;media (compact-flash   disk)&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code> <code>&lt;other-routing-engine&gt;</code> <code>&lt;partition (1   2   alternate)&gt;</code>
<b>Syntax (MX Series Router)</b>	<code>request system reboot</code> <code>&lt;all-members&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;local&gt;</code>

```

<media (external | internal)>
<member member-id>
<message "text">
<other-routing-engine>

```

<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Option <b>other-routing-engine</b> introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>sfc</b> introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Option <b>both-routing-engines</b> introduced in Junos OS Release 12.1.</p>
<b>Description</b>	Reboot the software.
<b>Options</b>	<p><b>none</b>—Reboot the software immediately.</p> <p><b>all-chassis</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all routers connected to the TX Matrix or TX Matrix Plus router, respectively.</p> <p><b>all-lcc</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively.</p> <p><b>all-members</b>—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on all members of the Virtual Chassis configuration.</p> <p><b>at <i>time</i></b>—(Optional) Time at which to reboot the software, specified in one of the following ways:</p> <ul style="list-style-type: none"> <li>• <b>now</b>—Stop or reboot the software immediately. This is the default.</li> <li>• <b>+<i>minutes</i></b>—Number of minutes from now to reboot the software.</li> <li>• <b><i>yymmddhhmm</i></b>—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.</li> <li>• <b><i>hh:mm</i></b>—Absolute time on the current day at which to stop the software, specified in 24-hour time.</li> </ul> <p><b>both-routing-engines</b>—(Optional) Reboot both Routing Engines at the same time.</p> <p><b>in <i>minutes</i></b>—(Optional) Number of minutes from now to reboot the software. This option is an alias for the <b>at +<i>minutes</i></b> option.</p> <p><b>lcc <i>number</i></b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis number.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> <li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li> <li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li> </ul>

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the local Virtual Chassis member.

**media (compact-flash | disk | removable-compact-flash | usb)**—(Optional) Boot medium for next boot. (The options **removable-compact-flash** and **usb** pertain to the J Series routers only.)

**media (external | internal)**—(EX Series switches and MX Series routers only) (Optional) Reboot the boot media:

- **external**—Reboot the external mass storage device.
- **internal**—Reboot the internal flash device.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

**message "*text*"**—(Optional) Message to display to all system users before stopping or rebooting the software.

**other-routing-engine**—(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is rebooted.

**partition**—(TX Matrix Plus routers only) (Optional) Reboot using the specified partition on the boot media. This option has the following suboptions:

- 1—Reboot from partition 1.
- 2—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition.

**scc**—(TX Matrix routers only) (Optional) Reboot the Routing Engine on the TX Matrix switch-card chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Reboot the Routing Engine on the TX Matrix Plus switch-fabric chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted. Replace ***number*** with 0.

**slice *slice***—(EX Series switches only) (Optional) Reboot a partition on the boot media. This option has the following suboptions:

- 1—Power off partition 1.
- 2—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

**Additional Information** Reboot requests are recorded in the system log files, which you can view with the **show log** command (see [show log](#)). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the **show system processes** command (see [show system processes](#)).

On a TX Matrix or TX Matrix Plus router, if you issue the **request system reboot** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are rebooted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are rebooted.



**NOTE:** Before issuing the **request system reboot** command on a TX Matrix Plus router with no options or the **all-chassis**, **all-lcc**, **lcc number**, or **sfc** options, verify that master Routing Engine for all routers in the routing matrix are in the same slot number. If the master Routing Engine for a line-card chassis is in a different slot number than the master Routing Engine for a TX Matrix Plus router, the line-card chassis might become logically disconnected from the routing matrix after the **request system reboot** command.



**NOTE:** To reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the master Routing Engine.

**Required Privilege Level** maintenance

**Related Documentation**

- [clear system reboot on page 84](#)
- [request system halt on page 101](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

[request system reboot on page 734](#)  
[request system reboot \(at 2300\) on page 734](#)  
[request system reboot \(in 2 Hours\) on page 734](#)  
[request system reboot \(Immediately\) on page 734](#)  
[request system reboot \(at 1:20 AM\) on page 734](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system reboot

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

### request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes
```

```
shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

### request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
user@host> request system reboot in 120
user@host> request system reboot at 19:00
```

### request system reboot (Immediately)

```
user@host> request system reboot at now
```

### request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

## request system reboot

<b>Syntax</b>	<pre>request system reboot &lt;all-members   local   member member-id&gt; &lt;at time&gt; &lt;in minutes&gt; &lt;media (external   internal)&gt; &lt;message "text"&gt; &lt;slice (1   2   alternate)&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>partition</b> changed to <b>slice</b> in Junos OS Release 10.0 for EX Series switches.</p>
<b>Description</b>	<p>Reboot the Junos OS.</p> <p>Reboot requests are recorded in the system log files, which you can view with the <b>show log</b> command. You can view the process names with the <b>show system processes</b> command.</p>
<b>Options</b>	<p><b>none</b>—Reboots the software immediately.</p> <p><b>all-members   local   member member-id</b>—(EX4200 switch only) (Optional) Specify which member of the Virtual Chassis to reboot:</p> <ul style="list-style-type: none"> <li>• <b>all-members</b>—Reboots each switch that is a member of the Virtual Chassis.</li> <li>• <b>local</b>—Reboots the local switch, meaning the switch you are logged into, only.</li> <li>• <b>member member-id</b>—Reboots the specified member switch of the Virtual Chassis.</li> </ul> <p><b>at time</b>—(Optional) Time at which to reboot the software, specified in one of the following ways:</p> <ul style="list-style-type: none"> <li>• <b>+minutes</b>—Number of minutes from now to reboot the software.</li> <li>• <b>hh:mm</b>—Absolute time on the current day at which to reboot the software, specified in 24-hour time.</li> <li>• <b>now</b>—Stop or reboot the software immediately. This is the default.</li> <li>• <b>yymmddhhmm</b>—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.</li> </ul> <p><b>in minutes</b>—(Optional) Number of minutes from now to reboot the software. This option is an alias for the <b>at +minutes</b> option.</p> <p><b>media (external   internal)</b>—(Optional) Boot medium for the next boot. The external option reboots the switch using a software package stored on an external boot source, such as a USB flash drive. The internal option reboots the switch using a software package stored in an internal memory source.</p> <p><b>message "text"</b>—(Optional) Message to display to all system users before rebooting the software.</p>

**slice (1 | 2 | alternate)**—(Optional) Reboot using the specified partition on the boot media. This option has the following suboptions:

- **1**—Reboot from partition 1.
- **2**—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition, which is the partition that did not boot the switch at the last bootup.

**Required Privilege Level** maintenance

**Related Documentation**

- [clear system reboot on page 84](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system reboot

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

### request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes

shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

### request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
user@host> request system reboot in 120
user@host> request system reboot at 19:00
```

### request system reboot (Immediately)

```
user@host> request system reboot at now
```

### request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```



## request system snapshot

<b>List of Syntax</b>	<a href="#">Syntax on page 737</a> <a href="#">Syntax (ACX Series Routers) on page 737</a> <a href="#">Syntax (EX Series Switches) on page 737</a> <a href="#">Syntax (MX Series Routers) on page 737</a> <a href="#">Syntax (TX Matrix Routers) on page 737</a> <a href="#">Syntax (TX Matrix Plus Routers) on page 737</a>
<b>Syntax</b>	request system snapshot <partition>
<b>Syntax (ACX Series Routers)</b>	request system snapshot <media type> <partition>
<b>Syntax (EX Series Switches)</b>	request system snapshot <all-members   local   member <i>member-id</i> > <media type> <partition> <re0   re1   routing-engine <i>routing-engine-id</i> > <slice alternate>
<b>Syntax (MX Series Routers)</b>	request system snapshot <all-members> <config-partition> <local> <member <i>member-id</i> > <media <i>usb-port-number</i> > <partition> <root-partition>
<b>Syntax (TX Matrix Routers)</b>	request system snapshot <all-chassis   all-lcc   lcc <i>number</i>   scc> <config-partition> <partition> <root-partition>
<b>Syntax (TX Matrix Plus Routers)</b>	request system snapshot <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> > <config-partition> <partition> <root-partition>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series switches.</p> <p>Options &lt;config-partition&gt; and &lt;root-partition&gt; introduced in Junos OS Release 13.1 for M, MX, T, TX Series switches.</p> <p>Option <b>media <i>usb-port-number</i></b> introduced in Junos OS Release 13.2 for MX104 routers.</p>

- Description**
- On the router, back up the currently running and active file system partitions to standby partitions that are not running. Specifically, the root file system (/) is backed up to **/altroot**, and **/config** is backed up to **/altconfig**. The root and **/config** file systems are on the router's flash drive, and the **/altroot** and **/altconfig** file systems are on the router's hard drive.
  - On the switch, take a snapshot of the files currently used to run the switch—the complete contents of the root (/) , **/altroot**, **/config**, **/var**, and **/var-tmp** directories, which include the running Junos OS, the active configuration, and log files.



**CAUTION:** After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

**Options** The specific options available depend upon the router or switch:

**none**—Back up the currently running software as follows:

- On the router, back up the currently running and active file system partitions to standby partitions that are not running. Specifically, the root file system (/) is backed up to **/altroot**, and **/config** is backed up to **/altconfig**. The root and **/config** file systems are on the router's flash drive, and the **/altroot** and **/altconfig** file systems are on the router's hard drive.
- On the switch, take a snapshot of the files currently used to run the switch and copy them to the media that the switch did not boot from. If the switch is booted from internal media, the snapshot is copied to external (USB) media. If the switch is booted from external (USB) media, the snapshot is copied to internal media.
- If the snapshot destination is external media but a USB flash drive is not connected, an error message is displayed.
- If the automatic snapshot procedure is already in progress, the command returns the following error: **Snapshot already in progress. Cannot start manual snapshot.** For additional information about the automatic snapshot feature, see [“Understanding Resilient Dual-Root Partitions on Switches” on page 695](#).

**all-chassis | all-lcc | lcc number** —(TX Matrix and TX Matrix Plus router only) (Optional)

- **all-chassis**—On a TX Matrix router, archive data and executable areas for all Routing Engines in the chassis. On a TX Matrix Plus router, archive data and executable areas for all Routing Engines in the chassis.
- **all-lcc**—On a TX Matrix router, archive data and executable areas for all T640 routers (or line-card chassis) connected to a TX Matrix router. On a TX Matrix Plus router, archive data and executable areas for all routers (or line-card chassis) connected to a TX Matrix Plus router.
- **lcc number**—On a TX Matrix router, archive data and executable areas for a specific T640 router (or line-card chassis) that is connected to a TX Matrix router. On a

TX Matrix Plus router, archive data and executable areas for a specific router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**all-members | local | member *member-id***—(EX Series switch Virtual Chassis and MX Series routers only) (Optional) Specify where to place the snapshot (archive data and executable areas) in a Virtual Chassis:

- **all-members**—Create a snapshot (archive data and executable areas) for all members of the Virtual Chassis.
- **local**—Create a snapshot (archive data and executable areas) on the member of the Virtual Chassis that you are currently logged into.
- **member *member-id***—Create a snapshot (archive data and executable areas) for the specified member of the Virtual Chassis.

**config-partition**—(M, MX, T, TX Series routers only) Create a snapshot of the configuration partition only and store it onto the default **/altconfig** on the hard disk device or an **/altconfig** on a USB device.

**media type**—(ACX Series, M320, T640, MX960 routers, and EX Series switches only)(Optional) Specify the boot device the software is copied to:

- **compact-flash**—Copy software to the primary compact flash drive.
- **external**—(Switches only) Copy software to an external mass storage device, such as a USB flash drive. If a USB drive is not connected, the switch displays an error message.
- **internal**—Copy software to an internal flash drive.
- **removable-compact-flash**—Copy software to the removable compact flash drive.
- **usb**—(ACX Series, M320, T640, MX960 routers only) Copy software to the device connected to the USB port.
- **usb0**—(MX104 routers only) Copy software to the device connected to the USB0 port.
- **usb1**—(MX104 routers only) Copy software to the device connected to the USB1 port.

**partition**—(Optional) Repartition the flash drive before a snapshot occurs. If the partition table on the flash drive is corrupted, the **request system snapshot** command fails and reports errors. The partition option is only supported for restoring the software image from the hard drive to the flash drive.

(Routers only) You cannot issue the **request system snapshot** command when you enable flash disk mirroring. We recommend that you disable flash disk mirroring when you upgrade or downgrade the software. For more information, see the *Junos OS Administration Library for Routing Devices*.

(EX Series switches only) If the snapshot destination is the media that the switch did not boot from, you must use the **partition** option.

**re0 | re1 | routing-engine routing-engine-id**—(EX6200 and EX8200 switches only) Specify where to place the snapshot in a redundant Routing Engine configuration.

- **re0**—Create a snapshot on Routing Engine 0.
- **re1**—Create a snapshot on Routing Engine 1.
- **routing-engine routing-engine-id**—Create a snapshot on the specified Routing Engine.

**root-partition**—(M, MX, T, TX Series routers only) Create a snapshot of the root partition only and store it onto the default **/altroot** on the hard disk device or an **/altroot** on a USB device.

**slice alternate**—(EX Series switches only) (Optional) Take a snapshot of the active root partition and copy it to the alternate slice on the boot media.

**scc**—(TX Matrix router only) (Optional) Archive data and executable areas for a TX Matrix router (or switch-card chassis).

**sfc number**—(TX Matrix Plus router only) (Optional) Archive data and executable areas for a TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

#### Additional Information

- (Routers only) Before upgrading the software on the router, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the **/altroot** and **/altconfig** file systems. After you have upgraded the software on the router and are satisfied that the new packages are successfully installed and running, issue the **request system snapshot** command again to back up the new software to the **/altroot** and **/altconfig** file systems.
- (Routers only) You cannot issue the **request system snapshot** command when you enable flash disk mirroring. We recommend that you disable flash disk mirroring when you upgrade or downgrade the software. For more information, see the *Junos OS Administration Library for Routing Devices*.
- (TX Matrix and TX Matrix Plus router only) On a routing matrix, if you issue the **request system snapshot** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are backed up. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are backed up.

**Required Privilege Level** maintenance

**Related Documentation**

- [show system snapshot on page 779](#)
- [show system auto-snapshot on page 763](#)

**List of Sample Output**

- [request system snapshot \(Routers\) on page 741](#)
- [request system snapshot \(EX Series Switches\) on page 741](#)
- [request system snapshot \(When the Partition Flag Is On\) on page 741](#)
- [request system snapshot \(MX104 routers when media device is missing\) on page 742](#)
- [request system snapshot \(When Mirroring Is Enabled\) on page 742](#)
- [request system snapshot all-lcc \(Routing Matrix\) on page 742](#)
- [request system snapshot all-members \(Virtual Chassis\) on page 742](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system snapshot (Routers)

```
user@host> request system snapshot
umount: /altroot: not currently mounted
Copying / to /altroot.. (this may take a few minutes)
umount: /altconfig: not currently mounted
Copying /config to /altconfig.. (this may take a few minutes)

The following filesystems were archived: / /config
```

### request system snapshot (EX Series Switches)

```
user@switch> request system snapshot partition
Clearing current label...
Partitioning external media (/dev/da1) ...
Partitions on snapshot:

  Partition  Mountpoint  Size    Snapshot argument
  s1a       /altroot    179M    none
  s2a       /           180M    none
  s3d       /var/tmp    361M    none
  s3e       /var       121M    none
  s4d       /config     60M     none
Copying '/dev/da0s1a' to '/dev/da1s1a' .. (this may take a few minutes)
Copying '/dev/da0s2a' to '/dev/da1s2a' .. (this may take a few minutes)
Copying '/dev/da0s3d' to '/dev/da1s3d' .. (this may take a few minutes)
Copying '/dev/da0s3e' to '/dev/da1s3e' .. (this may take a few minutes)
Copying '/dev/da0s4d' to '/dev/da1s4d' .. (this may take a few minutes)
The following filesystems were archived: /altroot / /var/tmp /var /config
```

### request system snapshot (When the Partition Flag Is On)

```
user@host> request system snapshot partition
Performing preliminary partition checks ...
Partitioning ad0 ...
umount: /altroot: not currently mounted
Copying / to /altroot.. (this may take a few minutes)

The following filesystems were archived: / /config
```

### request system snapshot (MX104 routers when media device is missing)

```
user@host > request system snapshot media usb0
error: usb0 media missing or invalid
```

### request system snapshot (When Mirroring Is Enabled)

```
user@host> request system snapshot
Snapshot is not possible since mirror-flash-on-disk is configured.
```

### request system snapshot all-lcc (Routing Matrix)

```
user@host> request system snapshot all-lcc
lcc0-re0:
-----
Copying '/' to '/altroot' .. (this may take a few minutes)
Copying '/config' to '/altconfig' .. (this may take a few minutes)
The following filesystems were archived: / /config

lcc2-re0:
-----
Copying '/' to '/altroot' .. (this may take a few minutes)
Copying '/config' to '/altconfig' .. (this may take a few minutes)
The following filesystems were archived: / /config
```

### request system snapshot all-members (Virtual Chassis)

```
user@switch> request system snapshot all-members media internal
fpc0:
-----
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc1:
-----
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc2:
-----
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc3:
-----
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc4:
-----
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc5:
-----
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /
```

## request system software add

<b>List of Syntax</b>	<a href="#">Syntax on page 743</a> <a href="#">Syntax (EX Series Switches) on page 743</a> <a href="#">Syntax (TX Matrix Router) on page 743</a> <a href="#">Syntax (TX Matrix Plus Router) on page 743</a> <a href="#">Syntax (MX Series Router) on page 744</a> <a href="#">Syntax (QFX Series) on page 744</a>
<b>Syntax</b>	<pre>request system software add <i>package-name</i> &lt;best-effort-load&gt; &lt;delay-restart&gt; &lt;force&gt; &lt;no-copy&gt; &lt;no-validate&gt; &lt;re0   re1&gt; &lt;reboot&gt; &lt;set [<i>package-name package-name</i>]&gt; &lt;unlink&gt; &lt;upgrade-with-config&gt; &lt;upgrade-with-config-format <i>format</i>&gt; &lt;validate&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>request system software add <i>package-name</i> &lt;best-effort-load&gt; &lt;delay-restart&gt; &lt;force&gt; &lt;no-copy&gt; &lt;no-validate&gt; &lt;re0   re1&gt; &lt;reboot&gt; &lt;set [<i>package-name package-name</i>]&gt; &lt;upgrade-with-config&gt; &lt;upgrade-with-config-format <i>format</i>&gt; &lt;validate&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>request system software add <i>package-name</i> &lt;best-effort-load&gt; &lt;delay-restart&gt; &lt;force&gt; &lt;lcc <i>number</i>   scc&gt; &lt;no-copy&gt; &lt;no-validate&gt; &lt;re0   re1&gt; &lt;reboot&gt; &lt;set [<i>package-name package-name</i>]&gt; &lt;unlink&gt; &lt;upgrade-with-config&gt; &lt;upgrade-with-config-format <i>format</i>&gt; &lt;validate&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>request system software add <i>package-name</i> &lt;best-effort-load&gt;</pre>

```
<delay-restart>
<force>
<lcc number | sfc number>
<no-copy>
<no-validate>
<re0 | re1>
<reboot>
<set [package-name package-name]>
<unlink>
<upgrade-with-config>
<upgrade-with-config-format format>
<validate>
```

**Syntax (MX Series Router)**

```
request system software add package-name
<best-effort-load>
<delay-restart>
<force>
<member member-id>
<no-copy>
<no-validate>
<re0 | re1>
<reboot>
<set [package-name package-name]>
<unlink>
<upgrade-with-config>
<upgrade-with-config-format format>
<validate>
```

**Syntax (QFX Series)**

```
request system software add package-name
<best-effort-load>
<component all>
<delay-restart>
<force>
<force-host>
<no-copy>
<no-validate>
<partition>
<reboot>
<unlink>
<upgrade-with-config>
<upgrade-with-config-format format>
<validate>
```

**Release Information**

Command introduced before Junos OS Release 7.4.

**best-effort-load** and **unlink** options added in Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

**sfc** option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Command introduced in Junos OS Release 11.1 for the QFX Series.

**set [*package-name package-name*]** option added in Junos OS Release 11.1 for EX Series switches.

**set [*package-name package-name*]** option added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways.





**NOTE:** On EX Series switches, the set `[package-name package-name]` option allows you to install only two software packages on a mixed EX4200 and EX4500 Virtual Chassis, whereas, on M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways, the set `[package-name package-name]` option allows you to install multiple software packages and software add-on packages at the same time.

`upgrade-with-config` and `upgrade-with-config-format` *format* options added in Junos OS Release 12.3 for M Series routers, MX Series routers, T Series routers, EX Series Ethernet switches, and QFX Series devices.

#### Description



**NOTE:** We recommend that you always download the software image to `/var/tmp` only. On EX Series and QFX Series switches, you must use the `/var/tmp` directory. Other directories are not supported.

Install a software package or bundle on the router or switch.



**WARNING:** Any configuration changes performed after inputting the `request system software add` command will be lost when the system reboots with an upgraded version of JUNOS.

**Options** `package-name`—Location from which the software package or bundle is to be installed.

For example:

- `/var/tmp/package-name`—For a software package or bundle that is being installed from a local directory on the router or switch.
- `protocol://hostname/pathname/package-name`—For a software package or bundle that is to be downloaded and installed from a remote location. Replace *protocol* with one of the following:
  - **ftp**—File Transfer Protocol.  
Use `ftp://hostname/pathname/package-name`. To specify authentication credentials, use `ftp://<username>:<password>@hostname/pathname/package-name`. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
  - **http**—Hypertext Transfer Protocol.  
Use `http://hostname/pathname/package-name`. To specify authentication credentials, use `http://<username>:<password>@hostname/pathname/package-name`. If a password is required and you omit it, you are prompted for it.

- **scp**—Secure copy (available only for Canada and U.S. version).  
Use **scp://hostname/pathname/package-name**. To specify authentication credentials, use  
**scp://<username>:<password>@hostname/pathname/package-name**.

**NOTE:**

- The **pathname** in the protocol is the relative path to the user's home directory on the remote system and not the root directory.
- Do not use the **scp** protocol in the **request system software add** command to download and install a software package or bundle from a remote location. The previous statement does not apply to the QFabric switch. The software upgrade is handled by the MGD process which does not support **scp**.  
Use the **file copy** command to copy the software package or bundle from the remote location to the **/var/tmp** directory on the hard disk:  
**file copy scp://source/package-name /var/tmp**  
Then install the software package or bundle using the **request system software add** command:  
**request system software add /var/tmp/package-name**
- On a J Series Services Router, when you install the software from a remote location, the package is removed at the earliest opportunity in order to make room for the installation to be completed. If you copy the software to a local directory on the router and then install the new package, use the **unlink** option to achieve the same effect and allow the installation to be completed.

---

**best-effort-load**—(Optional) Activate a partial load and treat parsing errors as warnings instead of errors.

**component all**—(QFabric systems only) (Optional) Install software package on all of the QFabric components.

**delay-restart**—(Optional) Install a software package or bundle, but do not restart software processes.

**force**—(Optional) Force the addition of the software package or bundle (ignore warnings).

**force-host**—(Optional) Force the addition of host software package or bundle (ignore warnings) on the QFX5100 device.

**lcc number** —(TX Matrix routers and TX Matrix Plus routers only) (Optional) In a routing matrix based on the TX Matrix router, install a software package or bundle on a T640 router that is connected to the TX Matrix router. In a routing matrix based on the TX Matrix Plus router, install a software package or bundle on a router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**member *member-id***—(MX Series routers only) (Optional) Install a software package on the specified Virtual Chassis member. Replace *member-id* with a value of 0 or 1.

**partition**—(QFX3500 switches only) (Optional) Format and repartition the media before installation.

**scc**—(TX Matrix routers only) (Optional) Install a software package or bundle on a Routing Engine on a TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Install a software package or bundle on a Routing Engine on a TX Matrix Plus router. Replace *number* with 0.

**no-copy**—(Optional) Install a software package or bundle, but do not save copies of the package or bundle files.

**no-validate**—(Optional) When loading a software package or bundle with a different release, suppress the default behavior of the **validate** option.

**re0 | re1**—(Optional) On routers or switches that support dual or redundant Routing Engines, load a software package or bundle on the Routing Engine in slot 0 (re0) or the Routing Engine in slot 1 (re1).

**reboot**—(Optional) After adding the software package or bundle, reboot the system. On a QFabric switch, the software installation is not complete until you reboot the component for which you have installed the software.

**set [*package-name package-name*]**—(Mixed EX4200 and EX4500 Virtual Chassis only) (Optional) Install two software packages—a package for an EX4200 switch and the same release of the package for an EX4500 switch—to upgrade all member switches in a mixed EX4200 and EX4500 Virtual Chassis.

**set [*package-name package-name*]**—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages and software add-on packages at the same time.

**unlink**—(Optional) On J Series Services Routers, this option ensures that the software package is removed at the earliest opportunity in order to make room for the installation to be completed. On M Series, T Series, and MX Series routers, use the

**unlink** option to remove the software package from this directory after a successful upgrade is completed.

**upgrade-with-config**—(Optional) Install one or more configuration files.

**upgrade-with-config-format *format***—(Optional) Specify the configuration file format, **text** or **xml**. The default format is **text**.



**NOTE:** The **upgrade-with-config** and **upgrade-with-config-format** options are only available locally on the router or switch. In a routing matrix, the configuration is applied only to the local router and is not propagated to other routers.

The options are validated during the validation process and applied to the router or switch during the upgrade process. If the upgrade process is successful, the options are removed from the configuration. If the upgrade process fails, the configuration file is renamed with the **.failed** suffix.

**validate**—(Optional) Validate the software package or bundle against the current configuration as a prerequisite to adding the software package or bundle. This is the default behavior when the software package or bundle being added is a different release.



**NOTE:** The **validate** option only works on systems that do not have **graceful-switchover (GRES)** enabled. To use the **validate** option on a system with GRES, either disable GRES for the duration of the installation, or install using the command **request system software in-service-upgrade**, which requires nonstop active routing (NSR) to be enabled when using GRES.

#### Additional Information

Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the **/altroot** and **/altconfig** file systems. After you have upgraded the software on the router or switch and are satisfied that the new package or bundle is successfully installed and running, issue the **request system snapshot** command again to back up the new software to the **/altroot** and **/altconfig** file systems.



**NOTE:** The **request system snapshot** command is currently not supported on the QFabric system. Also, you cannot add or install multiple packages on a QFabric system.

After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

If you are upgrading more than one package at the same time, delete the operating system package, jkernel, last. Add the operating system package, jkernel, first and the routing software package, jroute, last. If you are upgrading all packages at once, delete and add them in the following order:

```
user@host> request system software add /var/tmp/jbase
user@host> request system software add /var/tmp/jkernel
user@host> request system software add /var/tmp/jpfe
user@host> request system software add /var/tmp/jdocs
user@host> request system software add /var/tmp/jroute
user@host> request system software add /var/tmp/jcrypto
```

By default, when you issue the **request system software add *package-name*** command on a TX Matrix master Routing Engine, all the T640 master Routing Engines that are connected to it are upgraded to the same version of software. If you issue the same command on the TX Matrix backup Routing Engine, all the T640 backup Routing Engines that are connected to it are upgraded to the same version of software.

Likewise, when you issue the **request system software add *package-name*** command on a TX Matrix Plus master Routing Engine, all the T1600 or T4000 master Routing Engines that are connected to it are upgraded to the same version of software. If you issue the same command on the TX Matrix Plus backup Routing Engine, all the T1600 or T4000 backup Routing Engines that are connected to it are upgraded to the same version of software.

Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> <li>• <a href="#">request system software delete on page 752</a></li> <li>• <a href="#">request system software rollback on page 756</a></li> <li>• <a href="#">request system storage cleanup on page 125</a></li> <li>• <a href="#">Upgrading Software</a></li> <li>• <a href="#">Upgrading Software on a QFabric System</a></li> <li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li> </ul>
List of Sample Output	<a href="#">request system software add validate on page 750</a> <a href="#">request system software add (Mixed EX4200 and EX4500 Virtual Chassis) on page 750</a> <a href="#">request system software add component all (QFabric Systems) on page 751</a>
Output Fields	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system software add validate

```
user@host> request system software add validate /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Checking compatibility with configuration
Initializing...
Using jbase-7.1R2.2
Using /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Using /var/validate/tmp/jinstall-signed/jinstall-7.2R1.7-domestic.tgz
Using /var/validate/tmp/jinstall/jbundle-7.2R1.7-domestic.tgz
Checking jbundle requirements on /
Using /var/validate/tmp/jbundle/jbase-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jkernel-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jcrypto-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jpfe-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jdocs-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jroute-7.2R1.7.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Validating against /config/rescue.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-7.2R1.7-domestic-signed.tgz' ...
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Adding jinstall...

WARNING: This package will load JUNOS 7.2R1.7 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-7.2R1.7-domestic-signed.tgz ...
Saving state for rollback ...
```

## Sample Output

### request system software add (Mixed EX4200 and EX4500 Virtual Chassis)

```
user@switch> request system software add set
[/var/tmp/jinstall-ex-4200-11.1R1.1-domestic-signed.tgz
/var/tmp/jinstall-ex-4500-11.1R1.1-domestic-signed.tgz]
...
```

**request system software add component all (QFabric Systems)**

```
user@switch> request system software add /pbdata/packages/jinstall-qfabric-12.2X50-D1.3.rpm  
component all  
...
```

## request system software delete

---

<b>List of Syntax</b>	<a href="#">Syntax on page 752</a> <a href="#">Syntax (TX Matrix Router) on page 752</a> <a href="#">Syntax (TX Matrix Plus Router) on page 752</a>
<b>Syntax</b>	<code>request system software delete <i>software-package</i></code> <code>&lt;force&gt;</code> <code>&lt;reboot&gt;</code> <code>&lt;set [<i>package-name package-name</i>]&gt;</code>
<b>Syntax (TX Matrix Router)</b>	<code>request system software delete <i>software-package</i></code> <code>&lt;force&gt;</code> <code>&lt;lcc <i>number</i>   scc&gt;</code> <code>&lt;reboot&gt;</code> <code>&lt;set [<i>package-name package-name</i>]&gt;</code>
<b>Syntax (TX Matrix Plus Router)</b>	<code>request system software delete <i>software-package</i></code> <code>&lt;force&gt;</code> <code>&lt;lcc <i>number</i>   sfc <i>number</i>&gt;</code> <code>&lt;reboot&gt;</code> <code>&lt;set [<i>package-name package-name</i>]&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option <b>sfc</b> introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Option <b>set [<i>package-name package-name</i>]</b> added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Services Gateways. Option <b>reboot</b> introduced in Junos OS Release 12.3.
<b>Description</b>	Remove a software package or bundle from the router or switch.



**CAUTION:** Before removing a software package or bundle, make sure that you have already placed the new software package or bundle that you intend to load onto the router or switch.

---

- |                |   |
|----------------|---|
| <b>Options</b> | <p><b><i>software-package</i></b>—Software package or bundle name. You can delete any or all of the following software bundles or packages:</p> <ul style="list-style-type: none"><li>• <b>jbase</b>—(Optional) Junos base software suite</li><li>• <b>jcrypto</b>—(Optional, in domestic version only) Junos security software</li><li>• <b>jdocs</b>—(Optional) Junos online documentation file</li><li>• <b>jkernel</b>—(Optional) Junos kernel software suite</li><li>• <b>jpfe</b>—(Optional) Junos Packet Forwarding Engine support</li></ul> |
|----------------|---|



- **jroute**—(Optional) Junos routing software suite
- **junos**—(Optional) Junos base software



**NOTE:** On EX Series switches, some of the package names are different than those listed. To see the list of packages that you can delete on an EX Series switch, enter the command **show system software**.

**force**—(Optional) Ignore warnings and force removal of the software.

**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, remove an extension or upgrade package from a specific T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, remove an extension or upgrade package from a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**reboot**—As of Junos OS 12.3 and greater, automatically reboot upon completing the **request system software delete** command.

**scc**—(TX Matrix routers only) (Optional) Remove an extension or upgrade package from the TX Matrix router (or switch-card chassis).

**set [package-name package-name]**—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages or software add-on packages at the same time.

**sfc number**—(TX Matrix Plus routers only) (Optional) Remove an extension or upgrade package from the TX Matrix Plus router. Replace *number* with 0.

#### Additional Information

Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the `/altroot` and `/altconfig` file systems (on routers) or the `/`, `/altroot`, `/config`, `/var`, and `/var/tmp` file systems (on switches). After you have upgraded the software on the router or switch and are satisfied that the new packages are successfully installed and running, issue the **request system snapshot** command again to back up the new software to the `/altroot` and `/altconfig` file systems (on routers) or the `/`, `/altroot`,

/config, /var, and /var/tmp file systems (on switches). After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

**Required Privilege Level** maintenance

**Related Documentation**

- [request system software add on page 743](#)
- [request system software rollback on page 756](#)
- [request system software validate on page 760](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [request system software delete jdocs on page 754](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### [request system software delete jdocs](#)

The following example displays the system software packages before and after the **jdocs** package is deleted through the **request system software delete** command:

```
user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]

Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]

Information for jdocs:

Comment:
JUNOS Online Documentation [7.2R1.7]

Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

...

user@host> request system software delete jdocs
Removing package 'jdocs' ...

user@host> show system software
```

Information for jbase:

Comment:

JUNOS Base OS Software Suite [7.2R1.7]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [7.2R1.7]

Information for jkernel:

Comment:

JUNOS Kernel Software Suite [7.2R1.7]

...

## request system software rollback

---

<b>List of Syntax</b>	<a href="#">Syntax on page 756</a> <a href="#">Syntax (EX Series Switches) on page 756</a> <a href="#">Syntax (TX Matrix Router) on page 756</a> <a href="#">Syntax (TX Matrix Plus Router) on page 756</a> <a href="#">Syntax (MX Series Router) on page 756</a>
<b>Syntax</b>	request system software rollback
<b>Syntax (EX Series Switches)</b>	request system software rollback <all-members> <local> <member <i>member-id</i> > <reboot>
<b>Syntax (TX Matrix Router)</b>	request system software rollback <lcc <i>number</i>   scc> <reboot>
<b>Syntax (TX Matrix Plus Router)</b>	request system software rollback <lcc <i>number</i>   sfc <i>number</i> > <reboot>
<b>Syntax (MX Series Router)</b>	request system software rollback <all-members> <local> <member <i>member-id</i> > <reboot>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option <b>sfc</b> introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command behavior changed in Junos OS Release 12.1. Option <b>reboot</b> introduced in Junos OS Release 12.3.
<b>Description</b>	<p>For all versions of Junos OS up to and including Junos OS 11.4, revert to the software that was loaded at the last successful <b>request system software add</b> command.</p> <p>As of Junos OS 12.1 and greater, revert to the last known good state before the most recent <b>request system software (add   delete)</b> command. For example, using rollback in Junos OS 12.1 after using <b>request system software add</b> restores the system to a known good state prior to using the <b>add</b> command. Similarly, using rollback in Junos OS 12.1 after using <b>request system software delete</b> restores the system to a known good state prior to using the <b>delete</b> command.</p> <p>A software rollback fails if any required package (or a <b>bundle</b> package containing the required package) cannot be found in /var/sw/pkg.</p> <p><i>Additional Information</i></p>

- On M Series and T Series routers, if **request system software add <jinstall> reboot** was used for the previous installation, then **request system software rollback** has no effect. In this case, use **jinstall** to reinstall the required package.
- On M Series and T Series routers, if **request system software add <sdk1>** was used for the previous installation, then **request system software rollback** removes the last installed SDK package (**sdk1** in this example).
- On SRX Series devices with dual root systems, when **request system software rollback** is run, the system switches to the alternate root. Each root can have a different version of Junos OS. Rollback takes each root back to the previously installed image.
- On QFX3500 and QFX3600 devices in a mixed Virtual Chassis, when the **request system software rollback** command is issued, the system does not rollback to the image stored in the alternate partition.
- On QFX5100 switches, the **reboot** option has been removed. To reboot the switch after a software rollback, issue the **request system reboot** command as a separate, secondary command.

**Options** **all-members**—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on all members of the Virtual Chassis configuration.

**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, attempt to roll back to the previous set of packages on a T640 router connected to the TX Matrix router. On a TX Matrix Plus router, attempt to roll back to the previous set of packages on a connected router connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on the local Virtual Chassis member.

**member member-id**—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**none**—For all versions of Junos OS up to and including Junos OS 11.4, revert to the set of software as of the last successful **request system software add**. As of Junos OS 12.1 and greater, revert to the last known good state before the most recent **request system software (add | delete)** command.

**reboot**—As of Junos OS 12.3 and greater, automatically reboot upon completing the **request system software rollback** command.

**scc**—(TX Matrix routers only) (Optional) Attempt to roll back to the previous set of packages on the TX Matrix router (or switch-card chassis).

**sfc number**—(TX Matrix Plus routers only) (Optional) Attempt to roll back to the previous set of packages on the TX Matrix Plus router. Replace *number* with 0.

**Required Privilege Level** maintenance

**Related Documentation**

- [request system software abort](#)
- [request system software add on page 743](#)
- [request system software delete on page 752](#)
- [request system software validate on page 760](#)
- [request system configuration rescue delete on page 665](#)
- [request system configuration rescue save on page 666](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [request system software rollback on page 759](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system software rollback

```

user@host> request system software rollback
Verified SHA1 checksum of ./jbase-7.2R1.7.tgz
Verified SHA1 checksum of ./jdocs-7.2R1.7.tgz
Verified SHA1 checksum of ./jroute-7.2R1.7.tgz
Installing package './jbase-7.2R1.7.tgz' ...
Available space: 35495 require: 7335
Installing package './jdocs-7.2R1.7.tgz' ...
Available space: 35339 require: 3497
Installing package './jroute-7.2R1.7.tgz' ...
Available space: 35238 require: 6976
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...
Restarting aprobed ...
Restarting apsd ...
Restarting cosd ...
Restarting fsad ...
Restarting fud ...
Restarting gcdrd ...
Restarting ilmid ...
Restarting irsd ...
Restarting l2tpd ...
Restarting mib2d ...
Restarting nasd ...
Restarting pppoed ...
Restarting rdd ...
Restarting rmopd ...
Restarting rtspd ...
Restarting sampled ...
Restarting serviced ...
Restarting snmpd ...
Restarting spd ...
Restarting vrrpd ...

WARNING: cli has been replaced by an updated version:
CLI release 7.2R1.7 built by builder on 2005-04-22 02:03:44 UTC
Restart cli using the new version ? [yes,no] (yes) yes

Restarting cli ...
user@host

```

## request system software validate

---

<b>List of Syntax</b>	<a href="#">Syntax on page 760</a> <a href="#">Syntax (TX Matrix Router) on page 760</a> <a href="#">Syntax (TX Matrix Plus Router) on page 760</a> <a href="#">Syntax (MX Series Router) on page 760</a>
<b>Syntax</b>	<code>request system software validate <i>package-name</i></code> <code>&lt;set [<i>package-name package-name</i>]&gt;</code> <code>&lt;upgrade-with-config&gt;</code> <code>&lt;upgrade-with-config-format <i>format</i>&gt;</code>
<b>Syntax (TX Matrix Router)</b>	<code>request system software validate <i>package-name</i></code> <code>&lt;lcc <i>number</i>   scc&gt;</code> <code>&lt;set [<i>package-name package-name</i>]&gt;</code> <code>&lt;upgrade-with-config&gt;</code> <code>&lt;upgrade-with-config-format <i>format</i>&gt;</code>
<b>Syntax (TX Matrix Plus Router)</b>	<code>request system software validate <i>package-name</i></code> <code>&lt;lcc <i>number</i>   sfc <i>number</i>&gt;</code> <code>&lt;set [<i>package-name package-name</i>]&gt;</code> <code>&lt;upgrade-with-config&gt;</code> <code>&lt;upgrade-with-config-format <i>format</i>&gt;</code>
<b>Syntax (MX Series Router)</b>	<code>request system software validate <i>package-name</i></code> <code>&lt;member <i>member-id</i>&gt;</code> <code>&lt;set [<i>package-name package-name</i>]&gt;</code> <code>&lt;upgrade-with-config&gt;</code> <code>&lt;upgrade-with-config-format <i>format</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. <b>set [<i>package-name package-name</i>]</b> option added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways. <b>upgrade-with-config</b> and <b>upgrade-with-config-format <i>format</i></b> options added in Junos OS Release 12.3 for M Series routers, MX Series routers, and T Series routers.
<b>Description</b>	Validate candidate software against the current configuration of the router.
<b>Options</b>	<b>lcc <i>number</i></b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, validate the software bundle or package on a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, validate the software bundle or package for a specific router that is connected to the TX Matrix Plus router.  Replace <i>number</i> with the following values depending on the LCC configuration: <ul style="list-style-type: none"><li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li><li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li></ul>



- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**member *member-id***—(MX Series routers only) (Optional) Validate the software bundle or package on the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

***package-name***—Name of the software bundle or package to test.

**scc**—(TX Matrix routers only) (Optional) Validate the software bundle or package for the TX Matrix router (or switch-card chassis).

**set [*package-name package-name*]**—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages or software add-on packages at the same time.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Validate the software bundle or package for the TX Matrix Plus router.

**upgrade-with-config**—(Optional) Install one or more configuration files.

**upgrade-with-config-format *format***—(Optional) Specify the configuration file format, **text** or **xml**. The default format is **text**.



**NOTE:** The **upgrade-with-config** and **upgrade-with-config-format** options are only available locally on the router or switch. In a routing matrix, the configuration is applied only to the local router and is not propagated to other routers.

The options are validated during the validation process and applied to the router or switch during the upgrade process. If the upgrade process is successful, the options are removed from the configuration. If the upgrade process fails, the configuration file is renamed with the **.failed** suffix.

**Additional Information** By default, when you issue the **request system software validate** command on a TX Matrix master Routing Engine, all the T640 master Routing Engines that are connected to it are validated. If you issue the same command on the TX Matrix backup Routing Engine, all the T640 backup Routing Engines that are connected to it are upgraded to the same version of software.

Likewise, if you issue the **request system software validate** command on a TX Matrix Plus master Routing Engine, all the T1600 or T4000 master Routing Engines that are connected to it are validated. If you issue the same command on a TX Matrix Plus backup Routing Engine, all the T1600 or T4000 backup Routing Engines that are connected to it are upgraded to the same version of software.

Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"><li>• <i>request system software abort</i></li><li>• <a href="#">request system software add on page 743</a></li><li>• <a href="#">request system software delete on page 752</a></li><li>• <a href="#">request system software rollback on page 756</a></li><li>• <i>Routing Matrix with a TX Matrix Plus Router Solutions Page</i></li></ul>
List of Sample Output	<a href="#">request system software validate (Successful Case) on page 762</a> <a href="#">request system software validate (Failure Case) on page 762</a>
Output Fields	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system software validate (Successful Case)

```
user@host> request system software validate /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Checking compatibility with configuration
Initializing...
Using /packages/jbase-5.3I20020122_1901_sjg
Using /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jbase-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jkernel-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jcrypto-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jpfe-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jdocs-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jroute-5.3I20020124_0520_sjg.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete

WARNING: cli has been replaced by an updated version:
CLI release 5.3I0 built by sjg on 2002-01-24 05:23:53 UTC
Restart cli using the new version ? [yes,no] (yes)
```

### request system software validate (Failure Case)

```
user@host> request system software validate 6.3/
Pushing bundle to lcc0-re0
error: Failed to transfer package to lcc0-re0

user@host> request system software validate test
Pushing bundle to lcc0-re0
Pushing bundle to lcc2-re0

lcc0-re0:
gzip: stdin: not in gzip format
tar: child returned status 1
ERROR: Not a valid package: /var/tmp/test
```

## show system auto-snapshot

<b>Syntax</b>	<b>show system auto-snapshot</b>
<b>Release Information</b>	Command introduced in Junos OS Release 12.3 for EX Series switches. Command introduced in Junos OS Release 12.1X45-D10 for SRX Series devices.
<b>Description</b>	<p>Display automatic snapshot status information. When the automatic snapshot feature is enabled and the system reboots from the alternate root partition, the switch automatically takes a snapshot of the root file system in the alternate root partition and copies it onto the primary root partition. This automatic snapshot procedure takes place whenever the system reboots from the alternate partition, regardless of whether the reboot from the alternate partition is due to a command or due to a corruption of the primary partition.</p> <p>When the automatic snapshot procedure is in progress, you cannot run the manual snapshot command, <b>request system snapshot</b>.</p>
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Resilient Dual-Root Partitions on Switches on page 695</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show system auto-snapshot on page 764</a>
<b>Output Fields</b>	<a href="#">Table 72 on page 763</a> describes the output fields for the <b>show system auto-snapshot</b> command. Output fields are listed in the approximate order in which they appear.

**Table 72: show system auto-snapshot status Output Fields**

Field Name	Field Description
<b>Auto-snapshot configuration</b>	<p>Status of the configuration:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—If the system reboots from the alternate partition, the automatic snapshot feature automatically takes a snapshot of the alternate partition and copies it onto the primary partition.</li> <li>• <b>Disabled</b>—The system does not automatically take a snapshot of the alternate partition. You must use the manual snapshot command, <b>request system snapshot</b>, to take a snapshot of one partition and copy it onto the other.</li> </ul>
<b>Auto-snapshot state</b>	<p>Status of the automatic snapshot procedure:</p> <ul style="list-style-type: none"> <li>• <b>Completed</b>—The automatic snapshot procedure has completed copying the alternate partition to the primary partition and the alarm has been cleared.</li> <li>• <b>Disabled</b>—The automatic snapshot procedure is inactive.</li> <li>• <b>In progress</b>—The automatic snapshot procedure is in progress. It takes about 10 to 15 minutes to complete, depending upon disk size.</li> </ul>

## Sample Output

### show system auto-snapshot

```
user@switch> show system auto-snapshot
Auto-snapshot Configuration: Enabled
Auto-snapshot State: Disabled
```

## show system boot-messages

<b>List of Syntax</b>	<a href="#">Syntax on page 765</a> <a href="#">Syntax (EX Series Switches) on page 765</a> <a href="#">Syntax (TX Matrix Router) on page 765</a> <a href="#">Syntax (TX Matrix Plus Router) on page 765</a> <a href="#">Syntax (MX Series Router) on page 765</a> <a href="#">Syntax (QFX Series) on page 765</a>
<b>Syntax</b>	show system boot-messages
<b>Syntax (EX Series Switches)</b>	show system boot-messages <all-members> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	show system boot-messages <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system boot-messages <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system boot-messages <all-members> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show system boot-messages infrastructure <i>name</i>   interconnect-device <i>name</i>   node-group <i>name</i>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display initial messages generated by the system kernel upon startup. These messages are the contents of <code>/var/run/dmesg.boot</code> .
<b>Options</b>	<b>none</b> —Display all boot time messages.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display boot time messages for all of the chassis.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for all T640 routers connected to a TX Matrix router. On a TX Matrix Plus router, display boot time messages for all connected T1600 or T4000 LCCs.  <b>all-members</b> —(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on all members of the Virtual Chassis configuration.

**infrastructure *name***—(QFabric systems only) (Optional) Display boot time messages on the fabric control Routing Engine or fabric manager Routing engines.

**interconnect-device *name***—(QFabric systems only) (Optional) Display boot time messages on the Interconnect device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for a specific T640 router connected to a TX Matrix router. On a TX Matrix Plus router, display boot time messages for a specific router connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display boot time messages on the Node group.

**scc**—(TX Matrix routers only) (Optional) Display boot time messages for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display boot time messages for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system boot-messages** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

## Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

## List of Sample Output

- [show system boot-messages \(TX Matrix Router\) on page 767](#)
- [show system boot-messages lcc \(TX Matrix Router\) on page 768](#)
- [show system boot-messages \(TX Matrix Plus Router\) on page 769](#)
- [show system boot-messages \(QFX3500 Switch\) on page 769](#)

## Sample Output

### show system boot-messages (TX Matrix Router)

```

user@host> show system boot-messages
Copyright (c) 1992-1998 FreeBSD Inc.
Copyright (c) 1996-2000 Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1982, 1986, 1989, 1991, 1993
    The Regents of the University of California. All rights reserved.

JUNOS 4.1-20000216-Zf8469 #0: 2000-02-16 12:57:28 UTC
    tlim@single.juniper.net:/p/build/20000216-0905/4.1/release_kernel/sys/compil
e/GENERIC
CPU: Pentium Pro (332.55-MHz 686-class CPU)
    Origin = "GenuineIntel" Id = 0x66a Stepping=10
    Features=0x183f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,<b
16>,<b17>,MMX,<b24>>
Teknor CPU Card Recognized
real memory = 805306368 (786432K bytes)
avail memory = 786280448 (767852K bytes)
Probing for devices on PCI bus 0:
chip0 <generic PCI bridge (vendor=8086 device=7192 subclass=0)> rev 3 class 6000
0 on pci0:0:0
chip1 <Intel 82371AB PCI-ISA bridge> rev 1 class 60100 on pci0:7:0
chip2 <Intel 82371AB IDE interface> rev 1 class 10180 on pci0:7:1
chip3 <Intel 82371AB USB interface> rev 1 class c0300 int d irq 11 on pci0:7:2
smb0 <Intel 82371AB SMB controller> rev 1 class 68000 on pci0:7:3
pcic0 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int a irq 15 on pci0:13
:0
TI1131 PCI Config Reg: [pci only][FUNC0 pci int]
pcic1 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int b irq 12 on pci0:13
:1
TI1131 PCI Config Reg: [pci only][FUNC1 pci int]
fxp0 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 12 on

pci0:16:0
chip4 <generic PCI bridge (vendor=1011 device=0022 subclass=4)> rev 4 class 6040
0 on pci0:17:0
fxp1 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on

pci0:19:0
Probing for devices on PCI bus 1:
mcs0 <Miscellaneous Control Subsystem> rev 12 class ff0000 int a irq 12 on pci1:
13:0
fxp2 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on

pci1:14:0
Probing for devices on the ISA bus:
sc0 at 0x60-0x6f irq 1 on motherboard
sc0: EGA color <16 virtual consoles, flags=0x0>
ed0 not found at 0x300

```

```

ed1 not found at 0x280
ed2 not found at 0x340
psm0 not found at 0x60
sio0 at 0x3f8-0x3ff irq 4 flags 0x20010 on isa
sio0: type 16550A, console
sio1 at 0x3e8-0x3ef irq 5 flags 0x20000 on isa
sio1: type 16550A
sio2 at 0x2f8-0x2ff irq 3 flags 0x20000 on isa
sio2: type 16550A
pcic0 at 0x3e0-0x3e1 on isa
PC-Card ctlr(0) TI PCI-1131 [CardBus bridge mode] (5 mem & 2 I/O windows)
pcic0: slot 0 controller I/O address 0x3e0
npx0 flags 0x1 on motherboard
npx0: INT 16 interface
fdc0: direction bit not set
fdc0: cmd 3 failed at out byte 1 of 3
fdc0 not found at 0x3f0
wdc0 at 0x1f0-0x1f7 irq 14 on isa
wdc0: unit 0 (wd0): <SunDisk SQFXB-80>, single-sector-i/o
wd0: 76MB (156672 sectors), 612 cyls, 8 heads, 32 S/T, 512 B/S
wdc0: unit 1 (wd1): <IBM-DCXA-210000>
wd1: 8063MB (16514064 sectors), 16383 cyls, 16 heads, 63 S/T, 512 B/S
wdc1 not found at 0x170
wdc2 not found at 0x180
ep0 not found at 0x300
fxp0: Ethernet address 00:a0:a5:12:05:5a
fxp1: Ethernet address 00:a0:a5:12:05:59
fxp2: Ethernet address 02:00:00:00:00:01
swapon: adding /dev/wd1s1b as swap device
Automatic reboot in progress...
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd0s1e: clean, 9233 free (9 frags, 1153 blocks, 0.1% fragmentation)
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd1s1f: clean, 4301055 free (335 frags, 537590 blocks, 0.0% fragmentation)

```

### show system boot-messages lcc (TX Matrix Router)

```

user@host> show system boot-messages lcc 2
lcc2-re0:
-----
Copyright (c) 1996-2001, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2001 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 7.0-20040912.0 #0: 2004-09-12 09:16:32 UTC

builder@benten.juniper.net:/build/benten-b/7.0/20040912.0/obj-i386/sys/compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz
Timecounter "TSC" frequency 601368936 Hz
CPU: Pentium III/Pentium III Xeon/Celeron (601.37-MHz 686-class CPU)
    Origin = "GenuineIntel" Id = 0x68a Stepping = 10

Features=0x387f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,PAT,PSE36,PN,MMX,FXSR,SSE>
real memory = 2147467264 (2097136K bytes)
sio0: gdb debugging port
avail memory = 2084040704 (2035196K bytes)
Preloaded elf kernel "kernel" at 0xc06d9000.
DEVFS: ready for devices
Pentium Pro MTRR support enabled
md0: Malloc disk

```



```

DRAM Data Integrity Mode: ECC Mode with h/w scrubbing
npx0: <math processor> on motherboard
npx0: INT 16 interface
pcib0: <ServerWorks NB6635 3.0LE host to PCI bridge> on motherboard
pci0: <PCI bus> on pcib0
pcic-pci0: <TI PCI-1410 PCI-CardBus Bridge> irq 15 at device 1.0 on pci0
pcic-pci0: TI12XX PCI Config Reg: [pwr save][pci only]
fxp0: <Intel Embedded 10/100 Ethernet> port 0x1000-0x103f mem
0xfb800000-0xfb81ffff,0xfb820000-0xfb820fff irq 9 at device 3.0 on pci0
fxp1: <Intel Embedded 10/100 Ethernet> port 0x1040-0x107f mem
0xfb840000-0xfb85ffff,0xfb821000-0xfb821fff irq 11 at device 4.0 on pci0
...

```

### show system boot-messages (TX Matrix Plus Router)

```

user@host> show system boot-messages
sfc0-re0:
-----
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 9.6B3.3 #0: 2009-06-17 19:52:08 UTC

builder@lanath.juniper.net:/volume/build/junos/9.6/release/9.6B3.3/obj-i386/bsd/sys/compile/JUNIPER
MPTable: Timecounter "i8254" frequency 1193182 Hz quality 0 CPU: Intel(R) Xeon(R)
CPU          L5238 @ 2.66GHz (2660.01-MHz 686-class CPU)   Origin =
"GenuineIntel" Id = 0x1067a Stepping = 10   Features=0xbfebfbff
...
lcc1-re0:
-----
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 9.6-20090617.0 #0: 2009-06-17 04:15:14 UTC

builder@lanath.juniper.net:/volume/build/junos/9.6/production/20090617.0/obj-i386/bsd/sys/compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz quality 0
CPU: Intel(R) Xeon(R) CPU                               @ 1.86GHz (1862.01-MHz 686-class CPU)

Origin = "GenuineIntel" Id = 0x1067a Stepping = 10
Features=0xbfebfbff
...

```

### show system boot-messages (QFX3500 Switch)

```

user@switch> show sytem boot-messages
getmemsize: msgbufp[size=32768] = 0x81d07fe4

System physical memory distribution:
-----
Total physical memory: 4160749568 (3968 MB)
Physical memory used: 3472883712 (3312 MB)
Physical memory allocated to kernel: 2130706432 (2032 MB)
Physical memory allocated to user BTLB: 1342177280 (1280 MB)
-----

Copyright (c) 1996-2010, Juniper Networks, Inc.

```

All rights reserved.

Copyright (c) 1992-2006 The FreeBSD Project.

Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994

The Regents of the University of California. All rights reserved.

JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC

```
ssiano@svl-junos-pool125.juniper.net:/c/ssiano/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
WARNING: debug.mpsafenet forced to 0 as ipsec requires Giant
JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC
```

```
ssiano@svl-junos-pool125.juniper.net:/c/ssiano/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
real memory = 3472883712 (3312MB)
avail memory = 1708171264 (1629MB)
cpuid: 0, bt1b_cpumap:0xffffffff8
FreeBSD/SMP: Multiprocessor System Detected: 12 CPUs
ETHERNET SOCKET BRIDGE initialising
Initializing QFX platform properties ..
cpu0 on motherboard
: RMI's XLR CPU Rev. 0.3 with no FPU implemented
  L1 Cache: I size 32kb(32 line), D size 32kb(32 line), eight way.
  L2 Cache: Size 1024kb, eight way
pic_lbus0: <XLR Local Bus>
pic_lbus0: <XLR Local Bus> on motherboard
Enter qfx control ethernet probe addr:0xc5eeec00
gmac4: <XLR GMAC GE Ethernet> on pic_lbus0
me0: Ethernet address 00:1d:b5:f7:68:40
Enter qfx control ethernet probe addr:0xc5eeeb40
gmac5: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:41
Enter qfx control ethernet probe addr:0xc5eeea80
gmac6: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:42
sio0 on pic_lbus0
Entering sioattach
sio0: type 16550A, console
xls_setup_intr: skip irq 3, xlr regs are set up somewhere else.
gblmem0 on pic_lbus0
ehci0: <RMI XLS USB 2.0 controller> on pic_lbus0
ehci_bus_attach: allocated resource. tag=1, base=bef24000
xls_ehci_init: endian hardware swapping NOT enabled.
usb0: EHCI version 1.0
usb0 on ehci0
usb0: USB revision 2.0
uhub0: vendor 0x0000 EHCI root hub, class 9/0, rev 2.00/1.00, addr 1
uhub0: 2 ports with 2 removable, self powered
umass0: USB USBFlashDrive, rev 2.00/11.00, addr 2
pcib0: PCIe link 0 up
pcib0: PCIe link 2 up
pcib0: PCIe link 3 up
pcib0: <XLS PCI Host Controller> on pic_lbus0
pci0: <PCI bus> on pcib0
pcib1: <PCI-PCI bridge> at device 0.0 on pci0
pci1: <PCI bus> on pcib1
pci1: <network, ethernet> at device 0.0 (no driver attached)
pcib2: <PCI-PCI bridge> at device 1.0 on pci0
pcib3: <PCI-PCI bridge> at device 2.0 on pci0
pci2: <PCI bus> on pcib3
pci2: <network, ethernet> at device 0.0 (no driver attached)
pcib4: <PCI-PCI bridge> at device 3.0 on pci0
```

```

pci3: <PCI bus> on pcib4
pci3: <network, ethernet> at device 0.0 (no driver attached)
cfi device address space at 0xbc000000
cfi0: <AMD/Fujitsu - 8MB> on pic_lbus0
cfi device address space at 0xbc000000
i2c0: <I2C bus controller> on pic_lbus0
i2c1: <I2C bus controller> on pic_lbus0
qfx_fmn0 on pic_lbus0
pool offset 1503776768
xlr_lbus0: <XLR Local Bus Controller> on motherboard
qfx_bcpld_probe[124]
qfx_bcpld_probe[138]: dev_type=0x0
qfx_bcpld_probe[124]
qfx_bcpld0: QFX BCPLD probe success
qfx_bcpld0qfx_bcpld_attach[174]
qfx_bcpld_attach[207] : bus_space_tag=0x0, bus_space_handle=0xbd900000
qfx_bcpld_probe[124]
qfx_bcpld1: QFX BCPLD probe success
qfx_bcpld1qfx_bcpld_attach[174]
tor_bcpld_slave_attach[1245] : bus_space_tag=0x0, bus_space_handle=0xbda00000
Initializing product: 96 ..
bmeb: bmeb_lib_init done 0xc60a5000, addr 0x809c99a0
bme0:Virtual BME driver initializing
Timecounter "mips" frequency 1200000000 Hz quality 0
Timecounter "xlr_pic_timer" frequency 66666666 Hz quality 1
Timecounters tick every 1.000 msec
Loading the NETPFE fc module
IPsec: Initialized Security Association Processing.
SMP: AP CPU #3 Launched!
SMP: AP CPU #1 Launched!
SMP: AP CPU #2 Launched!
SMP: AP CPU #4 Launched!
SMP: AP CPU #5 Launched!
SMP: AP CPU #7 Launched!
SMP: AP CPU #6 Launched!
SMP: AP CPU #11 Launched!
SMP: AP CPU #10 Launched!
SMP: AP CPU #9 Launched!
SMP: AP CPU #8 Launched!
da0 at umass-sim0 bus 0 target 0 lun 0
da0: <USB USBFlashDrive 1100> Removable Direct Access SCSI-0 device
da0: 40.000MB/s transfers
da0: 3920MB (8028160 512 byte sectors: 255H 63S/T 499C)
Trying to mount root from ufs:/dev/da0s1a

```

## show system license

<b>Syntax</b>	show system license <installed   keys   usage>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 13.3 for the MX104 3D Universal Edge Routers.
<b>Description</b>	Display licenses and information about how they are used.
<b>Options</b>	<p><b>none</b>—Display all license information.</p> <p><b>installed</b>—(Optional) Display installed licenses only.</p> <p><b>keys</b>—(Optional) Display a list of license keys. Use this information to verify that each expected license key is present.</p> <p><b>usage</b>—(Optional) Display the state of licensed features.</p>
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">show system license on page 773</a> <a href="#">show system license installed on page 774</a> <a href="#">show system license keys on page 774</a> <a href="#">show system license usage on page 774</a> <a href="#">show system license (MX104 Routers) on page 774</a> <a href="#">show system license installed (MX104 Routers) on page 775</a> <a href="#">show system license keys (MX104 Routers) on page 775</a> <a href="#">show system license usage (MX104 Routers) on page 775</a> <a href="#">show system license (MX104 Routers) on page 775</a> <a href="#">show system license installed (MX104 Routers) on page 776</a> <a href="#">show system license keys (MX104 Routers) on page 776</a> <a href="#">show system license usage (MX104 Routers) on page 776</a> <a href="#">show system license (MX104 Routers) on page 777</a> <a href="#">show system license installed (MX104 Routers) on page 777</a> <a href="#">show system license keys (MX104 Routers) on page 777</a> <a href="#">show system license usage (MX104 Routers) on page 778</a> <a href="#">show system license (QFX Series) on page 778</a>
<b>Output Fields</b>	Table 73 on page 772 lists the output fields for the <b>show system license</b> command. Output fields are listed in the approximate order in which they appear.

**Table 73: show system license Output Fields**

Field Name	Field Description
<b>Feature name</b>	Name assigned to the configured feature. You use this information to verify that all the features for which you installed licenses are present.

Table 73: show system license Output Fields (*continued*)

Field Name	Field Description
<b>Licenses used</b>	<p>Number of licenses used by a router or switch. You use this information to verify that the number of licenses used matches the number configured. If a licensed feature is configured, the feature is considered used.</p> <p><b>NOTE:</b> In Junos OS Release 10.1 and later, the <b>Licenses used</b> column displays the actual usage count based on the number of active sessions or connections as reported by the corresponding feature daemons. This is applicable for scalable license-based features such as Subscriber Access (<b>scale-subscriber</b>), L2TP (<b>scale-l2tp</b>), Mobile IP (<b>scale-mobile-ip</b>), and so on.</p>
<b>Licenses installed</b>	<p>Information about the installed license key:</p> <ul style="list-style-type: none"> <li>• <b>License identifier</b>—Identifier associated with a license key.</li> <li>• <b>State</b>—State of the license key: <b>valid</b> or <b>invalid</b>. An <b>invalid</b> state indicates that the key was entered incorrectly or is not valid for the specific device.</li> <li>• <b>License version</b>—Version of a license. The version indicates how the license is validated, the type of signature, and the signer of the license key.</li> <li>• <b>Valid for device</b>—Device that can use a license key.</li> <li>• <b>Group defined</b>—Group membership of a device.</li> <li>• <b>Features</b>—Feature associated with a license, such as data link switching (DLSw).</li> </ul>
<b>Licenses needed</b>	Number of licenses required for features being used but not yet properly licensed.
<b>Expiry</b>	Amount of time left within the grace period before a license is required for a feature being used.

## Sample Output

### show system license

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

```
License usage:
```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-accounting	2	2	0	permanent
subscriber-authentication	1	2	0	permanent
subscriber-address-assignment	2	2	0	permanent
subscriber-vlan	2	2	0	permanent
subscriber-ip	0	2	0	permanent
scale-subscriber	2	3	0	permanent
scale-l2tp	4	5	0	permanent
scale-mobile-ip	1	2	0	permanent

```
Licenses installed:
```

```
License identifier: XXXXXXXXXX
```

```
License version: 2
```

```
Features:
```

```
subscriber-accounting - Per Subscriber Radius Accounting
permanent
subscriber-authentication - Per Subscriber Radius Authentication
permanent
subscriber-address-assignment - Radius/SRC Address Pool Assignment
permanent
subscriber-vlan - Dynamic Auto-sensed Vlan
```

```

    permanent
subscriber-ip    - Dynamic and Static IP
    permanent

```

### show system license installed

```

user@host> show system license installed
License identifier: XXXXXXXXXX
License version: 2
Features:
  subscriber-accounting - Per Subscriber Radius Accounting
    permanent
  subscriber-authentication - Per Subscriber Radius Authentication
    permanent
  subscriber-address-assignment - Radius/SRC Address Pool Assignment
    permanent
  subscriber-vlan - Dynamic Auto-sensed Vlan
    permanent
  subscriber-ip - Dynamic and Static IP
    permanent

```

### show system license keys

```

user@host> show system license keys
XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxx

```

### show system license usage

```

user@host> show system license usage
License usage:

```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-accounting	2	2	0	permanent
subscriber-authentication	1	2	0	permanent
subscriber-address-assignment	2	2	0	permanent
subscriber-vlan	2	2	0	permanent
subscriber-ip	0	2	0	permanent
scale-subscriber	2	3	0	permanent
scale-l2tp	4	5	0	permanent
scale-mobile-ip	1	2	0	permanent

### show system license (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host> show system license
License usage:

```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent

```

Licenses installed:
License identifier: XXXXXXXXXX
License version: 2
Features:

```

```

MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
    permanent

```

### show system license installed (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host > show system license installed
License identifier: XXXXXXXXXX
License version: 2
Features:
MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
    permanent

```

### show system license keys (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host > show system license keys

XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxx

```

### show system license usage (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host > show system license usage

```

Feature name	Licenses used	Licenses installed	Expiry needed	
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent

### show system license (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license
License usage:

```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

```

Licenses installed:
License identifier: XXXXXXXXXX
License version: 2

```

```

Features:
  MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
  permanent

License identifier: XXXXXXXXXX
License version: 2
Features:
  MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)
upgrade
  permanent

```

### show system license installed (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license installed
License identifier: XXXXXXXXXX
License version: 2
Features:
  MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
  permanent

License identifier: XXXXXXXXXX
License version: 2
Features:
  MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)
upgrade
  permanent

```

### show system license keys (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license keys

XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxx

XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxx

```

### show system license usage (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license usage

```

Feature name	Licenses used	Licenses installed	Expiry needed	
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent



scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

### show system license (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

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

License usage:

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

Licenses installed:

License identifier: XXXXXXXXXX

License version: 2

Features:

```

  MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
  permanent
  MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)
upgrade
  permanent
```

### show system license installed (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

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

License identifier: XXXXXXXXXX

License version: 2

Features:

```

  MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
  permanent
  MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)
upgrade
  permanent
```

### show system license keys (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

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

```

XXXXXXXX XXXXXX XXXXXX XXXXXX XXXXXX XXXXXX
XXXXXXXX XXXXXX XXXXXX XXXXXX XXXXXX
XXXXXXXX XXXXXX X
```

### show system license usage (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

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

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

### show system license (QFX Series)


```
user@switch> show system license
```

License usage:

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
qfx-edge-fab	1	1	1	permanent

Licenses installed:  
License identifier: JUNOS417988  
License version: 1  
Features:  
qfx-edge-fab - QFX3000 Series QF/Node feature license  
permanent

## show system snapshot

<b>List of Syntax</b>	<a href="#">Syntax on page 779</a> <a href="#">Syntax (EX Series Switches) on page 779</a>
<b>Syntax</b>	show system snapshot
<b>Syntax (EX Series Switches)</b>	show system snapshot <all-members local member <i>member-id</i> > <media (external   internal)>
<b>Release Information</b>	Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 10.0 for EX Series switches.
<b>Description</b>	Display information about the backup software: <ul style="list-style-type: none"> <li>• On the routers, display information about the backup software, which is located in the <code>/altroot</code>, and <code>/altconfig</code> file systems or on the alternate media.</li> <li>• On the switches, display information about the backup of the root file system (<code>/</code>) and directories <code>/altroot</code>, <code>/config</code>, <code>/var</code>, and <code>/var/tmp</code>, which are located either on an external USB flash drive or in internal flash memory.</li> </ul>
<div>  <b>NOTE:</b> To back up software, use the <code>request system snapshot</code> command.         </div>	
<b>Options</b>	<b>none</b> —Display information about the backup software.  <b>all-members   local   member <i>member-id</i></b> —(EX Series switch Virtual Chassis only) (Optional) Display the snapshot in a Virtual Chassis: <ul style="list-style-type: none"> <li>• <b>all-members</b>—Display the snapshot for all members of the Virtual Chassis.</li> <li>• <b>local</b>—Display the snapshot on the member of the Virtual Chassis that you are currently logged into.</li> <li>• <b>member <i>member-id</i></b>—Display the snapshot for the specified member of the Virtual Chassis.</li> </ul> <b>media (external   internal)</b> —(EX Series switch only) (Optional) Display the destination media location for the snapshot. The <b>external</b> option specifies the snapshot on an external mass storage device, such as a USB flash drive. The <b>internal</b> option specifies the snapshot on an internal memory source, such as internal flash memory. If no additional options are specified, the command displays the snapshot stored in both slices.
<b>Required Privilege Level</b>	view

**Related Documentation** • [request system snapshot on page 737](#)

**List of Sample Output** [show system snapshot \(Router\) on page 780](#)  
[show system snapshot media external \(Switch\) on page 780](#)  
[show system snapshot media internal \(Switch\) on page 781](#)

**Output Fields** [Table 74 on page 780](#) lists the output fields for the **show system snapshot** command. Output fields are listed in the approximate order in which they appear.

**Table 74: show system snapshot Output Fields**

Field Name	Field Description
Creation date	Date and time of the last snapshot.
JUNOS version on snapshot	Junos OS release number of individual software packages.

## Sample Output

### show system snapshot (Router)

```
user@host> show system snapshot
Information for snapshot on hard-disk
Creation date: Oct 5 13:53:29 2005
JUNOS version on snapshot:
  jbase   : 7.3R2.5
  jcrypto: 7.3R2.5
  jdocs   : 7.3R2.5
  jkernel: 7.3R2.5
  jpfe    : M40-7.3R2.5
  jroute  : 7.3R2.5
```

### show system snapshot media external (Switch)

```
user@switch> show system snapshot media external
Information for snapshot on      external (/dev/dals1a) (backup)
Creation date: Mar 19 03:37:18 2012
JUNOS version on snapshot:
  jbase   : ex-12.1I20120111_0048_user
  jcrypto-ex: 12.1I20120111_0048_user
  jdocs-ex: 12.1I20120111_0048_user
  jroute-ex: 12.1I20120111_0048_user
  jswitch-ex: 12.1I20120111_0048_user
  jweb-ex: 12.1I20120111_0048_user
Information for snapshot on      external (/dev/dals2a) (primary)
Creation date: Mar 19 03:38:25 2012
JUNOS version on snapshot:
  jbase   : ex-12.2I20120305_2240_user
  jcrypto-ex: 12.2I20120305_2240_user
  jdocs-ex: 12.2I20120305_2240_user
  jroute-ex: 12.2I20120305_2240_user
  jswitch-ex: 12.2I20120305_2240_user
  jweb-ex: 12.2I20120305_2240_user
```

### show system snapshot media internal (Switch)

```
user@switch> show system snapshot media internal
Information for snapshot on internal (/dev/da0s1a) (backup)
Creation date: Mar 14 05:01:02 2011
JUNOS version on snapshot:
  jbase : 11.1R1.9
  jcrypto-ex: 11.1R1.9
  jdocs-ex: 11.1R1.9
  jkernel-ex: 11.1R1.9
  jroute-ex: 11.1R1.9
  jswitch-ex: 11.1R1.9
  jweb-ex: 11.1R1.9
  jpfe-ex42x: 11.1R1.9
Information for snapshot on internal (/dev/da0s2a) (primary)
Creation date: Mar 30 08:46:27 2011
JUNOS version on snapshot:
  jbase : 11.2-20110330.0
  jcrypto-ex: 11.2-20110330.0
  jdocs-ex: 11.2-20110330.0
  jkernel-ex: 11.2-20110330.0
  jroute-ex: 11.2-20110330.0
  jswitch-ex: 11.2-20110330.0
  jweb-ex: 11.2-20110330.0
  jpfe-ex42x: 11.2-20110330.0
```

## show system storage partitions (EX Series Switches Only)

<b>Syntax</b>	show system storage partitions <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for EX Series switches.
<b>Description</b>	Display information about the disk partitions on EX Series switches.
<b>Options</b>	<p><b>none</b>—Display partition information.</p> <p><b>all-members</b>—(Virtual Chassis systems only) (Optional) Display partition information for all members of the Virtual Chassis.</p> <p><b>local</b>—(Virtual Chassis systems only) (Optional) Display partition information for the local Virtual Chassis member.</p> <p><b>member <i>member-id</i></b>—(Virtual Chassis systems only) (Optional) Display partition information for the specified member of the Virtual Chassis configuration.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Verifying Junos OS and Boot Loader Software Versions on an EX Series Switch on page 720</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show system storage partitions on page 783</a>
<b>Output Fields</b>	<a href="#">Table 75 on page 782</a> describes the output fields for the <b>show system storage partitions</b> command. Output fields are listed in the approximate order in which they appear.

**Table 75: show system storage partitions Output Fields**

Field Name	Field Description
Boot Media	Media (internal or external) from which the switch was booted.
Active Partition	Name of the active root partition.
Backup Partition	Name of the backup (alternate) root partition.
Currently booted from	Partition from which the switch was last booted.
Partitions information	Information about partitions on the boot media: <ul style="list-style-type: none"> <li>• Partition—Partition identifier.</li> <li>• Size—Size of partition.</li> <li>• Mountpoint—Directory on which the partition is mounted.</li> </ul>

## Sample Output

### show system storage partitions

```
user@switch> show system storage partitions
fpc0:
-----
Boot Media: internal (da0)
Active Partition: da0s1a
Backup Partition: da0s2a
Currently booted from: active (da0s1a)

Partitions information:
  Partition  Size  Mountpoint
  s1a        184M  /
  s2a        184M  altroot
  s3d        369M  /var/tmp
  s3e        123M  /var
  s4d         62M  /config
  s4e                unused (backup config)
```





# Troubleshooting Procedures

- [Troubleshooting Software Installation on page 785](#)
- [Troubleshooting a Switch That Has Booted from the Backup Junos OS Image on page 788](#)
- [Resilient Dual-Root Partitions Frequently Asked Questions on page 789](#)

## Troubleshooting Software Installation

---

This topic describes troubleshooting issues with software installations on EX Series switches.

- [Recovering from a Failed Software Upgrade on an EX Series Switch on page 785](#)
- [Rebooting from the Inactive Partition on page 786](#)
- [Freeing Disk Space for Software Installation on page 787](#)
- [Installation from the Boot Loader Generates 'cannot open package' Error on page 787](#)

## Recovering from a Failed Software Upgrade on an EX Series Switch

**Problem**    **Description:** If Junos OS loads but the CLI is not working, or if the switch has no software installed, use this recovery installation procedure to install Junos OS.

**Solution**    If there is already a Junos OS image on the system, you can either install the new Junos OS package in a separate partition and have both Junos OS images remain on the system, or you can wipe the disk clean before the new installation proceeds.

If there is no Junos OS image on the system, follow the instructions in "[Booting an EX Series Switch Using a Software Package Stored on a USB Flash Drive](#)" on page 712 to get an image on the system and boot the switch.

To perform a recovery installation:

1. Power on the switch. The loader script starts.

After the message **Loading /boot/defaults/loader.conf** displays, you are prompted with:

**Hit [Enter] to boot immediately, or space bar for command prompt.**

2. Press the space bar to enter the manual loader. The **loader>** prompt displays.

3. Enter the following command:

```
loader> install [--format] [--external] source
```

where:

- **format**—Use this option to wipe the installation media before installing the software package. If you do not include this option, the system installs the new Junos OS package in a different partition from the partition used by the most recently installed Junos OS package.
- **external**—Use this option to install the software package on an external medium.
- **source**—Represents the name and location of the Junos OS package either on a server on the network or as a file on the USB flash drive:
  - Network address of the server and the path on the server; for example, **tftp://192.171.28/junos/jinstall-ex-4200-9.4R1.5-domestic-signed.tgz**
  - The Junos OS package on a USB device is commonly stored in the root drive as the only file; for example, **file:///jinstall-ex-4200-9.4R1.5-domestic-signed.tgz**

The boot process proceeds as normal and ends with a login prompt.

## Rebooting from the Inactive Partition

**Problem**    **Description:** EX Series switches shipped with Junos OS Release 10.4R2 or earlier have Junos OS loaded on the system disk in partition 1. The first time you upgrade, the new software package is installed in partition 2. When you finish the installation and reboot, partition 2 becomes the active partition. Similarly, subsequent software packages are installed in the inactive partition, which becomes the active partition when you reboot at the end of the installation process.

On switches shipped with Release 10.4R3 and later, the same Junos OS image is loaded in each of the two root partitions, and you should copy the new software image to the alternate partition each time you upgrade.

If you performed an upgrade and rebooted, the system resets the active partition. You can use this procedure to manually boot from the inactive partition.



**NOTE:** If you have completed the installation of the software image but have not yet rebooted, issue the **request system software rollback** command to return to the original software installation package.

**Solution**    Reboot from the inactive partition:

```
user@switch> request system reboot slice alternate
```



**NOTE:** If you cannot access the CLI, you can reboot from the inactive partition using the following procedure from the loader script prompt:

1. Unload and clear the interrupted boot from the active partition:

```
loader> unload
loader> unset vfs.root.mountfrom
```

2. Select the new (inactive) partition to boot from:

```
loader> set currdev=diskxsy:
```

where *x* is either 0 (internal) or 1 (external) and the *y* indicates the number of the inactive partition, either 1 or 2.

You must include the colon (:) at the end of this command.

3. Boot Junos OS from the inactive partition:

```
loader> boot
```

## Freeing Disk Space for Software Installation

**Problem**    **Description:** The software installation process requires a certain amount of unused disk space. If there is not enough space, you might receive an error message such as:

```
fetch: /var/tmp/incoming-package.tgz: No space left on device
```

**Solution**    Identify and delete unnecessary files by using the [request system storage cleanup](#) command.

## Installation from the Boot Loader Generates 'cannot open package' Error

**Problem**    **Description:** When installing a Junos OS software image from the loader prompt, a "cannot open package error" is generated:

```
loader> install - -format
tftp://10.204.33.248/images/Flash_corr/official/jinstall-ex-4200-10.4I2011012-domestic-signed.tgz
Speed: 1000, full duplex
bootp: no reply
No response for RARP request
net_open: RARP failed
cannot open package (error 5)
```

**Solution**    This might be due to the IP address, gateway IP address, netmask address, or server IP address not being properly set. You can set these values either from the shell or from the u-boot prompt.

To set these values from the shell:

```
% nvram setenv ipaddr 10.204.35.235
% nvram setenv netmask 255.255.240.0
```

```
% nvram setenv gatewayip 10.204.47.254
```

```
% nvram setenv serverip 10.204.33.248
```

To set these values from the u-boot prompt, log in to a console connection, reboot, and stop at the u-boot prompt (Cntrl+c):

```
=> setenv ipaddr 10.204.35.235
```

```
=> setenv gatewayip 10.204.47.254
```

```
=> setenv serverip 10.204.33.248
```

```
=> setenv netmask 255.255.240.0
```

```
=> saveenv
```

```
=> printenv Verify whether variables are set properly or not
```

```
=> boot
```

#### Related Documentation

- *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*
- *Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)*
- *Installing Software on EX Series Switches (J-Web Procedure)*
- [Understanding Software Installation on EX Series Switches on page 691](#)
- [show system storage partitions \(EX Series Switches Only\) on page 782](#)

---

## Troubleshooting a Switch That Has Booted from the Backup Junos OS Image

**Problem** **Description:** The switch boots from the backup root file partition. It is possible that the primary copy of JUNOS OS failed to boot properly, which could indicate that it is corrupted. This event is flagged in two ways:

- Upon login through the console or management port, the following warning message is displayed:

```
WARNING: THIS DEVICE HAS BOOTED FROM THE BACKUP JUNOS IMAGE
```

It is possible that the primary copy of JUNOS failed to boot up properly, and so this device has booted from the backup copy.

Please re-install JUNOS to recover the primary copy in case it has been corrupted.

- The following alarm message is generated:

```
user@switch> show chassis alarms
1 alarms currently active
Alarm time           Class  Description
2011-02-17 05:48:49 PST  Minor  Host 0 Boot from backup root
```

If the switch is in a Virtual Chassis, the switch member number appears in the **Description** field, where the switch is called a host.

**Solution** Install a new Junos OS image on the partition that had the corruption, or take a snapshot (use [request system snapshot](#)) of the currently active partition and use it to replace the image in the alternate partition:

If the switch is a standalone switch or a Virtual Chassis master switch, enter this command:

```
user@switch> request system snapshot slice alternate
```

If the switch is a Virtual Chassis member switch (not the master), enter this command on the Virtual Chassis:

```
user@switch> request system snapshot slice alternate member member-id
```

where *member-id* is the Virtual Chassis member ID number.

#### Related Documentation

- [Verifying Junos OS and Boot Loader Software Versions on an EX Series Switch on page 720](#)
- [Troubleshooting Software Installation on page 785](#)
- [show system storage partitions \(EX Series Switches Only\) on page 782](#)

## Resilient Dual-Root Partitions Frequently Asked Questions



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Resilient Dual-Root Partitions Frequently Asked Questions*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

This FAQ addresses questions regarding resilient dual-root partitions on EX Series switches. The resilient dual-root partition feature was introduced on EX Series switches at Junos OS Release 10.4R3. It provides additional resiliency for EX Series switches.

This FAQ covers the following questions:

- [What Happens to My Files If the System Detects a File System Corruption and Automatic Snapshot is Enabled? on page 789](#)
- [What Happens to My Files If the System Detects a File System Corruption and Automatic Snapshot is Not Enabled? on page 790](#)
- [How Will I Be Informed If My Switch Boots from the Alternate Slice Because of Corruption in the Root File System? on page 791](#)

### What Happens to My Files If the System Detects a File System Corruption and Automatic Snapshot is Enabled?

If the automatic snapshot feature is enabled during a reboot, the system automatically takes a snapshot of Junos OS from the alternate root partition (Slice 2) and copies it onto the primary root partition (Slice 1). The system checks each file system partition for corruption. [Table 76 on page 790](#) shows the action the system takes if corruption is detected and the corrective action that you can take.

**Table 76: Actions If Corrupt Files Are Found and Automatic Snapshot is Enabled**

Slice 1	Slice 2	Slice 3		Slice 4
s1a	s2a	s3e	s3d	s4d
/	/	/var	/var/tmp	/config
(root Junos OS)	(root Junos OS)			
If a root directory (/) is corrupted, the corrupted file system is not mounted. The switch automatically takes a snapshot of the Junos OS root file system and copies it onto the primary root partition. It boots from the alternate slice, but the next reboot happens from the primary slice.		During early boot, the integrity of /var, /var/tmp, and /config files is verified. If they are corrupted, the corrupted slice is reformatted and the file directory in that slice is lost.		
Corrective action: No corrective action is required.		Corrective action: Restore the /var or /config files from the external backup.		

### What Happens to My Files If the System Detects a File System Corruption and Automatic Snapshot is Not Enabled?

During a reboot, the system checks each file system partition for corruption. [Table 77 on page 790](#) shows the action the system takes if corruption is detected and the corrective action that you can take.

**Table 77: Actions If Corrupt Files Are Found**

Slice 1	Slice 2	Slice 3		Slice 4
s1a	s2a	s3e	s3d	s4d
/	/	/var	/var/tmp	/config
(root Junos OS)	(root Junos OS)			
If a root directory (/) is corrupted, the corrupted file system is not mounted and the switch boots from the alternate slice.		During early boot, the integrity of /var, /var/tmp, and /config files is verified. If they are corrupted, the corrupted slice is reformatted and the file directory in that slice is lost.		
Corrective action: Issue a <a href="#">request system snapshot</a> command from the good root directory to the corrupted slice.		Corrective action: Restore the /var or /config files from the external backup.		

## How Will I Be Informed If My Switch Boots from the Alternate Slice Because of Corruption in the Root File System?

If the switch detects corruption in the primary root file system, it boots from the alternate root partition. When this occurs, the type of notification depends on whether you have enabled the automatic snapshot feature or not:

- If the automatic snapshot feature is not enabled:

- If you are logged in through the console port or the management port:

```
WARNING: THIS DEVICE HAS BOOTED FROM THE BACKUP JUNOS IMAGE
```

It is possible that the primary copy of JUNOS failed to boot up properly, and so this device has booted from the backup copy.

Please re-install JUNOS to recover the primary copy in case it has been corrupted.

- The following message is displayed when you issue **show chassis alarms**:

```
user@switch> show chassis alarms
1 alarms currently active
Alarm time          Class  Description
2011-02-17 05:48:49 PST  Minor  Host 0 Boot from backup root
```

- If the automatic snapshot feature is enabled:
  - A banner message appears, indicating that an automatic snapshot operation is in progress. The banner message disappears when the snapshot operation is complete.
  - No alarm is issued to indicate that the switch has been rebooted from the alternate partition. However, the switch does log the event.

### Related Documentation

- [Verifying Junos OS and Boot Loader Software Versions on an EX Series Switch on page 720](#)
- [Troubleshooting Software Installation on page 785](#)
- [Troubleshooting a Switch That Has Booted from the Backup Junos OS Image on page 788](#)
- [Verifying Junos OS and Boot Loader Software Versions on an EX Series Switch on page 720](#)





## PART 7

# System Monitoring

- [Overview on page 795](#)
- [Configuration on page 821](#)
- [Administration on page 839](#)



## CHAPTER 21

# Overview

- [Software Overview on page 795](#)
- [Alarms Overview on page 797](#)
- [Dashboard Overview on page 798](#)
- [Hardware/CLI Terminology Mapping Overview on page 817](#)

## Software Overview

---

- [Understanding Software Infrastructure and Processes on page 795](#)

### Understanding Software Infrastructure and Processes

Each switch runs the Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches on its general-purpose processors. Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the chassis.

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

With the J-Web interface and the command-line interface (CLI) to the Junos OS, you configure switching features and routing protocols and set the properties of network interfaces on your switch. After activating a software configuration, use either the J-Web or CLI user interface to monitor the switch, manage operations, and diagnose protocol and network connectivity problems.

- [Routing Engine and Packet Forwarding Engine on page 795](#)
- [Junos OS Processes on page 796](#)

### Routing Engine and Packet Forwarding Engine

---

A switch has two primary software processing components:

- **Packet Forwarding Engine**—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.
- **Routing Engine**—Provides three main functions:

- Creates the packet forwarding switch fabric for the switch, providing route lookup, filtering, and switching on incoming data packets, then directing outbound packets to the appropriate interface for transmission to the network
- Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
- Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

### Junos OS Processes

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

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

[Table 6 on page 30](#) describes the primary Junos OS processes.

**Table 78: Junos OS Processes**

Process	Name	Description
Chassis process	chassisd	<p>Detects hardware on the system that is used to configure network interfaces.</p> <p>Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered.</p> <p>Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.</p>
Ethernet switching process	eswd	<p>Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree protocol and access port security. The process is also responsible for managing Ethernet switching interfaces, VLANs, and VLAN interfaces.</p> <p>Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.</p>
Forwarding process	pfem	<p>Defines how routing protocols operate on the switch. The overall performance of the switch is largely determined by the effectiveness of the forwarding process.</p>
Interface process	dcd	<p>Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.</p>

Table 78: Junos OS Processes (*continued*)

Process	Name	Description
Management process	mgd	<p>Provides communication between the other processes and an interface to the configuration database.</p> <p>Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured.</p> <p>Interacts with the other processes when commands are issued through one of the user interfaces on the switch.</p> <p>If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.</p>
Routing protocol process	rpd	Defines how routing protocols such as RIP, OSPF, and BGP operate on the device, including selecting routes and maintaining forwarding tables.

#### Related Documentation

- [For more information about processes, see \*Junos OS Network Operations Guide\*](#)
- [For more information about basic system parameters, supported protocols, and software processes, see \*Junos OS System Basics Configuration Guide\*](#)

## Alarms Overview

- [Understanding Alarm Types and Severity Levels on EX Series Switches on page 797](#)

### Understanding Alarm Types and Severity Levels on EX Series Switches



**NOTE:** This topic applies only to the J-Web Application package.

Alarms alert you to conditions that might prevent normal operation of the switch. Before monitoring alarms on a Juniper Networks EX Series Ethernet switch, become familiar with the terms defined in [Table 79 on page 797](#).

Table 79: Alarm Terms

Term	Definition
<b>alarm</b>	Signal alerting you to conditions that might prevent normal operation. On a switch, the alarm signal is the <b>ALM</b> LED lit on the front of the chassis.
<b>alarm condition</b>	Failure event that triggers an alarm.
<b>alarm severity</b>	Seriousness of the alarm. If the Alarm ( <b>ALM</b> ) LED is red, this indicates a major alarm. If the Alarm LED is yellow, this indicates a minor alarm. If the Alarm LED is unlit, there is no alarm or the switch is halted.

Table 79: Alarm Terms (*continued*)

Term	Definition
<b>chassis alarm</b>	Preset alarm triggered by a physical condition on the switch such as a power supply failure, excessive component temperature, or media failure.
<b>system alarm</b>	Preset alarm triggered by a missing rescue configuration or failure to install a license for a licensed software feature.  <b>NOTE:</b> On EX6200 switches, a system alarm can be triggered by an internal link error.

### Alarm Types

The switch supports these alarms:

- Chassis alarms indicate a failure on the switch or one of its components. Chassis alarms are preset and cannot be modified.
- System alarms indicate a missing rescue configuration. System alarms are preset and cannot be modified, although you can configure them to appear automatically in the J-Web interface display or the CLI display.

### Alarm Severity Levels

Alarms on switches have two severity levels:

- Major (red)—Indicates a critical situation on the switch that has resulted from one of the following conditions. A red alarm condition requires immediate action.
  - One or more hardware components have failed.
  - One or more hardware components have exceeded temperature thresholds.
  - An alarm condition configured on an interface has triggered a critical warning.
- Minor (yellow or amber)—Indicates a noncritical condition on the switch that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.

A missing rescue configuration generates a yellow system alarm.

- Related Documentation**
- [Checking Active Alarms with the J-Web Interface on page 842](#)
  - [Dashboard for EX Series Switches on page 799](#)

## Dashboard Overview

- [Dashboard for EX Series Switches on page 799](#)

## Dashboard for EX Series Switches



**NOTE:** This topic applies only to the J-Web Application package.

When you log in to the J-Web user interface, the dashboard for the Juniper Networks EX Series Ethernet Switches appears. Use the dashboard to view system information.

The Update Available window appears if there is a latest update of the J-Web Application package available on the Juniper Networks server. This window is enabled by the auto update feature of J-Web.



**NOTE:**

- The Update Available window will *not* appear when you log in, if you have not selected the Check for updates automatically on every login in the *Update Preference* section in the Maintain > Update J-Web side pane. By default, the *Check for update automatically on every login* is selected.
- If you choose *Update Later*, you can update to the latest J-Web Application package by clicking the orange icon next to *Update Available* on the top pane of the J-Web interface or through Maintain > Update J-Web.

The dashboard comprises a graphical chassis viewer and four panels.

This topic describes:

- [Graphical Chassis Viewer on page 799](#)
- [System Information Panel on page 801](#)
- [Health Status Panel on page 803](#)
- [Capacity Utilization Panel on page 805](#)
- [Alarms Panel on page 805](#)
- [File System Usage on page 806](#)
- [Chassis Viewer on page 806](#)

### Graphical Chassis Viewer

The Dashboard panel displays a graphical view of the chassis of a switch. In a Virtual Chassis, it displays a graphical view of each member switch.

In a Virtual Chassis, the default values are shown on the Dashboard panel when no chassis image is clicked. The panel displays the value for a switch if you click its image.



**NOTE:** If the member switch is not present, inactive, or not provisioned, you cannot expand the member switch image.

Table 80 on page 800 lists the details that are displayed on each member switch.

**Table 80: Details of a Virtual Chassis Member Switch**

Details	Example
Model number of the member switch	<b>EX3300</b>
Assigned ID that applies to the entire Virtual Chassis configuration	<b>ID 2</b>  <b>NOTE:</b> If the member switch is not provisioned, the serial number of the switch is displayed instead of its ID.
Role of the member switch	<b>Master</b>  Possible roles are: <b>Master</b> , <b>Backup</b> , or <b>Linecard</b>
Status of the member switch	<b>Prsnt</b>  Possible statuses are: <b>Prsnt</b> , <b>NotPrsnt</b> , <b>Inactive</b> , or <b>Unprvsnd</b>

The status of the member switch is displayed on the image of the switch. If the member switch appears dimmed, it means the switch is not present, is inactive, or is not provisioned in the Virtual Chassis. If the member switch does not appear dimmed, it means the switch is present and is active.

Table 81 on page 800 describes the possible status of a member switch.

**Table 81: Status of a Member Switch in a Virtual Chassis**

If the member switch is	It appears as	It means the member switch
Present	<b>Prsnt</b>	Has established physical and logical connections with Virtual Chassis member switches.
Not present	dimmed and <b>NotPrsnt</b>	Has been disconnected from the existing Virtual Chassis.
Inactive	dimmed and <b>Inactive</b>	Has established physical connections, but is unable to establish logical connections.
Not provisioned	dimmed and <b>Unprvsnd</b>	Cannot synchronize with the existing preprovisioned Virtual Chassis.

Click **Rear View** for a graphical view of the rear panel of the switch.

Click **Preferences** to choose which panels must be displayed and set the refresh interval for chassis viewer information. Click **OK** to save your changes and return to the dashboard or click **Cancel** to return to the dashboard without saving changes.





**NOTE:** You can drag the various panels to different locations in the J-Web window.

### System Information Panel

Table 82: System Information

Field	Description
System name	Indicates the local name of the EX Series switch. The local name of the EX Series switches changes when an individual image is clicked.
Device model	<p>Indicates the model of the EX Series switch. In a Virtual Chassis configuration, to indicate the model of a switch, click the image of that switch.</p> <p><b>NOTE:</b> In a Virtual Chassis setup for an EX6210, EX8208, or EX8216 switch, the Device model field displays details of the master Routing Engine. To view details of a member, select it.</p>

Table 82: System Information (*continued*)

Field	Description
Inventory details	<p>Indicates the following:</p> <ul style="list-style-type: none"> <li>For EX3200 switches; and for EX2200, EX2200-C, EX3300, EX4200, EX4300, EX4500, and EX4550 switches that are not configured as Virtual Chassis, the value displayed in Inventory details field is always 1 FPC. FPC is a legacy term for a slot in a large Juniper Networks chassis; which simply refers to the standalone switch.</li> <li>For EX2200 and EX2200-C switches configured as a Virtual Chassis, the value displayed in the Inventory details field is 1–4 FPC, with the number corresponding to the number of member switches.</li> <li>For EX3300 switches configured as a Virtual Chassis, the value displayed in the Inventory details field is 1–6 FPC, with the number corresponding to the number of member switches.</li> </ul> <p><b>NOTE:</b> For Junos OS Release 14.1X53-D10 and later, EX3300 switches configured as a Virtual Chassis display the value 1–10 FPC in the Inventory details field.</p> <ul style="list-style-type: none"> <li>For EX4200, EX4500, and EX4550 switches configured as a Virtual Chassis, the value displayed in the Inventory details field is 1–10 FPC, with the number corresponding to the number of member switches.</li> <li>For EX6210 switches, the values displayed in the Inventory details field are 1–2 CB and 1–9 FPC. CB, or Control Board, refers to the SRE module. FPC refers to line cards and the FPC within the CB.</li> <li>For an EX8208 switch, the values displayed in Inventory details field are 1–3 CB and 0–8 FPC. CB, or Control Board, refers to SRE and SF modules. FPC refers to line cards.</li> <li>For EX8216 switches, the values displayed in Inventory details field are 1–2 CB and 0–16 FPC. CB, or Control Board, refers to RE modules and FPC refers to line cards.</li> <li>For an XRE200 External Routing Engine in an EX8200 Virtual Chassis, the value displayed in Inventory details is 1 XRE. XRE refers to RE modules. For XRE200 External Routing Engines configured as a Virtual Chassis, the values displayed in Inventory details are 1–2 XRE and 0–4 LCC, where LCC refers to the EX8200 line card chassis.</li> </ul>
Junos image	Indicates the version of the Junos OS image. In a Virtual Chassis configuration, the Junos OS image of the master switch is displayed by default. To display the Junos OS image of a specific switch, click the image of that switch.
Boot image	Indicates the version of the boot image that is used. In a Virtual Chassis configuration, the boot image of the master switch is displayed by default. To display the boot image of a specific switch, click the image of that switch.

Table 82: System Information (*continued*)

Field	Description
Device uptime	Indicates the time since the last reboot. In a Virtual Chassis configuration, to display the uptime of the specific switch, click the image of that switch.
Last configured time	Indicates the time when the switch was last configured.

### Health Status Panel

Table 83: Health Status

Field	Description
<b>EX2200, EX2200-C, EX3200, EX3300, EX4200, and EX4300 Switches</b>	
Memory util.	<p>Indicates the memory used in the Routing Engine. In a Virtual Chassis configuration, the memory utilization value of the master Routing Engine is displayed.</p> <p><b>NOTE:</b> In EX4300 Virtual Chassis, to display the Routing Engine memory utilization of the master or backup, click the respective image.</p>
Flash	<p>Indicates the usage and capacity of internal flash memory and any external USB flash drive.</p> <p><b>NOTE:</b> In EX4300 Virtual Chassis, the flash memory utilization of the master switch is displayed by default. To display the flash memory utilization along with the internal and external flash memory utilization details for each switch or line card, mouse over individual switch or line card images.</p>
Temp.	<p>Indicates the chassis temperature status. Temperatures are listed in Celsius and the corresponding Fahrenheit values.</p> <p><b>NOTE:</b> The <b>Temp</b> field is unavailable for a standalone EX2200-C switch.</p> <p>The <b>Temp</b> field is dynamically available for an EX2200 Virtual Chassis switch based on the model of the member clicked.</p> <p><b>NOTE:</b> In EX4300 Virtual Chassis, the temperature of the master Routing Engine is displayed by default. To display the temperature of the Routing Engine of any switch, click the image of that switch.</p>
CPU load	<p>Indicates the average CPU usage over 15 minutes. In a Virtual Chassis configuration, on loading the master or backup switch, the CPU load for that switch's Routing Engine is displayed by default. To display the CPU load for a specific switch's Routing Engine, click the image of that switch.</p>
Fan status	<p>Indicates the status of the fans in the fan tray. The possible values are <b>OK</b>, <b>Failed</b>, and <b>Absent</b>. In a Virtual Chassis configuration, the fan status of the master switch is displayed by default. To display the fan status for any switch, click the image of that switch.</p> <p><b>NOTE:</b> The <b>Fan status</b> field is unavailable for a standalone EX2200-C switch.</p> <p>The <b>Fan status</b> field is dynamically available for an EX2200 Virtual Chassis switch based on the model of the member clicked.</p>

### EX4500 and EX4550 Switches

Table 83: Health Status (*continued*)

Field	Description
Memory util.	Indicates the memory used in the Routing Engine. In a Virtual Chassis configuration, the memory utilization value of the master Routing Engine is displayed.
Flash	Indicates the usage and capacity of internal flash memory and any external USB flash drive.
Temp.	Indicates the chassis temperature status. Temperatures in the dashboard are listed in Celsius and the corresponding Fahrenheit values.  <b>NOTE:</b> The <b>Temp</b> field is unavailable for an EX4500 switch.
CPU load	Indicates the average CPU usage over 15 minutes.
Fan status	Indicates the status of the fans in the fan tray. The possible values are <b>OK</b> , <b>Failed</b> , and <b>Absent</b> . This field also indicates the direction of airflow of the fan tray. The possible values are <b>Front to back</b> and <b>Back to front</b> .
<b>EX6210 Switches</b>	
Memory util.	Indicates the memory used in the master Routing Engine. Click the <b>backup Routing Engine</b> to view the memory used in the backup Routing Engine.
CPU load	Indicates the average CPU usage over 15 minutes.
Flash	Indicates the usage and capacity of internal flash memory and any external USB flash drive.
Fan status	Indicates the status of the fans in the fan tray. The possible values are <b>OK</b> , <b>Failed</b> , and <b>Absent</b> .
<b>EX8208 Switches</b>	
Memory util.	Indicates the memory used in the external Routing Engine. In an EX8200 Virtual Chassis, the memory utilization value of the XRE200 External Routing Engine in the master role is displayed. Click the <b>XRE200 External Routing Engine</b> in the backup role to view the memory used in the backup external Routing Engine.
CPU load	Indicates the average CPU usage over 15 minutes.
Flash	Indicates the usage and capacity of internal flash memory and any external USB flash drive.
<b>EX8216 Switches</b>	
Memory util.	Indicates the memory used in the external Routing Engine. In an EX8200 Virtual Chassis, the memory utilization value of the XRE200 External Routing Engine in the master role is displayed. Click the <b>XRE200 External Routing Engine</b> in the backup role to view the memory used in the backup external Routing Engine.
CPU load	Indicates the average CPU usage over 15 minutes.
Flash	Indicates the usage and capacity of internal flash memory and any external USB flash drive.
<b>XRE200 External Routing Engines</b>	

Table 83: Health Status (*continued*)

Field	Description
Memory util.	Indicates the memory used in the external Routing Engine. In an EX8200 Virtual Chassis, the memory utilization value of the XRE200 External Routing Engine in the master role is displayed. Click the backup XRE200 External Routing Engine to view the memory used in backup external Routing Engine.
CPU load	Indicates the average CPU usage over 15 minutes.
Flash	Indicates the usage and capacity of internal flash memory and any external USB flash drive.
Fan Status	Indicates the status of the fans in the fan tray. The possible values are <b>OK</b> , <b>Failed</b> , and <b>Absent</b> .

### Capacity Utilization Panel

Table 84: Capacity Utilization

Field	Description
Number of active ports	Indicates the number of active ports in the switch. Configured Virtual Chassis ports (VCPs) are considered as active ports.
Total number of ports	Indicates the number of ports in the switch.  <b>NOTE:</b> In EX3300 Virtual Chassis, the total number of ports of all of the switches is displayed.
Used-up MAC-Table entries	Indicates the number of MAC table entries.
Supported MAC-Table entries	Indicates the maximum number of MAC table entries permitted.
Number of VLANs configured	Indicates the number of VLANs configured.  <b>NOTE:</b> Only tagged VLANs are counted.
Number of VLANs supported	Indicates the maximum number of VLANs supported.

### Alarms Panel

Displays information about the last five alarms raised in the system. For example, if there are 5 major alarms, then details of all 5 major alarms are displayed. If there are 4 major alarms and 3 minor alarms, then details of the 4 major alarms and 1 minor alarm are displayed. Major alarms are displayed in red and minor alarms are displayed in yellow.

In an EX8200 Virtual Chassis, the top 5 alarms for the master external Routing Engine are displayed by default. If you select an EX8200 member switch of the Virtual Chassis, the top 5 alarms for that member switch are displayed.

## File System Usage

To display the file system storage details of a switch in the backup or linecard role, click the image of that switch.

## Chassis Viewer

Click the **Rear View** button to see the back of the chassis image. Click the **Front View** button to see the front of the chassis image. In a Virtual Chassis configuration, the **Rear View** button is disabled if the switch is not selected.

- [Table 85 on page 806](#)—Describes the chassis viewer for EX2200 switches.
- [Table 86 on page 807](#)—Describes the chassis viewer for EX2200-C switches.
- [Table 87 on page 807](#)—Describes the chassis viewer for EX3200, EX3300, and EX4200 switches.
- [Table 88 on page 809](#)—Describes the chassis viewer for EX4300 switches.
- [Table 89 on page 810](#)—Describes the chassis viewer for EX4500 switches.
- [Table 90 on page 811](#)—Describes the chassis viewer for EX4550 switches.
- [Table 91 on page 813](#)—Describes the chassis viewer for EX6210 switches.
- [Table 92 on page 813](#)—Describes the chassis viewer for EX8208 switches.
- [Table 93 on page 815](#)—Describes the chassis viewer for EX8216 switches.
- [Table 94 on page 815](#)—Describes the chassis viewer for the XRE200 External Routing Engines.

**Table 85: Chassis Viewer for EX2200 Switches**

Field	Description
<b>Front View</b>	
Interface status	<p>In the image, the following colors denote the interface status:</p> <ul style="list-style-type: none"> <li>• Green—Interface is up and operational.</li> <li>• Yellow—Interface is up but is nonoperational.</li> <li>• Gray—Interface is down and nonoperational.</li> </ul> <p>Mouse over the interface (port) to view more information.</p>
<b>Rear View</b>	
Management (me0) port	The management port is used to connect the switch to a management device for out-of-band management.
Console port	The console port is used to connect the switch to a management console or to a console server. (You might do this for initial switch configuration.)

Table 85: Chassis Viewer for EX2200 Switches (*continued*)

Field	Description
USB port	Indicates the USB port for the switch.  <b>NOTE:</b> We recommend that you use USB flash drives purchased from Juniper Networks for your EX Series switch.
Fan tray	Mouse over the fan tray icon to display name, status, and description information.
Power supply	Mouse over the power outlet icon to display name, status, and description information.

Table 86: Chassis Viewer for EX2200-C Switches

Field	Description
<b>Front View</b>	
Interface status	In the image, the following colors denote the interface status: <ul style="list-style-type: none"> <li>Green—Interface is up and operational.</li> <li>Yellow—Interface is up but is nonoperational.</li> <li>Gray—Interface is down and nonoperational.</li> </ul> Mouse over the interface (port) to view more information.
Management (me0) port	The management port is used to connect the switch to a management device for out-of-band management.
Console port	The console port is used to connect the switch to a management console or to a console server. (You might do this for initial switch configuration.)
USB port	Indicates the USB port for the switch.  <b>NOTE:</b> We recommend that you use USB flash drives purchased from Juniper Networks for your EX Series switch.
<b>Rear View</b>	
Power supply	Mouse over the power outlet icon to display name, status, and description information.

Table 87: Chassis Viewer for EX3200, EX3300, and EX4200 Switches

Field	Description
<b>Front View</b>	

Table 87: Chassis Viewer for EX3200, EX3300, and EX4200 Switches (*continued*)

Field	Description
Interface status	<p>In the image, the following colors denote the interface status:</p> <ul style="list-style-type: none"> <li>Green—Interface is up and operational.</li> <li>Yellow—Interface is up but is nonoperational.</li> <li>Gray—Interface is down and nonoperational.</li> </ul> <p>Mouse over the interface (port) to view more information.</p> <p>For a Virtual Chassis configuration, select the switch to view the interface status.</p> <p>If an SFP+ uplink module is installed in the switch, mouse over the port icon to display whether the module is configured to operate in 1-gigabit mode or in 10-gigabit mode. If the module is configured to operate in 1-gigabit mode, the tool tip information is displayed for all 4 ports. If the module is configured to operate in 10-gigabit mode, the tool tip information is displayed only for 2 ports.</p> <p>On an EX3300 switch with the 4x GE/XE SFP+ module, mouse over the port icon to display whether the module is configured to operate in 1-gigabit mode or 10-gigabit mode.</p> <p>For SFP, SFP+, and XFP ports, the interfaces appear dimmed if no transceiver is inserted. The chassis viewer displays <b>Transceiver not plugged-in</b> when you mouse over the port icon.</p>
LCD panel	LCD panel configured for the LEDs on the ports. Mouse over the icon to view the current character display.
<b>Rear View of the EX3200 Switch</b>	
Management (me0) port	The management port is used to connect the switch to a management device for out-of-band management.
Console port	The console port is used to connect the switch to a management console or to a console server. (You might do this for initial switch configuration.)
USB port	<p>Indicates the USB port for the switch.</p> <p><b>NOTE:</b> We recommend that you use USB flash drives purchased from Juniper Networks for your EX Series switch.</p>
Fan tray	Mouse over the fan tray icon to display name, status, and description information.
Power supply	Mouse over the power supply icon to display name, status, and description information.
<b>Rear View of the EX3300 and EX4200 Switch</b>	
Fan tray	Mouse over the fan tray icon to display name, status, and description information. For a Virtual Chassis, the status of the fans of the selected member switch is displayed.
Virtual Chassis port	<p>Displayed only when EX4200 switches are configured as a Virtual Chassis. The following colors denote the Virtual Chassis port (VCP) status:</p> <ul style="list-style-type: none"> <li>Green—VCP is up and operational.</li> <li>Yellow—VCP is up but is nonoperational.</li> <li>Gray—VCP is down and nonoperational.</li> </ul>



Table 87: Chassis Viewer for EX3200, EX3300, and EX4200 Switches (*continued*)

Field	Description
USB port	Indicates the USB port for the switch.  <b>NOTE:</b> We recommend that you use USB flash drives purchased from Juniper Networks for your EX Series switch.
Management (me0) port	The management port is used to connect the switch to a management device for out-of-band management.
Console port	The console port is used to connect the switch to a management console or to a console server. (You might do this for initial switch configuration.)
Power supplies	Mouse over the power supply icons to display name, status, and description information.

Table 88: Chassis Viewer for EX4300 Switches

Field	Description
<b>Front View</b>	
Interface status	In the image, the colors listed below denote the interface status for both copper and fiber media type of ports: <ul style="list-style-type: none"> <li>Green—Interface is up and operational.</li> <li>Yellow—Interface is up but is nonoperational.</li> <li>Gray—Interface is down and nonoperational.</li> </ul> Mouse over the interface (port) to view more information.
LCD panel	LCD panel configured for the LEDs on the ports. Mouse over the icon to view the current character display.
Mini USB console	The mini console port is used to connect the switch to the management console.
PIC 2 slot	You can install an uplink module in the PIC 2 slot. Mouse over the ports in the module to view the details of the ports in module.  24-port and 48-port EX4300 switches support the 4-port 10-Gigabit SFP+ uplink module.  EX4300-32F switches support the 2-port 40-Gigabit QSFP+ uplink module and the 8-port 10-Gigabit SFP+ uplink module.  When you install a transceiver in the port, the following colors denote the interface status: <ul style="list-style-type: none"> <li>Green—Interface is up and operational.</li> <li>Yellow—Interface is up but is not operational.</li> <li>Gray—Interface is down and not operational.</li> </ul>

**NOTE:** In EX4300 switches the LEDs are seen in the front panel, these are not active.

#### Rear View of the EX4300 Switch

Management port	The management port is used to connect the switch to a management device for out-of-band management.
-----------------	--

Table 88: Chassis Viewer for EX4300 Switches (*continued*)

Field	Description
Console port	The Console port (RJ-45) is used to connect the switch to a management console or to a console server.
USB port	Indicates the USB port for the switch.  <b>NOTE:</b> We recommend that you use USB flash drives purchased from Juniper Networks for your EX Series switch.
Fan tray	Mouse over the fan tray icons to display name, status, and description information.
Power supplies	Mouse over the power supply icons to display name, status, and description information.
PIC 1 slot	<p>The rear panel of a 24-port and a 48-port EX4300 switch has four (built-in) 40-Gigabit QSFP+ ports, and the rear panel of an EX4300-32F switch has two (built-in) 40-Gigabit QSFP+ ports, in which you can install QSFP+ transceivers. Mouse over the ports to view the details of the ports.</p> <p>After you install a transceiver in the port, the following colors denote the interface status:</p> <ul style="list-style-type: none"> <li>• Green—Interface is up and operational.</li> <li>• Yellow—Interface is up but is not operational.</li> <li>• Gray—Interface is down and not operational.</li> </ul> <p>For QSFP+ ports, the interfaces appear dimmed if no transceiver is inserted. The chassis viewer displays <b>Transceiver not plugged in</b> when you mouse over the port.</p> <p>When a QSFP+ port is configured as a Virtual Chassis Port (VCP), the following colors denote the VCP status:</p> <ul style="list-style-type: none"> <li>• Green—VCP is up and operational.</li> <li>• Yellow—VCP is up but is not operational.</li> <li>• Gray—VCP is down and not operational.</li> </ul>

Table 89: Chassis Viewer for EX4500 Switches

Field	Description
<b>Front View</b>	
Interface status	<p>In the image, the colors listed below denote the interface status:</p> <ul style="list-style-type: none"> <li>• Green—Interface is up and operational.</li> <li>• Yellow—Interface is up but is nonoperational.</li> <li>• Gray—Interface is down and nonoperational.</li> </ul> <p>Mouse over the interface (port) to view more information.</p> <p>For a Virtual Chassis configuration, select the switch to view the interface status.</p> <p>If an SFP+ uplink module is installed in the switch, mouse over the interface (ports) on the module for more information.</p> <p>For SFP and SFP+ ports, the interfaces appear dimmed if no transceiver is inserted. The chassis viewer displays <b>Transceiver not plugged-in</b> when you mouse over the port icon.</p>

Table 89: Chassis Viewer for EX4500 Switches (*continued*)

Field	Description
LCD panel	LCD panel configured for the LEDs on the ports. Mouse over the icon to view the current character display.
Console port	The console port is used to connect the switch to a management console or to a console server.
Management (me0) port	The management port is used to connect the switch to a management device for out-of-band management. Use this port for initial switch configuration.
USB port	Indicates the USB port for the switch.  <b>NOTE:</b> We recommend that you use USB flash drives purchased from Juniper Networks for your EX Series switch.
<b>Rear View of the EX4500 Switch</b>	
Fan tray	Mouse over the fan tray icon to display status of the fans and airflow direction information. For a Virtual Chassis, the status of the fans of the selected member switch is displayed.
Virtual Chassis port	Displayed only when switches are configured as a Virtual Chassis. The colors listed below denote the Virtual Chassis port (VCP) status: <ul style="list-style-type: none"> <li>• Green—VCP is up and operational.</li> <li>• Yellow—VCP is up but is nonoperational.</li> <li>• Gray—VCP is down and nonoperational.</li> </ul>
Power supplies	Mouse over the power supply icons to display name, status, and description information.
Intraconnect module	Mouse over the module to display details of the intraconnect module. The intraconnect module helps the switch achieve line rate on all its ports.
Virtual Chassis module	Mouse over to display details of the switches in the Virtual Chassis configuration.

Table 90: Chassis Viewer for EX4550 Switches

Field	Description
<b>Front View</b>	

Table 90: Chassis Viewer for EX4550 Switches (*continued*)

Field	Description
Interface status	<p>In the image, the colors listed below denote the interface status:</p> <ul style="list-style-type: none"> <li>Green—Interface is up and operational.</li> <li>Yellow—Interface is up but is nonoperational.</li> <li>Gray—Interface is down and nonoperational.</li> </ul> <p>Mouse over the interface (port) to view more information.</p> <p>For a Virtual Chassis configuration, select the switch to view the interface status.</p> <p>If an expansion module or a Virtual Chassis module is installed in the switch, mouse over the interface (ports) on the module for more information.</p> <p>On an EX4550-32F switch, for SFP and SFP+ ports, the interfaces appear dimmed if no transceiver is inserted. The chassis viewer displays <b>Transceiver (1G/10G) not plugged in</b> when you mouse over the port icon.</p>
LCD panel	LCD panel configured for the LEDs on the ports. Mouse over the icon to view the current character display.
Console port	The console port is used to connect the switch to a management console or to a console server.
Mini Console port	The mini console port is used to connect the switch to the management console.
Management (me0) port	The management port is used to connect the switch to a management device for out-of-band management. Use this port for initial switch configuration.
PIC1 slot	You can insert an uplink module or a Virtual Chassis module in the PIC1 slot. Mouse over to display the details of the module inserted (uplink or Virtual Chassis).
USB port	<p>Indicates the USB port for the switch.</p> <p><b>NOTE:</b> We recommend that you use USB flash drives purchased from Juniper Networks for your EX Series switch.</p>
<b>Rear View of the EX4550 Switch</b>	
Fan tray	Mouse over the fan tray icon to display the status of the fans and airflow direction information. For a Virtual Chassis, the status of the fans of the selected member switch is displayed.
Virtual Chassis port	<p>Displayed only when switches are configured as a Virtual Chassis. In the image, the colors listed below denote the Virtual Chassis port (VCP) status:</p> <ul style="list-style-type: none"> <li>Green—VCP is up and operational.</li> <li>Yellow—VCP is up but is nonoperational.</li> <li>Gray—VCP is down and nonoperational.</li> </ul>
Power supplies	Mouse over the power supply icons to display name, status, and description information.
PIC2 slot	You can insert an uplink module or a Virtual Chassis module into the PIC2 slot. Mouse over to display the details of the module inserted (uplink or Virtual Chassis).

Table 91: Chassis Viewer for EX6210 Switches

Field	Description
<b>Front View</b>	
Temperature	Mouse over the temperature icon to display the temperature of the CB or line card.
Interface status	<p>Select the CB or line card.</p> <p>In the image, the colors listed below denote the interface status:</p> <ul style="list-style-type: none"> <li>• Green—Interface is up and operational.</li> <li>• Yellow—Interface is up but is nonoperational.</li> <li>• Gray—Interface is down and nonoperational.</li> </ul> <p>Mouse over the interface (port) to view more information.</p> <p>You can view status for the following ports on the SRE module:</p> <ul style="list-style-type: none"> <li>• USB port—Indicates the USB port for the switch.</li> </ul> <p><b>NOTE:</b> We recommend that you use USB flash drives purchased from Juniper Networks for your EX Series switch.</p> <ul style="list-style-type: none"> <li>• Management (<b>me0</b>) port—The management port is used to connect the switch to a management device for out-of-band management. There are 2 management ports: fiber and copper. The same status is displayed for both the <b>me0</b> ports.</li> <li>• Console port—The console port is used to connect the switch to a management console or to a console server. (You might do this for initial switch configuration.)</li> </ul> <p>CBs support 4 SFP+ uplink ports. Mouse over the interface on the CB for more information.</p> <p>For SFP and SFP+ ports, the interfaces appear dimmed if no transceiver is inserted. The chassis viewer displays <b>Transceiver not plugged-in</b> when you mouse over the port icon.</p>
Power supplies	Mouse over the power supply icons to display name, status, and description information.
LCD panel	LCD panel configured for the LEDs on the ports. Mouse over the icon to view the current character display of the master Routing Engine. The EX6210 switch has 2 LCD panels, one for each Routing Engine. The backup Routing Engine LCD displays <b>Backup</b> .
<b>Rear View of the EX6210 Switch</b>	
Fan tray	Mouse over the fan tray icon to display information regarding the cooling fans.

Table 92: Chassis Viewer for EX8208 Switches

Field	Description
<b>Front View</b>	

Table 92: Chassis Viewer for EX8208 Switches (*continued*)

Field	Description
Interface status	<p>In the image, click any line card, SRE module, or SF module to view the front view of the selected component. In the image, the colors listed below denote the interface status:</p> <ul style="list-style-type: none"> <li>• Green—Interface is up and operational.</li> <li>• Yellow—Interface is up but is nonoperational.</li> <li>• Gray—Interface is down and nonoperational.</li> </ul> <p>Mouse over the interface (port) to view more information.</p> <p>You can view status for the following ports on the SRE module:</p> <ul style="list-style-type: none"> <li>• USB port—Indicates the USB port for the switch.</li> </ul> <p><b>NOTE:</b> We recommend that you use USB flash drives purchased from Juniper Networks for your EX Series switch.</p> <ul style="list-style-type: none"> <li>• Auxiliary port—This port is unavailable.</li> <li>• Management (<b>me0</b>) port—The management port is used to connect the switch to a management device for out-of-band management.</li> <li>• Console port—The console port is used to connect the switch to a management console or to a console server. (You might do this for initial switch configuration.)</li> </ul> <p>Because the SF module has no ports, no status information is displayed.</p>
Slot numbers	<p>Slots on the switch are labeled, from the top of the switch down:</p> <ul style="list-style-type: none"> <li>• 0–3 (line cards)</li> <li>• SRE0, SF, SRE1 (SRE and SF modules)</li> <li>• 4–7 (line cards)</li> </ul>
Temperature	<p>The active slots contain a gray temperature icon. Mouse over the icon to display temperature information for the slot.</p>
Fan status	<p>Mouse over the fan tray icon to display name, status, and description information.</p>
Power supplies	<p>Mouse over the power supply icons to display name, status, and description information.</p>
LCD panel	<p>LCD panel configured for the LEDs on the ports. Mouse over the icon to view the current character display.</p>
Rear View	<p>The EX8208 switch does not have any components on the rear of the chassis.</p>

Table 93: Chassis Viewer for EX8216 Switches

Field	Description
<b>Front View</b>	
Interface status	<p>In the image, click any line card or RE module to display the front view of the selected component. In the image, the colors listed below denote the interface status:</p> <ul style="list-style-type: none"> <li>Green—Interface is up and operational.</li> <li>Yellow—Interface is up but is nonoperational.</li> <li>Gray—Interface is down and nonoperational.</li> </ul> <p>Mouse over the interface (port) to view more information.</p> <p>You can view status for the following ports on the RE module:</p> <ul style="list-style-type: none"> <li>USB port—Indicates the USB port for the switch.</li> </ul> <p><b>NOTE:</b> We recommend that you use USB flash drives purchased from Juniper Networks for your EX Series switch.</p> <ul style="list-style-type: none"> <li>Auxiliary port—This port is unavailable.</li> <li>Management (<b>me0</b>) port—The management port is used to connect the switch to a management device for out-of-band management.</li> <li>Console port—The console port is used to connect the switch to a management console or to a console server. (You might do this for initial switch configuration.)</li> </ul>
Slot numbers	<p>Slots on the switch are labeled, from the top of the switch down:</p> <ul style="list-style-type: none"> <li>RE0 (RE module)</li> <li>RE1 (RE module)</li> <li>0–15 (line cards)</li> </ul>
Temperature	The active slots contain a gray temperature icon. Mouse over the icon to display temperature information for the slot.
Fan status	Mouse over the fan tray icon to display consolidated information about the fans.
Power supplies	Mouse over the power supply icons to display name, status, and description information.
LCD panel	LCD panel configured for the LEDs on the ports. Mouse over the icon to view the current character display.
<b>Rear View</b>	
SF modules	Mouse over the SF module icons in their respective slots to display information. Slots are numbered SF7–SF0, from left to right.

Table 94: Chassis Viewer for XRE200 External Routing Engines

Field	Description
<b>Front View</b>	

Table 94: Chassis Viewer for XRE200 External Routing Engines (*continued*)

Field	Description
Interface status	<p>In the image, the colors listed below denote the interface status:</p> <ul style="list-style-type: none"> <li>Green—Interface is up and operational.</li> <li>Yellow—Interface is up but is nonoperational.</li> <li>Gray—Interface is down and nonoperational.</li> </ul> <p>Mouse over the interface (port) to view more information.</p> <p>For a Virtual Chassis configuration, select the switch to view the interface status.</p>
Console port	The console port is used to connect the switch to a management console or to a console server.
Management (me0) port	The management port is used to connect the switch to a management device for out-of-band management. Use this port for initial switch configuration.
Virtual Chassis port	<p>In the image, the colors listed below denote the Virtual Chassis port (VCP) status:</p> <ul style="list-style-type: none"> <li>Green—VCP is up and operational.</li> <li>Yellow—VCP is up but is nonoperational.</li> <li>Gray—VCP is down and nonoperational.</li> </ul> <p>Mouse over the interface (port) to view more information.</p>
LCD panel	LCD panel configured for the LEDs on the ports. Mouse over the icon to view the current character display.
Temperature	The active slots contain a gray temperature icon. Mouse over the icon to display temperature information for the slot.
USB port	<p>Indicates the USB port for the switch.</p> <p><b>NOTE:</b> We recommend that you use USB flash drives purchased from Juniper Networks for your EX Series switch.</p>
PIC1 slot	You can install a Virtual Chassis module in the PIC1 slot. Mouse over the Virtual Chassis ports to display the port status details.
PIC2 slot	You can install a Virtual Chassis module in the PIC2 slot. Mouse over the Virtual Chassis ports to display the port status details.
<b>Rear View of the XRE200 External Routing Engine</b>	
Fan modules	Mouse over the fan modules to display status of the fans and airflow direction information. For a Virtual Chassis, the status of the fans of the selected member switch is displayed.
Power supplies	Mouse over the power supply icons to display name, status, and description information.

**Related Documentation**

- *J-Web User Interface for EX Series Switches Overview*
- *EX2200 Switches Hardware Overview*
- *EX3200 Switches Hardware Overview*



- [EX3300 Switches Hardware Overview](#)
- [EX4200 Switches Hardware Overview](#)
- [EX4500 Switches Hardware Overview](#)
- [EX6210 Switch Hardware Overview](#)
- [EX8208 Switch Hardware Overview](#)
- [EX8216 Switch Hardware Overview](#)
- [Checking Active Alarms with the J-Web Interface on page 842](#)
- [XRE200 External Routing Engine Documentation](#)

## Hardware/CLI Terminology Mapping Overview

- [EX4300 Switch Hardware and CLI Terminology Mapping on page 817](#)

### EX4300 Switch Hardware and CLI Terminology Mapping

This topic describes the hardware terms used in EX4300 switch documentation and the corresponding terms used in the Junos OS CLI. See [Table 95 on page 817](#).

**Table 95: CLI Equivalents of Terms Used in the Documentation for EX4300 Switches**

Hardware Item (CLI)	Description (CLI)	Value	Item In Documentation	Additional Information
Chassis	One of the following: <ul style="list-style-type: none"> <li>• EX4300-24T</li> <li>• EX4300-24P</li> <li>• EX4300-32F</li> <li>• EX4300-48T</li> <li>• EX4300-48P</li> </ul>	–	Switch chassis	<i>Identifying EX4300 Switch Models</i>
Routing Engine (n)	One of the following: <ul style="list-style-type: none"> <li>• EX4300-24T</li> <li>• EX4300-24P</li> <li>• EX4300-32F</li> <li>• EX4300-48T</li> <li>• EX4300-48P</li> </ul>	n is a value in the range 0 through 9. <ul style="list-style-type: none"> <li>• On a standalone switch, the default value is 0.</li> <li>• On a Virtual Chassis configuration, the values correspond to the member IDs of switches configured in the master role and the backup role in the Virtual Chassis.</li> </ul>	Routing Engine	–

Table 95: CLI Equivalents of Terms Used in the Documentation for EX4300 Switches (*continued*)

Hardware Item (CLI)	Description (CLI)	Value	Item In Documentation	Additional Information
FPC ( <i>n</i> )	Abbreviated name of the Flexible PIC Concentrator (FPC)	<i>n</i> is a value in the range 0 through 9.		<a href="#">“Understanding Interface Naming Conventions on EX Series Switches” on page 2580</a>
	One of the following:	On a standalone switch, the default value is 0.	In this case, FPC refers to the switch itself.	
	<ul style="list-style-type: none"> <li>EX4300-24T</li> <li>EX4300-24P</li> <li>EX4300-32F</li> <li>EX4300-48T</li> <li>EX4300-48P</li> </ul>	On a Virtual Chassis configuration, the values correspond to the assigned member IDs of switches in the Virtual Chassis.	In this case, the FPC number refers to the member ID assigned to the switch.	
PIC ( <i>n</i> )	Abbreviated name of the Physical Interface Card (PIC)	<i>n</i> is a value in the range 0 through 2.		<a href="#">“Understanding Interface Naming Conventions on EX Series Switches” on page 2580</a>
	One of the following:	PIC 0	<ul style="list-style-type: none"> <li>In 24-port and 48-port switches—PIC 0 stands for built-in network ports numbered 0 through 23 or 0 or through 47.</li> <li>In 32-port switches—PIC 0 stands for built-in network ports followed by SFP+ uplink ports. In the CLI output, the network ports are numbered 0 through 31 and the SFP+ uplink ports are numbered 32 through 35.</li> </ul>	
	<ul style="list-style-type: none"> <li>For 24-port and 48-port switches: 4x40GE</li> <li>For 32-port switches: 2x40GE</li> </ul>	PIC 1	Built-in QSFP+ ports	
	<ul style="list-style-type: none"> <li>For 24-port and 48-port switches: 4x 1G/10G SFP/SFP+</li> <li>For 32-port switches, one of the following:               <ul style="list-style-type: none"> <li>8x1G/10G SFP/SFP+</li> <li>2x40GE</li> </ul> </li> </ul>	PIC 2	Uplink module installed in the switch	

Table 95: CLI Equivalents of Terms Used in the Documentation for EX4300 Switches (*continued*)

Hardware Item (CLI)	Description (CLI)	Value	Item In Documentation	Additional Information
Xcvr ( <i>n</i> )	Abbreviated name of the transceiver	<i>n</i> is a value equivalent to the number of the port in which the transceiver is installed.	Optical transceivers	<i>Pluggable Transceivers Supported on EX4300 Switches</i>
Power supply ( <i>n</i> )	One of the following: <ul style="list-style-type: none"> <li>JPSU-350-AC-AFO-A</li> <li>JPSU-350-AC-AFI-A</li> <li>JPSU-550-DC-AFO-A</li> <li>JPSU-550-DC-AFI-A</li> <li>JPSU-715-AC-AFO-A</li> <li>JPSU-1100-AC-AFO-A</li> </ul>	<i>n</i> has a value 0 or 1, corresponding to the power supply slot number.	AC power supply or DC power supply	<ul style="list-style-type: none"> <li>AC Power Supply in EX4300 Switches</li> <li>DC Power Supply in EX4300 Switches</li> </ul>
<p><b>CAUTION:</b> Do not mix:</p> <ul style="list-style-type: none"> <li>AC and DC power supplies in the same chassis.</li> <li>Power supplies with different airflow labels (<b>AIR IN (AFI)</b> and <b>AIR OUT (AFO)</b>) in the same chassis.</li> <li>Fan modules with different airflow labels (<b>AIR IN (AFI)</b> and <b>AIR OUT (AFO)</b>) in the same chassis.</li> <li>Power supplies and fan modules with different airflow labels (<b>AIR IN (AFI)</b> and <b>AIR OUT (AFO)</b>) in the same chassis.</li> </ul> <p><b>NOTE:</b> The 32-port EX4300 switches support fan modules and power supplies with the <b>AIR OUT (AFO)</b> label only.</p>				
Fan tray	One of the following: <ul style="list-style-type: none"> <li>Fan Module, Airflow In (AFI)</li> <li>Fan Module, Airflow Out (AFO)</li> </ul>	<i>n</i> has a value 0 or 1, corresponding to the fan module slot number.	Fan module	<i>Cooling System and Airflow in an EX4300 Switch</i>
<p><b>CAUTION:</b> Do not mix:</p> <ul style="list-style-type: none"> <li>Fan modules with different airflow labels (<b>AIR IN (AFI)</b> and <b>AIR OUT (AFO)</b>) in the same chassis.</li> <li>Power supplies with different airflow labels (<b>AIR IN (AFI)</b> and <b>AIR OUT (AFO)</b>) in the same chassis.</li> <li>Power supplies and fan modules with different airflow labels (<b>AIR IN (AFI)</b> and <b>AIR OUT (AFO)</b>) in the same chassis.</li> <li>AC and DC power supplies in the same chassis.</li> </ul> <p><b>NOTE:</b> The 32-port EX4300 switches support fan modules and power supplies with the <b>AIR OUT (AFO)</b> label only.</p>				

**Related Documentation** • *EX4300 Switches Hardware Overview*



## CHAPTER 22

# Configuration

- [Configuration Statements on page 821](#)

### Configuration Statements

---

- [facility-override on page 822](#)
- [file \(System Logging\) on page 823](#)
- [files on page 824](#)
- [host \(System\) on page 825](#)
- [interface \(Accounting or Sampling\) on page 827](#)
- [log-prefix \(System\) on page 827](#)
- [match on page 828](#)
- [size \(System\) on page 829](#)
- [structured-data on page 830](#)
- [syslog \(System\) on page 831](#)
- [time-format on page 833](#)
- [time-zone on page 834](#)
- [user \(System Logging\) on page 836](#)
- [world-readable \(System\) on page 837](#)

## facility-override

---

<b>Syntax</b>	<code>facility-override <i>facility</i>;</code>
<b>Hierarchy Level</b>	[edit system syslog host]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Substitute an alternate facility for the default facilities used when messages are directed to a remote destination.
<b>Options</b>	<i>facility</i> —Alternate facility to substitute for the default facilities. For a list of the possible facilities, see <i>Alternate Facilities for System Log Messages Directed to a Remote Destination</i> .
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Changing the Alternative Facility Name for System Log Messages Directed to a Remote Destination</i></li><li>• <i>Junos OS System Log Messages Reference</i></li></ul>

## file (System Logging)

<b>Syntax</b>	<pre> file <i>filename</i> {     <i>facility severity</i>;     archive {         <i>files number</i>;         <i>size size</i>;         (no-world-readable   world-readable);     }     explicit-priority;     match "<i>regular-expression</i>";     structured-data {         brief;     } } </pre>
<b>Hierarchy Level</b>	[edit system <a href="#">syslog</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure the logging of system messages to a file.
<b>Options</b>	<p><b><i>facility</i></b>—Class of messages to log. To specify multiple classes, include multiple <b><i>facility severity</i></b> statements. For a list of the facilities, see <i>Junos OS System Logging Facilities and Message Severity Levels</i>.</p> <p><b>file <i>filename</i></b>—File in the <b>severity</b> directory in which to log messages from the specified facility. To log messages to more than one file, include more than one <b>file</b> statement.</p> <p><b><i>severity</i></b>—Severity of the messages that belong to the facility specified by the paired <b><i>facility</i></b> name. Messages with severities of the specified level and higher are logged. For a list of the severities, see <i>Junos OS System Logging Facilities and Message Severity Levels</i>.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Directing System Log Messages to a Log File</i></li> <li>• <i>Junos OS System Log Messages Reference</i></li> </ul>

## files

---

<b>Syntax</b>	<code>files <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit system syslog archive], [edit system syslog file <i>filename</i> archive]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for QFX Series switches.
<b>Description</b>	Configure the maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file <i>logfile</i> , it closes the file, compresses it, and renames it to <i>logfile.0.gz</i> (for information about the maximum file size, see <a href="#">size</a> ). The utility then opens and writes to a new file called <i>logfile</i> . When the new file reaches the maximum size, the <i>logfile.0.gz</i> file is renamed to <i>logfile.1.gz</i> , and the new file is closed, compressed, and renamed <i>logfile.0.gz</i> . By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).
<b>Options</b>	<i>number</i> —Maximum number of archived files. <b>Range:</b> 1 through 1000 <b>Default:</b> 10 files
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Specifying Log File Size, Number, and Archiving Properties</i></li><li>• <i>Junos OS System Log Messages Reference</i></li><li>• <a href="#">size on page 829</a></li></ul>



## host (System)

<b>Syntax</b>	<pre> host (hostname   other-routing-engine) {     facility severity;     exclude-hostname     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     source-address source-address;     structured-data {         brief;     } } </pre>
<b>QFX Series</b>	<pre> host (hostname {     facility severity;     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     port;     source-address source-address; } </pre>
<b>TX Matrix Router and EX Series Switches</b>	<pre> host (hostname   other-routing-engine   scc-master) {     facility severity;     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     port;     source-address source-address; } </pre>
<b>TX Matrix Plus Router</b>	<pre> host (hostname   other-routing-engine   sfc0-master) {     facility severity;     allow-duplicates;     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     port;     source-address source-address; } </pre>
<b>Hierarchy Level</b>	<pre> [edit logical-systems logical-system-name system syslog], [edit system syslog] </pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure the logging of system messages to a remote destination.

**Options** *facility*—Class of messages to log. To specify multiple classes, include multiple *facility severity* statements. For a list of the facilities, see *Junos OS System Logging Facilities and Message Severity Levels*.

*hostname*—IPv4 address, IPv6 address, or fully qualified hostname of the remote machine to which to direct messages. To direct messages to multiple remote machines, include a **host** statement for each one.

**other-routing-engine**—Direct messages to the other Routing Engine on a router or switch with two Routing Engines installed and operational.



**NOTE:** The **other-routing-engine** option is not applicable to the QFX Series.

---

**port**—Port number of the remote syslog server that can be modified.

**scc-master**—(TX Matrix routers only) On a T640 router that is part of a routing matrix, direct messages to the TX Matrix router.

*severity*—Severity of the messages that belong to the facility specified by the paired *facility* name. Messages with severities of the specified level and higher are logged. For a list of the severities, see *Junos OS System Logging Facilities and Message Severity Levels*.

**sfc0-master**—(TX Matrix Plus routers only) On a T1600 or T4000 router that is part of a routing matrix, direct messages to the TX Matrix Plus router.

The remaining statements are explained separately.

<b>Required Privilege</b>	system—To view this statement in the configuration.
<b>Level</b>	system-control—To add this statement to the configuration.

- |                              |   |
|------------------------------|---|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Directing System Log Messages to a Remote Machine or the Other Routing Engine</i></li><li>• <i>Directing Messages to a Remote Destination from the Routing Matrix Based on the TX Matrix Router</i></li><li>• <i>Directing Messages to a Remote Destination from the Routing Matrix Based on a TX Matrix Plus Router</i></li><li>• <i>Junos OS System Log Messages Reference</i></li></ul> |
|------------------------------|---|

## interface (Accounting or Sampling)

<b>Syntax</b>	interface <i>interface-name</i> { engine-id <i>number</i> ; engine-type <i>number</i> ; source-address <i>address</i> ; }
<b>Hierarchy Level</b>	[edit forwarding-options accounting <i>group-name</i> output], [edit forwarding-options sampling family <i>family-name</i> output]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the output interface for sending copies of packets elsewhere to be analyzed.
<b>Options</b>	<b>engine-id <i>number</i></b> —Identity of the accounting interface.  <b>engine-type <i>number</i></b> —Type of this accounting interface.  <b><i>interface-name</i></b> —Name of the accounting interface.  <b>source-address <i>address</i></b> —Address used for generating packets.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Discard Accounting</i></li> <li>• <i>Configuring the Output File for Traffic Sampling</i></li> </ul>

## log-prefix (System)

<b>Syntax</b>	log-prefix <i>string</i> ;
<b>Hierarchy Level</b>	[edit system syslog host]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Include a text string in each message directed to a remote destination.
<b>Options</b>	<b><i>string</i></b> —Text string to include in each message.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Adding a Text String to System Log Messages Directed to a Remote Destination</i></li> <li>• <i>Junos OS System Log Messages Reference</i></li> </ul>

## match

---


<b>Syntax</b>	<code>match "regular-expression";</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> system syslog file <i>filename</i> ], [edit logical-systems <i>logical-system-name</i> system syslog user ( <i>username</i>   *)], [edit system syslog file <i>filename</i> ], [edit system syslog host <i>hostname</i>   other-routing-engine  scc-master)], [edit system syslog user ( <i>username</i>   *)]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify a text string that must (or must not) appear in a message for the message to be logged to a destination.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Using Regular Expressions to Refine the Set of Logged Messages</i></li></ul>

## size (System)

<b>Syntax</b>	<code>size size;</code>
<b>Hierarchy Level</b>	[edit system syslog archive], [edit system syslog file <i>filename</i> archive]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the maximum amount of data that the Junos OS logging utility writes to a log file <b>logfile</b> before archiving it (closing it, compressing it, and changing its name to <b>logfile.0.gz</b> ). The utility then opens and writes to a new file called <b>logfile</b> . For information about the number of archive files that the utility creates in this way, see <a href="#">files</a> .
<b>Options</b>	<p><b>size</b>—Maximum size of each system log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).</p> <p><b>Syntax:</b> <b>xk</b> to specify the number of kilobytes, <b>xm</b> for the number of megabytes, or <b>xg</b> for the number of gigabytes</p> <p><b>Range:</b> 64 KB through 1 GB</p> <p><b>Default:</b> 1 MB for MX Series routers and the QFX Series</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Specifying Log File Size, Number, and Archiving Properties</i></li> <li>• <i>Junos OS System Log Messages Reference</i></li> <li>• <a href="#">files on page 824</a></li> </ul>

## structured-data

---

<b>Syntax</b>	structured-data { brief; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> system syslog file <i>filename</i> ], [edit system syslog file <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Write system log messages to the log file in structured-data format, which complies with Internet draft draft-ietf-syslog-protocol-23, <i>The syslog Protocol</i> ( <a href="http://tools.ietf.org/html/draft-ietf-syslog-protocol-23">http://tools.ietf.org/html/draft-ietf-syslog-protocol-23</a> ).
<div> <b>NOTE:</b> When this statement is included, other statements that specify the format for messages written to the file are ignored (the <code>explicit-priority</code> statement at the [edit system syslog file <i>filename</i>] hierarchy level and the <code>time-format</code> statement at the [edit system syslog] hierarchy level).</div>	
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Logging Messages in Structured-Data Format</i></li><li>• <i>Junos OS System Log Messages Reference</i></li><li>• <i>explicit-priority</i></li><li>• <a href="#">time-format on page 833</a></li></ul>

## syslog (System)

```

Syntax  syslog {
        allow-duplicates;
        archive {
            (binary-data| no-binary-data);
            files number;
            size maximum-file-size;
            start-time "YYYY-MM-DD.hh:mm";
            transfer-interval minutes;
            (world-readable | no-world-readable);
        }
        console {
            facility severity;
        }
        file filename {
            facility severity;
            explicit-priority;
            match "regular-expression";
            archive {
                (binary-data| no-binary-data);
                files number;
                size maximum-file-size;
                start-time "YYYY-MM-DD.hh:mm";
                transfer-interval minutes;
                (world-readable | no-world-readable);
            }
            structured-data {
                brief;
            }
        }
        host (hostname | other-routing-engine | scc-master) {
            facility severity;
            explicit-priority;
            facility-override facility;
            log-prefix string;
            match "regular-expression";
            source-address source-address;
            structured-data {
                brief;
            }
            port port number;
        }
        log-rotate-frequency frequency;
        server server name;
        source-address source-address;
        time-format (millisecond | year | year millisecond);
        user (username | *) {
            facility severity;
            match "regular-expression";
        }
    }


```

Hierarchy Level [edit system]

<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the types of system log messages to send to files, to a remote destination, to user terminals, or to the system console.  The remaining statements are explained separately.
<b>Options</b>	<b>archive</b> —Define parameters for archiving log messages.  <b>console</b> —Send log messages of a specified class and severity to the console.  <b>file</b> —Send log messages to a named file.  <b>host</b> —Remote location to be notified of specific log messages.  <b>log-rotate-frequency</b> —Configure the interval for checking logfile size and archiving messages.  <b>server</b> —Name of the system log server in the inet.0 routing instance.  <b>source-address</b> —Include a specified address as the source address for log messages.  <b>time-format</b> —Additional information to include in the system log time stamp.  <b>user</b> —Notify a specific user of the log event.
<b>Required Privilege Level</b>	<b>system</b> —To view this statement in the configuration. <b>system-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Junos OS System Log Overview</i></li><li>• <i>Junos OS System Log Messages Reference</i></li><li>• <i>Overview of Single-Chassis System Logging Configuration</i></li></ul>



## time-format

<b>Syntax</b>	<code>time-format (year   millisecond   year millisecond);</code>
<b>Hierarchy Level</b>	<code>[edit system syslog]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Include the year, the millisecond, or both, in the timestamp on every standard-format system log message. The additional information is included for messages directed to each destination configured by a <b>file</b>, <b>console</b>, or <b>user</b> statement at the <code>[edit system syslog]</code> hierarchy level. As of Junos OS Release 11.4, the additional time information is also sent to destinations configured by a <b>host</b> statement.</p> <p>By default, the timestamp specifies the month, date, hour, minute, and second when the message was logged—for example, <b>Aug 21 12:36:30</b>. However, the timestamp for traceoption messages is specified in milliseconds by default, and is independent of the <code>[edit system syslog time-format]</code> statement.</p>
	<p> <b>NOTE:</b> When the <code>structured-data</code> statement is included at the <code>[edit system syslog file filename]</code> hierarchy level, this statement is ignored for the file.</p>
<b>Options</b>	<p><b>millisecond</b>—Include the millisecond in the timestamp.</p> <p><b>year</b>—Include the year in the timestamp.</p>
<b>Required Privilege Level</b>	<p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Including the Year or Millisecond in Timestamps</i></li> <li>• <i>Junos OS System Log Messages Reference</i></li> <li>• <a href="#">structured-data on page 830</a></li> </ul>

## time-zone

<b>Syntax</b>	<code>time-zone (GMT <i>hour-offset</i>   <i>time-zone</i>);</code>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. <b>GMT <i>hour-offset</i></b> option added in Junos OS Release 7.4.
<b>Description</b>	Set the local time zone. To have the time zone change take effect for all processes running on the router or switch, you must reboot the router or switch.
<b>Default</b>	UTC
<b>Options</b>	<p><b>GMT <i>hour-offset</i></b>—Set the time zone relative to UTC time.</p> <p><b>Range:</b> –14 through +12</p> <p><b>Default:</b> 0</p> <p><b><i>time-zone</i></b>—Specify the time zone as <b>UTC</b>, which is the default time zone, or as a string such as PDT (Pacific Daylight Time), or use one of the following continents and major cities:</p> <p>Africa/Abidjan, Africa/Accra, Africa/Addis_Ababa, Africa/Algiers, Africa/Asmera, Africa/Bamako, Africa/Bangui, Africa/Banjul, Africa/Bissau, Africa/Blantyre, Africa/Brazzaville, Africa/Bujumbura, Africa/Cairo, Africa/Casablanca, Africa/Ceuta, Africa/Conakry, Africa/Dakar, Africa/Dar_es_Salaam, Africa/Djibouti, Africa/Douala, Africa/El_Aaiun, Africa/Freetown, Africa/Gaborone, Africa/Harare, Africa/Johannesburg, Africa/Kampala, Africa/Khartoum, Africa/Kigali, Africa/Kinshasa, Africa/Lagos, Africa/Libreville, Africa/Lome, Africa/Luanda, Africa/Lubumbashi, Africa/Lusaka, Africa/Malabo, Africa/Maputo, Africa/Maseru, Africa/Mbabane, Africa/Mogadishu, Africa/Monrovia, Africa/Nairobi, Africa/Ndjamena, Africa/Niamey, Africa/Nouakchott, Africa/Ouagadougou, Africa/Porto-Novo, Africa/Sao_Tome, Africa/Timbuktu, Africa/Tripoli, Africa/Tunis, Africa/Windhoek</p> <p>America/Adak, America/Anchorage, America/Anguilla, America/Antigua, America/Aruba, America/Asuncion, America/Barbados, America/Belize, America/Bogota, America/Boise, America/Buenos_Aires, America/Caracas, America/Catamarca, America/Cayenne, America/Cayman, America/Chicago, America/Cordoba, America/Costa_Rica, America/Cuiaba, America/Curacao, America/Dawson, America/Dawson_Creek, America/Denver, America/Detroit, America/Dominica, America/Edmonton, America/El_Salvador, America/Ensenada, America/Fortaleza, America/Glace_Bay, America/Godthab, America/Goose_Bay, America/Grand_Turk, America/Grenada, America/Guadeloupe, America/Guatemala, America/Guayaquil, America/Guyana, America/Halifax, America/Havana, America/Indiana/Knox, America/Indiana/Marengo, America/Indiana/Vevay, America/Indianapolis, America/Inuvik, America/Iqaluit, America/Jamaica, America/Jujuy, America/Juneau, America/La_Paz, America/Lima, America/Los_Angeles, America/Louisville, America/Maceio, America/Managua, America/Manaus, America/Martinique, America/Mazatlan, America/Mendoza, America/Menominee, America/Mexico_City, America/Miquelon, America/Montevideo, America/Montreal, America/Montserrat, America/Nassau, America/New_York, America/Nipigon, America/Nome, America/Noronha, America/Panama, America/Pangnirtung, America/Paramaribo, America/Phoenix, America/Port-au-Prince,</p>

America/Port\_of\_Spain, America/Porto\_Acre, America/Puerto\_Rico, America/Rainy\_River, America/Rankin\_Inlet, America/Regina, America/Rosario, America/Santiago, America/Santo\_Domingo, America/Sao\_Paulo, America/Scoresbysund, America/Shiprock, America/St\_Johns, America/St\_Kitts, America/St\_Lucia, America/St\_Thomas, America/St\_Vincent, America/Swift\_Current, America/Tegucigalpa, America/Thule, America/Thunder\_Bay, America/Tijuana, America/Tortola, America/Vancouver, America/Whitehorse, America/Winnipeg, America/Yakutat, America/Yellowknife

Antarctica/Casey, Antarctica/DumontDURville, Antarctica/Mawson, Antarctica/McMurdo, Antarctica/Palmer, Antarctica/South\_Pole

Arctic/Longyearbyen

Asia/Aden, Asia/Alma-Ata, Asia/Amman, Asia/Anadyr, Asia/Aqttau, Asia/Aqtobe, Asia/Ashkhabad, Asia/Baghdad, Asia/Bahrain, Asia/Baku, Asia/Bangkok, Asia/Beirut, Asia/Bishkek, Asia/Brunei, Asia/Calcutta, Asia/Chungking, Asia/Colombo, Asia/Dacca, Asia/Damascus, Asia/Dubai, Asia/Dushanbe, Asia/Gaza, Asia/Harbin, Asia/Hong\_Kong, Asia/Irkutsk, Asia/Ishigaki, Asia/Jakarta, Asia/Jayapura, Asia/Jerusalem, Asia/Kabul, Asia/Kamchatka, Asia/Karachi, Asia/Kashgar, Asia/Katmandu, Asia/Krasnoyarsk, Asia/Kuala\_Lumpur, Asia/Kuching, Asia/Kuwait, Asia/Macao, Asia/Magadan, Asia/Manila, Asia/Muscat, Asia/Nicosia, Asia/Novosibirsk, Asia/Omsk, Asia/Phnom\_Penh, Asia/Pyongyang, Asia/Qatar, Asia/Rangoon, Asia/Riyadh, Asia/Saigon, Asia/Seoul, Asia/Shanghai, Asia/Singapore, Asia/Taipei, Asia/Tashkent, Asia/Tbilisi, Asia/Tehran, Asia/Thimbu, Asia/Tokyo, Asia/Ujung\_Pandang, Asia/Ulan\_Bator, Asia/Urumqi, Asia/Vientiane, Asia/Vladivostok, Asia/Yakutsk, Asia/Yekaterinburg, Asia/Yerevan

Atlantic/Azores, Atlantic/Bermuda, Atlantic/Canary, Atlantic/Cape\_Verde, Atlantic/Faeroe, Atlantic/Jan\_Mayen, Atlantic/Madeira, Atlantic/Reykjavik, Atlantic/South\_Georgia, Atlantic/St\_Helena, Atlantic/Stanley

Australia/Adelaide, Australia/Brisbane, Australia/Broken\_Hill, Australia/Darwin, Australia/Hobart, Australia/Lindeman, Australia/Lord\_Howe, Australia/Melbourne, Australia/Perth, Australia/Sydney

Europe/Amsterdam, Europe/Andorra, Europe/Athens, Europe/Belfast, Europe/Belgrade, Europe/Berlin, Europe/Bratislava, Europe/Brussels, Europe/Bucharest, Europe/Budapest, Europe/Chisinau, Europe/Copenhagen, Europe/Dublin, Europe/Gibraltar, Europe/Helsinki, Europe/Istanbul, Europe/Kaliningrad, Europe/Kiev, Europe/Lisbon, Europe/Ljubljana, Europe/London, Europe/Luxembourg, Europe/Madrid, Europe/Malta, Europe/Minsk, Europe/Monaco, Europe/Moscow, Europe/Oslo, Europe/Paris, Europe/Prague, Europe/Riga, Europe/Rome, Europe/Samara, Europe/San\_Marino, Europe/Sarajevo, Europe/Simferopol, Europe/Skopje, Europe/Sofia, Europe/Stockholm, Europe/Tallinn, Europe/Tirane, Europe/Vaduz, Europe/Vatican, Europe/Vienna, Europe/Vilnius, Europe/Warsaw, Europe/Zagreb, Europe/Zurich

Indian/Antananarivo, Indian/Chagos, Indian/Christmas, Indian/Cocos, Indian/Comoro, Indian/Kerguelen, Indian/Mahe, Indian/Maldives, Indian/Mauritius, Indian/Mayotte, Indian/Reunion

Pacific/Apia, Pacific/Auckland, Pacific/Chatham, Pacific/Easter, Pacific/Efate, Pacific/Enderbury, Pacific/Fakaofu, Pacific/Fiji, Pacific/Funafuti, Pacific/Galapagos, Pacific/Gambier, Pacific/Guadalcanal, Pacific/Guam, Pacific/Honolulu, Pacific/Johnston, Pacific/Kiritimati, Pacific/Kosrae, Pacific/Kwajalein, Pacific/Majuro, Pacific/Marquesas, Pacific/Midway, Pacific/Nauru, Pacific/Niue, Pacific/Norfolk, Pacific/Noumea, Pacific/Pago\_Pago, Pacific/Palau, Pacific/Pitcairn, Pacific/Ponape, Pacific/Port\_Moresby, Pacific/Rarotonga, Pacific/Saipan, Pacific/Tahiti, Pacific/Tarawa, Pacific/Tongatapu, Pacific/Truk, Pacific/Wake, Pacific/Wallis, Pacific/Yap

<b>Required Privilege</b>	system—To view this statement in the configuration.
<b>Level</b>	system-control—To add this statement to the configuration.

- Related Documentation**
- *Modifying the Default Time Zone for a Router or Switch Running Junos OS*
  - *System Management Configuration Statements*

---

## user (System Logging)

---

<b>Syntax</b>	<pre>user (username   *) {     facility severity;     match "regular-expression"; }</pre>
<b>Hierarchy Level</b>	[edit system syslog]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the logging of system messages to user terminals.
<b>Options</b>	<p><b>*</b> (the asterisk)—Log messages to the terminal sessions of all users who are currently logged in.</p> <p><b>facility</b>—Class of messages to log. To specify multiple classes, include multiple <b>facility severity</b> statements. For a list of the facilities, see <i>Junos OS System Logging Facilities and Message Severity Levels</i>.</p> <p><b>severity</b>—Severity of the messages that belong to the facility specified by the paired <b>facility</b> name. Messages with severities the specified level and higher are logged. For a list of the severities, see <i>Junos OS System Logging Facilities and Message Severity Levels</i>.</p> <p><b>username</b>—Junos OS login name of the user whose terminal session is to receive system log messages. To log messages to more than one user's terminal session, include more than one <b>user</b> statement.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Directing System Log Messages to a User Terminal</i></li><li>• <i>Junos OS System Logging Facilities and Message Severity Levels</i></li><li>• <i>Junos OS System Log Messages Reference</i></li></ul>

---

## world-readable (System)

---

<b>Syntax</b>	world-readable   no-world-readable;
<b>Hierarchy Level</b>	[edit system <a href="#">syslog</a> archive], [edit system <a href="#">syslog file filename</a> archive]
<b>Release Information</b>	Statement introduced before OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Grant all users permission to read log files, or restrict the permission only to the <b>root</b> user and users who have the Junos <b>maintenance</b> permission.
<b>Default</b>	no-world-readable
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Specifying Log File Size, Number, and Archiving Properties</i></li><li>• <i>Junos System Log Messages Reference</i></li></ul>



## CHAPTER 23

# Administration

- [Routine Monitoring on page 839](#)
- [Operational Commands on page 853](#)

## Routine Monitoring

---

- [Monitoring System Log Messages on page 839](#)
- [Checking Active Alarms with the J-Web Interface on page 842](#)
- [Monitoring Switch Control Traffic on page 843](#)
- [Monitoring System Properties on page 846](#)
- [Monitoring Chassis Information on page 847](#)
- [Monitoring System Process Information on page 850](#)
- [Managing Log, Temporary, and Crash Files on the Switch \(J-Web Procedure\) on page 850](#)

## Monitoring System Log Messages

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to filter and view system log messages for EX Series switches.

### Action

To view events in the J-Web interface, select **Monitor > Events and Alarms > View Events**.

Apply a filter or a combination of filters to view messages. You can use filters to display relevant events. [Table 96 on page 840](#) describes the different filters, their functions, and the associated actions.

To view events in the CLI, enter the following command:

```
show log
```

Table 96: Filtering System Log Messages

Field	Function	Your Action
System Log File	<p>Specifies the name of a system log file for which you want to display the recorded events.</p> <p>Lists the names of all the system log files that you configure.</p> <p>By default, a log file, <b>messages</b>, is included in the <b>/var/log/</b> directory.</p>	<p>To specify events recorded in a particular file, select the system log filename from the list—for example, <b>messages</b>.</p> <p>Select <b>Include archived files</b> to include archived files in the search.</p>
Process	<p>Specifies the name of the process generating the events you want to display.</p> <p>To view all the processes running on your system, enter the CLI command <b>show system processes</b>.</p> <p>For more information about processes, see the <a href="#">Junos OS Installation and Upgrade Guide</a>.</p>	<p>To specify events generated by a process, type the name of the process.</p> <p>For example, type <b>mgd</b> to list all messages generated by the management process.</p>
Date From To	<p>Specifies the time period in which the events you want displayed are generated.</p> <p>Displays a calendar that allows you to select the year, month, day, and time. It also allows you to select the local time.</p> <p>By default, the messages generated in the last hour are displayed. End Time shows the current time and Start Time shows the time one hour before End Time.</p>	<p>To specify the time period:</p> <ul style="list-style-type: none"> <li>Click the <b>Calendar</b> icon and select the year, month, and date—for example, <b>02/10/2007</b>.</li> <li>Click the <b>Calendar</b> icon and select the year, month, and date—for example, <b>02/10/2007</b>.</li> <li>Click to select the time in hours, minutes, and seconds.</li> </ul>
Event ID	<p>Specifies the event ID for which you want to display the messages.</p> <p>Allows you to type part of the ID and completes the remainder automatically.</p> <p>An event ID, also known as a system log message code, uniquely identifies a system log message. It begins with a prefix that indicates the generating software process or library.</p>	<p>To specify events with a specific ID, type the partial or complete ID—for example, <b>TFTPD_AF_ERR</b>.</p>
Description	<p>Specifies text from the description of events that you want to display.</p> <p>Allows you to use regular expressions to match text from the event description.</p> <p><b>NOTE:</b> Regular expression matching is case-sensitive.</p>	<p>To specify events with a specific description, type a text string from the description with regular expression.</p> <p>For example, type <b>^Initial*</b> to display all messages with lines beginning with the term <i>Initial</i>.</p>
Search	<p>Applies the specified filter and displays the matching messages.</p>	<p>To apply the filter and display messages, click <b>Search</b>.</p>
Reset	<p>Resets all the fields in the Events Filter box.</p>	<p>To reset the field values that are listed in the Events Filter box, click <b>Reset</b>.</p>



Table 96: Filtering System Log Messages (*continued*)

Field	Function	Your Action
Generate Raw Report  <b>NOTE:</b> <ul style="list-style-type: none"> <li>The Generate Raw Report button is enabled once the event log messages start loading in the Events Detail table.</li> <li>After the log messages are completely loaded in the Events Detail table, Generate Raw Report changes to Generate Report.</li> </ul>	Generates a list of event log messages in nontabular format  <b>NOTE:</b> Generate Raw Report is disabled until event log messages start loading in the Events Detail table.	To generate a raw report: <ol style="list-style-type: none"> <li>Click <b>Generate Raw Report</b>.  The <i>Opening filteredEvents.html</i> window appears.</li> <li>Select <b>Open with</b> to open the html file or select <b>Save File</b> to save the file.</li> <li>Click <b>OK</b>.</li> </ol>
Generate Report  <b>NOTE:</b> The Generate Report button appears only after event log messages are completely loaded in the Events Detail table. The Generate Raw Report button is displayed while event log messages are being loaded.	Generates a list of event log messages in tabular format, which shows system details, events filter criteria, and event details.	To generate a formatted report: <ol style="list-style-type: none"> <li>Click <b>Generate Report</b>.  The <i>Opening Report.html</i> window appears.</li> <li>Select <b>Open with</b> to open the html file or select <b>Save File</b> to save the file.</li> <li>Click <b>OK</b>.</li> </ol>

**Meaning** Table 97 on page 841 describes the Event Summary fields.



**NOTE:** By default, the View Events page in the J-Web interface displays the most recent 25 events, with severity levels highlighted in different colors. After you specify the filters, Event Summary displays the events matching the specified filters. Click the **First**, **Next**, **Prev**, and **Last** links to navigate through messages.

Table 97: Viewing System Log Messages

Field	Function	Additional Information
Process	Displays the name and ID of the process that generated the system log message.	The information displayed in this field is different for messages generated on the local Routing Engine than for messages generated on another Routing Engine (on a system with two Routing Engines installed and operational). Messages from the other Routing Engine also include the identifiers <b>re0</b> and <b>re1</b> to identify the Routing Engine.

Table 97: Viewing System Log Messages (*continued*)

Field	Function	Additional Information
Severity	<p>Severity level of a message is indicated by different colors.</p> <ul style="list-style-type: none"> <li>• <b>Unknown</b>—Gray—Indicates no severity level is specified.</li> <li>• <b>Debug/Info/Notice</b>—Green—Indicates conditions that are not errors but are of interest or might warrant special handling.</li> <li>• <b>Warning</b>—Yellow—Indicates conditions that warrant monitoring.</li> <li>• <b>Error</b>—Blue—Indicates standard error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels.</li> <li>• <b>Critical</b>—Pink—Indicates critical conditions, such as hard-drive errors.</li> <li>• <b>Alert</b>—Orange—Indicates conditions that require immediate correction, such as a corrupted system database.</li> <li>• <b>Emergency</b>—Red—Indicates system panic or other conditions that cause the switch to stop functioning.</li> </ul>	<p>A severity level indicates how seriously the triggering event affects switch functions. When you configure a location for logging a facility, you also specify a severity level for the facility. Only messages from the facility that are rated at that level or higher are logged to the specified file.</p>
Event ID	<p>Displays a code that uniquely identifies the message.</p> <p>The prefix on each code identifies the message source, and the rest of the code indicates the specific event or error.</p>	<p>The event ID begins with a prefix that indicates the generating software process.</p> <p>Some processes on a switch do not use codes. This field might be blank in a message generated from such a process.</p> <p>An event can belong to one of the following type categories:</p> <ul style="list-style-type: none"> <li>• <b>Error</b>—Indicates an error or failure condition that might require corrective action.</li> <li>• <b>Event</b>—Indicates a condition or occurrence that does not generally require corrective action.</li> </ul>
Event Description	Displays a more detailed explanation of the message.	
Time	Displays the time at which the message was logged.	

- Related Documentation**
- [Checking Active Alarms with the J-Web Interface on page 842](#)
  - [Understanding Alarm Types and Severity Levels on EX Series Switches on page 797](#)

## Checking Active Alarms with the J-Web Interface

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to view alarm information for the EX Series switches including alarm type, alarm severity, and a brief description for each active alarm on the switching platform.

**Action** To view the active alarms:

1. Select **Monitor > Events and Alarms > View Alarms** in the J-Web interface.
2. Select an alarm filter based on alarm type, severity, description, and date range.
3. Click **Go**.

All the alarms matching the filter are displayed.



**NOTE:** When the switch is reset, the active alarms are displayed.

**Meaning** Table 98 on page 843 lists the alarm output fields.

**Table 98: Summary of Key Alarm Output Fields**

Field	Values
Type	Category of the alarm: <ul style="list-style-type: none"> <li>• Chassis—Indicates an alarm condition on the chassis (typically an environmental alarm such as one related to temperature).</li> <li>• System—Indicates an alarm condition in the system.</li> </ul>
Severity	Alarm severity—either major (red) or minor (yellow).
Description	Brief synopsis of the alarm.
Time	Date and time when the failure was detected.

**Related Documentation**

- [Monitoring System Log Messages on page 839](#)
- [Dashboard for EX Series Switches on page 799](#)
- [Understanding Alarm Types and Severity Levels on EX Series Switches on page 797](#)

## Monitoring Switch Control Traffic

**Purpose**



**NOTE:** This topic applies only to the J-Web Application package.

Use the packet capture feature when you need to quickly capture and analyze switch control traffic on a switch. The packet capture feature allows you to capture traffic destined for or originating from the Routing Engine.

**Action** To use the packet capture feature in the J-Web interface, select **Troubleshoot > Packet Capture**.

To use the packet capture feature in the CLI, enter the following CLI command:

**monitor traffic**

**Meaning** You can use the packet capture feature to compose expressions with various matching criteria to specify the packets that you want to capture. You can decode and view the captured packets in the J-Web interface as they are captured. The packet capture feature does not capture transient traffic.

**Table 99: Packet Capture Field Summary**

Field	Function	Your Action
Interface	Specifies the interface on which the packets are captured. If you select default, packets on the Ethernet management port 0, are captured.	From the list, select an interface—for example, <b>ge-0/0/0</b> .
Detail level	Specifies the extent of details to be displayed for the packet headers. <ul style="list-style-type: none"> <li>Brief—Displays the minimum packet header information. This is the default.</li> <li>Detail—Displays packet header information in moderate detail.</li> <li>Extensive—Displays the maximum packet header information.</li> </ul>	From the list, select <b>Detail</b> .
Packets	Specifies the number of packets to be captured. Values range from 1 to <b>1000</b> . Default is <b>10</b> . Packet capture stops capturing packets after this number is reached.	From the list, select the number of packets to be captured—for example, <b>10</b> .
Addresses	Specifies the addresses to be matched for capturing the packets using a combination of the following parameters: <ul style="list-style-type: none"> <li>Direction—Matches the packet headers for IP address, hostname, or network address of the source, destination or both.</li> <li>Type—Specifies if packet headers are matched for host address or network address.</li> </ul> You can add multiple entries to refine the match criteria for addresses.	Select address-matching criteria. For example: <ol style="list-style-type: none"> <li>From the Direction list, select <b>source</b>.</li> <li>From the Type list, select <b>host</b>.</li> <li>In the Address box, type <b>10.1.40.48</b>.</li> <li>Click <b>Add</b>.</li> </ol>
Protocols	Matches the protocol for which packets are captured. You can choose to capture TCP, UDP, or ICMP packets or a combination of TCP, UDP, and ICMP packets.	From the list, select a protocol—for example, <b>tcp</b> .
Ports	Matches packet headers containing the specified source or destination TCP or UDP port number or port name.	Select a direction and a port. For example: <ul style="list-style-type: none"> <li>From the Type list, select <b>src</b>.</li> <li>In the Port box, type <b>23</b>.</li> </ul>
Advanced Options		

Table 99: Packet Capture Field Summary (*continued*)

Field	Function	Your Action
Absolute TCP Sequence	Specifies that absolute TCP sequence numbers are to be displayed for the packet headers.	To display absolute TCP sequence numbers in the packet headers, select this check box.
Layer 2 Headers	Specifies that link-layer packet headers are to be displayed.	To include link-layer packet headers while capturing packets, select this check box.
Non-Promiscuous	Specifies not to place the interface in promiscuous mode, so that the interface reads only packets addressed to it. In promiscuous mode, the interface reads every packet that reaches it.	To read all packets that reach the interface, select this check box.
Display Hex	Specifies that packet headers, except link-layer headers, are to be displayed in hexadecimal format.	To display the packet headers in hexadecimal format, select this check box.
Display ASCII and Hex	Specifies that packet headers are to be displayed in hexadecimal and ASCII format.	To display the packet headers in ASCII and hexadecimal formats, select this check box.
Header Expression	Specifies the match condition for the packets to be captured. The match conditions you specify for Addresses, Protocols, and Ports are displayed in expression format in this field.	You can enter match conditions directly in this field in expression format or modify the expression composed from the match conditions you specified for Addresses, Protocols, and Ports. If you change the match conditions specified for Addresses, Protocols, and Ports again, packet capture overwrites your changes with the new match conditions.
Packet Size	Specifies the number of bytes to be displayed for each packet. If a packet header exceeds this size, the display is truncated for the packet header. The default value is 96 bytes.	Type the number of bytes you want to capture for each packet header—for example, <b>256</b> .
Don't Resolve Addresses	Specifies that IP addresses are not to be resolved into hostnames in the packet headers displayed.	To prevent packet capture from resolving IP addresses to hostnames, select this check box.
No Timestamp	Suppresses the display of packet header timestamps.	To stop displaying timestamps in the captured packet headers, select this check box.
Write Packet Capture File	Writes the captured packets to a file in PCAP format in /var/tmp. The files are named with the prefix jweb-pcap and the extension .pcap. If you select this option, the decoded packet headers are not displayed on the packet capture page.	To decode and display the packet headers on the J-Web page, clear this check box.

**Related Documentation** • [Using the J-Web CLI Terminal on page 317](#)

## Monitoring System Properties

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to view system properties such as the name and IP address of the switch and resource usage.

### Action

To monitor system properties in the J-Web interface, select **Monitor > System View > System Information**.

To monitor system properties in the CLI, enter the following commands:

- **show system uptime**
- **show system users**
- **show system storage**

### Meaning

[Table 100 on page 846](#) summarizes key output fields in the system properties display.

**Table 100: Summary of Key System Properties Output Fields**

Field	Values	Additional Information
<b>General Information</b>		
Serial Number	Serial number for the switch.	
Junos OS Version	Version of Junos OS active on the switch.	
Hostname	The name of switch.	
IP Address	The IP address of the switch.	
Loopback Address	The loopback address.	
Domain Name Server	The address of the domain name server.	
Time Zone	The time zone on the switch.	
<b>Time</b>		
Current Time	Current system time, in Coordinated Universal Time (UTC).	

Table 100: Summary of Key System Properties Output Fields (*continued*)

Field	Values	Additional Information
System Booted Time	Date and time when the switch was last booted and how long it has been running.	
Protocol Started Time	Date and time when the switching protocols were last started and how long they have been running.	
Last Configured Time	Date and time when a configuration was last committed. This field also shows the name of the user who issued the last <b>commit</b> command, through either the J-Web interface or the CLI.	
Load Average	The CPU load average for 1, 5, and 15 minutes.	
<b>Storage Media</b>		
Internal Flash Memory	Memory usage details of internal flash.	
External Flash Memory	Usage details of external flash memory.	
<b>Logged in Users Details</b>		
User	Username of any user logged in to the switching platform.	
Terminal	Terminal through which the user is logged in.	
From	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.	
Login Time	Time when the user logged in.	This is the <b>LOGIN@</b> field in <b>show system users</b> command output.
Idle Time	How long the user has been idle.	

- Related Documentation**
- [Monitoring System Process Information on page 850](#)
  - [Understanding J-Web User Interface Sessions](#)

## Monitoring Chassis Information

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to view chassis properties such as general switch information, temperature and fan status, and resource information for the EX Series switch.

**Action** To view chassis properties in the J-Web interface, select **Monitor > System View > Chassis Information**. For an EX8200 Virtual Chassis configuration, select the Virtual Chassis member from the list.

To view chassis properties in the CLI, enter the following commands:

- **show chassis environment**
- **show chassis fpc**
- **show chassis hardware**

**Meaning** [Table 101 on page 848](#) gives information about the key output fields for chassis information.



**NOTE:** For an EX2200, EX2200-C, EX3200, or EX4500 switch or an EX4200, EX4300, or EX4550 standalone switch, the FPC slot number refers to the switch itself and is always 0. In a Virtual Chassis configuration, the FPC slot number refers to the member ID.

**Table 101: Summary of the Key Output Fields for Chassis Information**

Field	Values
Routing Engine Details	Select the <b>Master</b> tab to view details about the master Routing Engine or select <b>Backup</b> to view details about the backup Routing Engine.
Name/Value	<p>This table displays the following details of the master Routing Engine:</p> <ul style="list-style-type: none"> <li>• Routing Engine module</li> <li>• Model</li> <li>• Version</li> <li>• Part number</li> <li>• Serial number</li> <li>• Memory utilization</li> <li>• Temperature</li> <li>• Start time</li> <li>• CPU load average for 1, 5, and 15 minutes</li> </ul>
Power and Fan Tray Details	
Power	Select the <b>Power</b> tab to view details of the power supplies.
Name/Value	Displays the status and model number of each power supply.
Fan	Select the <b>Fan</b> tab to view details about the fans.



Table 101: Summary of the Key Output Fields for Chassis Information (*continued*)

Field	Values
Name/Value	Displays the status of each fan in the corresponding FPC.
Chassis Component Details	
Select component	Select an FPC to view general, temperature, resource, and subcomponent details.
General	Select the <b>General</b> tab to view the general information about the chassis components.
Name/Value	Displays general information: <ul style="list-style-type: none"> <li>• Version—Revision level. Supply the version number when reporting hardware problems to customer support.</li> <li>• Part number</li> <li>• Serial number—Supply the serial number when contacting customer support about the switch chassis.</li> <li>• Description—Brief text description.</li> </ul>
Temperature	Select the <b>Temperature</b> tab to view the temperature details of the components in the selected FPC.
Name/Value	Displays the temperature details of the sensors present in the selected FPC.
Resource	Select the <b>Resource</b> tab to view the resource details of the selected FPC.
Name/Value	Displays resource details: <ul style="list-style-type: none"> <li>• State:               <ul style="list-style-type: none"> <li>• <b>Dead</b>—Held in reset because of errors.</li> <li>• <b>Diag</b>—The FPC is running diagnostics.</li> <li>• <b>Dormant</b>—Held in reset.</li> <li>• <b>Empty</b>—No FPC is present.</li> <li>• <b>Online</b>—The FPC is online and running.</li> <li>• <b>Probed</b>—Probe is complete. The FPC is awaiting restart of the Packet Forwarding Engine.</li> <li>• <b>Probe-wait</b>—The FPC is waiting for the probe operation to start.</li> </ul> </li> <li>• <b>Total CPU DRAM</b>—Total DRAM, in megabytes, available to the FPC.</li> <li>• <b>Start time</b>—Date and time the switch was last rebooted.</li> </ul>

- Related Documentation**
- [Monitoring System Process Information on page 850](#)
  - [Monitoring System Properties on page 846](#)
  - [Dashboard for EX Series Switches on page 799](#)

## Monitoring System Process Information

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to view the processes running on the switch.

### Action

To view the software processes running on the switch in the J-Web interface, select **Monitor > System View > Process Details**.

To view the software processes running on the switch in the CLI, enter the following command.

```
show system processes
```

### Meaning

[Table 102 on page 850](#) summarizes the output fields in the system process information display.

The display includes the total CPU load and total memory utilization.

**Table 102: Summary of System Process Information Output Fields**

Field	Values	Additional Information
PID	Identifier of the process.	
Name	Owner of the process.	
State	Current state of the process.	
CPU Load	Percentage of the CPU that is being used by the process.	
Memory Utilization	Amount of memory that is being used by the process.	
Start Time	Time of day when the process started.	

### Related Documentation

- [Monitoring System Properties on page 846](#)
- For more information about show system properties command, see [show system uptime on page 221](#).

## Managing Log, Temporary, and Crash Files on the Switch (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You can use the J-Web interface to rotate log files and delete unnecessary log, temporary, and crash files on the switch.

1. [Cleaning Up Files on page 851](#)
2. [Downloading Files on page 851](#)
3. [Deleting Files on page 852](#)

---

### Cleaning Up Files

If you are running low on storage space, use the file cleanup procedure to quickly identify files to delete.

The file cleanup procedure performs the following tasks:

- Rotates log files—Archives the current log files, and creates fresh log files.
- Deletes log files in **/var/log**—Deletes files that are not currently being written to.
- Deletes temporary files in **/var/tmp**—Deletes files that have not been accessed within two days.
- Deletes all crash files in **/var/crash**—Deletes core files that the switch has written during an error.

To rotate log files and delete unnecessary files with the J-Web interface:

1. Select **Maintain > Files**.
2. In the Clean Up Files section, click **Clean Up Files**. The switching platform rotates log files and identifies files that can be safely deleted.

The J-Web interface displays the files that you can delete and the amount of space that will be freed on the file system.

3. Click one of the following options:
  - To delete the files and return to the Files page, click **OK**.
  - To cancel your entries and return to the list of files in the directory, click **Cancel**.

---

### Downloading Files

You can use the J-Web interface to download a copy of an individual log, temporary, or crash file from the switching platform. When you download a file, it is not deleted from the file system.

To download files with the J-Web interface:

1. In the J-Web interface, select **Maintain > Files**.
2. In the Download and Delete Files section, Click one of the following options:
  - Log Files—Log files in the **/var/log** directory on the switch.
  - Temporary Files—Lists the temporary files in the **/var/tmp** directory on the switching platform.

- Jailed Temporary Files (Install, Session, and so on)—Lists the files in the **/var/jail/tmp** directory on the switching platform.
- Crash (Core) Files—Lists the core files in the **/var/crash** directory on the switching platform.

The J-Web interface displays the files located in the directory.

3. Select the files that you want to download and click **Download**.
4. Choose a location for the saved file.

The file is saved as a text file, with a **.txt** file extension.

---

### Deleting Files

You can use the J-Web interface to delete an individual log, temporary, and crash file from the switching platform. When you delete the file, it is permanently removed from the file system.



**CAUTION:** If you are unsure whether to delete a file from the switching platform, we recommend using the Clean Up Files tool described in Cleaning Up Files. This tool determines which files can be safely deleted from the file system.

---

To delete files with the J-Web interface:

1. Select **Maintain > Files**.
2. In the Download and Delete Files section, Click one of the following options:
  - Log Files—Lists the log files in the **/var/log** directory on the switching platform.
  - Temporary Files—Lists the temporary files in the **/var/tmp** directory on the switching platform.
  - Jailed Temporary Files (Install, Session, etc)—Lists the files in the **/var/jail/tmp** directory on the switching platform.
  - Crash (Core) Files—Lists the core files in the **/var/crash** directory on the switching platform.

The J-Web interface displays the files in the directory.

3. Select the box next to each file you plan to delete.
4. Click **Delete**.

The J-Web interface displays the files you can delete and the amount of space that will be freed on the file system.

5. Click one of the following buttons on the confirmation page:
  - To delete the files and return to the Files page, click **OK**.
  - To cancel your entries and return to the list of files in the directory, click **Cancel**.

**Related Documentation** • *J-Web User Interface for EX Series Switches Overview*

## Operational Commands

---

- clear log
- file archive
- file checksum md5
- file checksum sha1
- file checksum sha-256
- file compare
- file copy
- file delete
- file list
- file rename
- file show
- monitor list
- monitor start
- monitor stop
- request chassis cb
- request chassis fabric plane
- request chassis fpc
- request system configuration rescue delete
- request system configuration rescue save
- request system scripts refresh-from commit
- request system scripts refresh-from event
- request system scripts refresh-from op
- show chassis alarms
- show chassis environment
- show chassis environment cb
- show chassis environment fpc
- show chassis environment routing-engine
- show chassis ethernet-switch
- show chassis fabric fpcs
- show chassis fabric map
- show chassis fabric plane
- show chassis fabric plane-location
- show chassis fabric summary

- `show chassis fpc`
- `show chassis led`
- `show chassis location`
- `show chassis pic`
- `show chassis routing-engine`
- `show log`
- `show pfe next-hop`
- `show pfe route`
- `show pfe terse`
- `show system alarms`
- `show system audit`
- `show system buffers`
- `show system connections`
- `show system core-dumps`
- `show system directory-usage`
- `show system processes`

## clear log

<b>Syntax</b>	<code>clear log <i>filename</i></code> <code>&lt;all&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Remove contents of a log file.
<b>Options</b>	<i>filename</i> —Name of the specific log file to delete.  <code>all</code> —(Optional) Delete the specified log file and all archived versions of it.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show log on page 1249</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear log on page 855</a>
<b>Output Fields</b>	See <a href="#">file list</a> for an explanation of output fields.

## Sample Output

### clear log

The following sample commands list log file information, clear the contents of a log file, and then display the updated log file information:

```

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel          26450 Jun 23 18:47 /var/log/sampled
total 1

user@host> clear log lcc0-re0:sampled
lcc0-re0:
-----

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel           57 Sep 15 03:44 /var/log/sampled
total 1

```

## file archive

---

<b>Syntax</b>	<code>file archive destination <i>destination</i> source <i>source</i> &lt;compress&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.
<b>Options</b>	<p><b>destination <i>destination</i></b>—Destination of the archived file or files. Specify the destination as a URL or filename. The Junos OS adds one of the following suffixes if the destination filename does not already have it:</p> <ul style="list-style-type: none"><li>• For archived files—The suffix <b>.tar</b></li><li>• For archived and compressed files—The suffix <b>.tgz</b></li></ul> <p><b>source <i>source</i></b>—Source of the original file or files. Specify the source as a URL or filename.</p> <p><b>compress</b>—(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix <b>.tgz</b>.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Format for Specifying Filenames and URLs in Junos OS CLI Commands</i></li></ul>
<b>List of Sample Output</b>	<a href="#">file archive (Multiple Files) on page 856</a> <a href="#">file archive (Single File) on page 856</a> <a href="#">file archive (with Compression) on page 857</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file archive (Multiple Files)

The following sample command archives all message files in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages* destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

### file archive (Single File)

The following sample command archives one message file in the local directory `/var/log/messages` as the single file `messages-archive.tar`.



```
user@host> file archive source /var/log/messages destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host
```

### file archive (with Compression)

The following sample command archives and compresses all message files in the local directory **/var/log/messages** as the single file **messages-archive.tgz**.

```
user@host> file archive compress source /var/log/messages* destination
/var/log/messages-archive.tgz
/usr/bin/tar: Removing leading / from absolute path names in the archive.
```

## file checksum md5

---

<b>Syntax</b>	<code>file checksum md5 &lt;pathname&gt; filename</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Calculate the Message Digest 5 (MD5) checksum of a file.
<b>Options</b>	<b>pathname</b> —(Optional) Path to a filename. <b>filename</b> —Name of a local file for which to calculate the MD5 checksum.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <a href="#">file checksum sha-256 on page 652</a></li><li>• <a href="#">file checksum sha1 on page 651</a></li><li>• <a href="#">op on page 90</a></li></ul>
<b>List of Sample Output</b>	<a href="#">file checksum md5 on page 858</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = 2a3b69e43f9bd4893729cc16f505a0f5
```

## file checksum sha1

<b>Syntax</b>	<code>file checksum sha1 &lt;pathname&gt; filename</code>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.
<b>Options</b>	<p><b>pathname</b>—(Optional) Path to a filename.</p> <p><b>filename</b>—Name of a local file for which to calculate the SHA-1 checksum.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <a href="#">file checksum md5 on page 650</a></li> <li>• <a href="#">file checksum sha-256 on page 652</a></li> <li>• <a href="#">op on page 90</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">file checksum sha1 on page 859</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax
```

```
SHA1 (/var/db/scripts/commitscript.slax) = ba9e47120c7ce55cff29afd73eacd370e162c676
```

## file checksum sha-256

---

<b>Syntax</b>	<code>file checksum sha-256 &lt;pathname&gt; filename</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.
<b>Options</b>	<b>pathname</b> —(Optional) Path to a filename.  <b>filename</b> —Name of a local file for which to calculate the SHA-256 checksum.
<b>Required Privilege Level</b>	maintenance view view-configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <a href="#">file checksum md5 on page 650</a></li><li>• <a href="#">file checksum sha1 on page 651</a></li><li>• <a href="#">op on page 90</a></li></ul>
<b>List of Sample Output</b>	<a href="#">file checksum sha-256 on page 860</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### file checksum sha-256

```
user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =
94c2b061fb55399e15babd2529453815601a602b5c98e5c12ed929c9d343dd71
```

## file compare

<b>Syntax</b>	<pre>file compare (files <i>filename filename</i>) &lt;context   unified&gt; &lt;ignore-white-space&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Compare two local files and describe the differences between them in default, context, or unified output styles:</p> <ul style="list-style-type: none"> <li>• <b>Default</b>—In the first line of output, <b>c</b> means lines were changed between the two files, <b>d</b> means lines were deleted between the two files, and <b>a</b> means lines were added between the two files. The numbers preceding this alphabetical marker represent the first file, and the lines after the alphabetical marker represent the second file. A left angle bracket (&lt;) in front of output lines refers to the first file. A right angle bracket (&gt;) in front of output lines refers to the second file.</li> <li>• <b>Context</b>—The display is divided into two parts. The first part is the first file; the second part is the second file. Output lines preceded by an exclamation point (!) have changed. Additions are marked with a plus sign (+), and deletions are marked with a minus sign (-).</li> <li>• <b>Unified</b>—The display is preceded by the line number from the first and the second file (xx,xxx,x). Before the line number, additions to the file are marked with a plus sign (+), and deletions to the file are marked with a minus sign (-). The body of the output contains the affected lines. Changes are viewed as additions plus deletions.</li> </ul>
<b>Options</b>	<p><b>files <i>filename</i></b>—Names of two local files to compare.</p> <p><b>context</b>—(Optional) Display output in context format.</p> <p><b>ignore-white-space</b>—(Optional) Ignore changes in the amount of white space.</p> <p><b>unified</b>—(Optional) Display output in unified format.</p>
<b>Required Privilege Level</b>	none
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Format for Specifying Filenames and URLs in Junos OS CLI Commands</i></li> <li>• <i>Viewing Core Files from Junos OS Processes</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">file compare files on page 862</a></p> <p><a href="#">file compare files context on page 862</a></p> <p><a href="#">file compare files unified on page 862</a></p> <p><a href="#">file compare files unified ignore-white-space on page 862</a></p>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file compare files

```
user@host> file compare files /tmp/one /tmp/two
100c100
<          full-name "File 1";
---
>          full-name "File 2";
102c102
<          class foo; # 'foo' is not defined
---
>          class super-user;
```

### file compare files context

```
user@host> file compare files /tmp/one /tmp/two context
*** /tmp/one   Wed Dec  3 17:12:50 2003
--- /tmp/two   Wed Dec  3 09:13:14 2003
*****
*** 97,104 ****
        }
    }
    user bill {
!         full-name "Bill Smith";
!         class foo; # 'foo' is not defined
        authentication {
            encrypted-password SECRET;
        }
--- 97,105 ----
    }
    user bill {
!         full-name "Bill Smith";
!         uid 1089;
!         class super-user;
        authentication {
            encrypted-password SECRET;
        }
    }
```

### file compare files unified

```
user@host> file compare files /tmp/one /tmp/two unified
--- /tmp/one   Wed Dec  3 17:12:50 2003
+++ /tmp/two   Wed Dec  3 09:13:14 2003
@@ -97,8 +97,9 @@
    }
}
user bill {
-     full-name "Bill Smith";
-     class foo; # 'foo' is not defined
+     full-name "Bill Smith";
+     uid 1089;
+     class super-user;
    authentication {
        encrypted-passwordSECRET;
    }
}
```

### file compare files unified ignore-white-space

```
user@host> file compare files /tmp/one /tmp/two unified ignore-white-space
```

```
--- /tmp/one    Wed Dec  3 09:13:10 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -99,7 +99,7 @@
     user bill {
         full-name "Bill Smith";
         uid 1089;
-        class foo; # 'foo' is not defined
+        class super-user;
         authentication {
             encrypted-password <SECRET>; # SECRET-DATA
         }
     }
```

## file copy

---

<b>Syntax</b>	<code>file copy <i>source destination</i></code> <code>&lt;source-address <i>address</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <b>source-address</b> option added in Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series switches.
<b>Description</b>	Copy files from one location to another location on the local device or to a location on a remote device reachable by the local device.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Format for Specifying Filenames and URLs in Junos OS CLI Commands</i></li><li>• <i>Default Directories for Junos OS File Storage on the Router or Switch</i></li><li>• <i>Copying a Configuration File from One Routing Engine to the Other</i></li></ul>
<b>List of Sample Output</b>	<a href="#">Copy a File from the Local Device to a Personal Computer on page 864</a> <a href="#">Copy a Configuration File between Routing Engines on page 864</a> <a href="#">Copy a Log File between Routing Engines on page 864</a> <a href="#">Copy a File from a TX Matrix Plus Router to a T1600 Router Connected to the TX Matrix Plus on page 865</a> <a href="#">Copy a File Using File Transfer Protocol on page 865</a> <a href="#">Copy a File Using File Transfer Protocol and Requiring a Password on page 865</a> <a href="#">Copy a File Using Secure Copy Protocol (scp) on page 865</a>

## Sample Output

The following are examples of a variety of file copy scenarios.

### Copy a File from the Local Device to a Personal Computer

```
user@host> file copy /var/tmp/rpd.core.4 mypc:/c/junipero/tmp
...transferring.file..... |           0 KB |    0.3 kB/s | ETA: 00:00:00 | 100%
```

### Copy a Configuration File between Routing Engines

The following sample command copies a configuration file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy /config/juniper.conf re1:/var/tmp/copied-juniper.conf
```

### Copy a Log File between Routing Engines

The following sample command copies a log file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy lcc0-re0:/var/log/chassisd lcc0-re1:/var/tmp
```



### Copy a File from a TX Matrix Plus Router to a T1600 Router Connected to the TX Matrix Plus

The following sample command copies a text file from Routing Engine 1 on the switch-fabric chassis sfc0 to Routing Engine 1 on the line-card chassis lcc0:

```
user@host> file copy sfc0-re1:/tmp/sample.txt lcc0-re1:/var/tmp
```

### Copy a File Using File Transfer Protocol

To use anonymous FTP to copy a local file to a remote system, enter the following command:

```
user@host> file copy filename ftp://hostname/filename
```

In the following example, `/config/juniper.conf` is the local file and `hostname` is the FTP server:

```
user@host> file copy /config/juniper.conf ftp://hostname/juniper.conf
Receiving ftp: //hostname/juniper.conf (2198 bytes): 100%
2198 bytes transferred in 0.0 seconds (2.69 MBps)
```

### Copy a File Using File Transfer Protocol and Requiring a Password

To use FTP where you require more privacy and are prompted for a password, enter the following command:

```
root@host> file copy filename ftp://user@hostname/filename
```

In the following example, `/config/juniper.conf` is the local file and `hostname` is the FTP server:

```
root@host> file copy /config/juniper.conf ftp://user@hostname/juniper.conf
Password for user@hostname: *****
Receiving ftp: //user@hostname/juniper.conf (2198 bytes): 100%
2198 bytes transferred in 0.0 seconds (2.69 MBps)
```

### Copy a File Using Secure Copy Protocol (scp)

To use scp to copy a local file to a remote system, enter the following command:

```
root@host> file copy filename scp://user@hostname/path/filename
```

In the following example, `/config/juniper.conf` is the local file, `user` is the username, and `ssh-host` is the scp server:

```
root@host> file copy /config/juniper.conf scp://user@ssh-host/tmp/juniper.conf
user@ssh-host's password: *****
juniper.conf          100%
|*****|
2198          00:00
```

## file delete

---

<b>Syntax</b>	<code>file delete <i>filename</i></code> <code>&lt;purge&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Delete a file on the local router or switch.
<b>Options</b>	<b><i>filename</i></b> —Name of the file to delete. For a routing matrix, include chassis information in the filename if the file to be deleted is not local to the Routing Engine from which the command is issued.  <b><i>purge</i></b> —(Optional) Overwrite regular files before deleting them.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file delete on page 866</a> <a href="#">file delete (Routing Matrix) on page 866</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file delete

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

### file delete (Routing Matrix)

```
user@host> file list lcc0-re0:/var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete lcc0-re0:/var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

## file list

<b>Syntax</b>	file list <detail   recursive> <filename>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display a list of files on the local router or switch.
<b>Options</b>	<p><b>none</b>—Display a list of all files for the current directory.</p> <p><b>detail   recursive</b>—(Optional) Display detailed output or descend recursively through the directory hierarchy, respectively.</p> <p><b>filename</b>—(Optional) Display a list of files. For a routing matrix, the filename must include the chassis information.</p>
<b>Additional Information</b>	The default directory is the home directory of the user logged in to the router or switch. To view available directories, enter a space and then a backslash (/) after the <b>file list</b> command. To view files within a specific directory, include a backslash followed by the directory and, optionally, subdirectory name after the <b>file list</b> command.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file list on page 867</a> <a href="#">file list (Routing Matrix) on page 867</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file list

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
```

### file list (Routing Matrix)

```
user@host> file list lcc0-re0:var/tmp
lcc0-re0:
-----
/var/tmp/:
.gdbinit
.pccardd
Test/
chassisd*
chassisd.nathan*
check_time*
```

```
cores/  
diagTestPrep*  
diagtest*  
diagtest.regress*  
do_switchovers*  
dump_test*  
err.manoj.log  
esw_clearstats*  
esw_counter*  
esw_debug*  
esw_debug_ge*  
esw_filt_test*  
esw_filter_tnp_addr*  
esw_getstats*  
esw_phy*  
esw_stats*
```

## file rename

<b>Syntax</b>	<code>file rename <i>source destination</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Rename a file on the local router or switch.
<b>Options</b>	<i>destination</i> —New name for the file.  <i>source</i> —Original name of the file. For a routing matrix, the filename must include the chassis information.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file rename on page 869</a> <a href="#">file rename (Routing Matrix) on page 869</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### file rename

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

#### file rename (Routing Matrix)

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
-----

/var/tmp:
.pccardd
sartre.conf
snmpd
syslogd.core-tarball.0.tgz
```

```
user@host> file rename lcc0-re0:/var/tmp/snmpd /var/tmp/snmpd.rr
```

```
user@host> file list lcc0-re1:/var/tmp
```

```
lcc0-re1:
```

```
-----
```

```
/var/tmp:
```

```
.pccardd
```

```
sartre.conf
```

```
snmpd.rr
```

```
syslogd.core-tarball.0.tgz
```

## file show

<b>Syntax</b>	<code>file show <i>filename</i></code> <code>&lt;encoding (base64   raw)&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the contents of a file.
<b>Options</b>	<b><i>filename</i></b> —Name of a file. For a routing matrix, the filename must include the chassis information.  <b><code>encoding (base64   raw)</code></b> —(Optional) Encode file contents with base64 encoding or show raw text.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file show on page 871</a> <a href="#">file show (Routing Matrix) on page 871</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file show

```
user@host> file show /var/log/messages
Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
...
```

### file show (Routing Matrix)

```
user@host> file show lcc0-re0:/var/tmp/gdbinit
lcc0-re0:
-----
#####
# Settings
#####

set print pretty

#####
# Basic stuff
#####

define msgbuf
    printf "%s", msgbufp->msg_ptr
end
```

```
# hex dump of a block of memory
# usage: dump address length
define dump
  p $arg0, $arg1
  set $ch = $arg0
  set $j = 0
  set $n = $arg1
  while ($j < $n)
    #printf "%x %x ",&$ch[$j],$ch[$j]
    printf "%x ",$ch[$j]
    set $j = $j + 1
    if (!($j % 16))
      printf "\n"
    end
  end
end
end
```



## monitor list

<b>Syntax</b>	monitor list
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Display the status of monitored log and trace files.
<b>Options</b>	This command has no options.
<b>Additional Information</b>	Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are configured with the <b>syslog</b> statement at the <b>[edit system]</b> hierarchy level and the <b>options</b> statement at the <b>[edit routing-options]</b> hierarchy level. The trace files generated by the routing protocol process are those configured with <b>traceoptions</b> statements at the <b>[edit routing-options]</b> , <b>[edit interfaces]</b> , and <b>[edit protocols protocol]</b> hierarchy levels.
<b>Required Privilege Level</b>	trace
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">monitor start on page 874</a></li> <li>• <a href="#">monitor stop on page 876</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">monitor list on page 873</a>
<b>Output Fields</b>	<a href="#">Table 103 on page 873</a> describes the output fields for the <b>monitor list</b> command. Output fields are listed in the approximate order in which they appear.

Table 103: monitor list Output Fields

Field Name	Field Description
<b>monitor start</b>	Indicates the file is being monitored.
<b>"filename"</b>	Name of the file that is being monitored.
<b>Last changed</b>	Date and time at which the file was last modified.

## Sample Output

### monitor list

```
user@host> monitor list
monitor start "vrrpd" (Last changed Dec 03:11:06 20)
monitor start "cli-commands" (Last changed Nov 07:3)
```

## monitor start

<b>Syntax</b>	<code>monitor start <i>filename</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Start displaying the system log or trace file and additional entries being added to those files.
<b>Options</b>	<i>filename</i> —Specific log or trace file.
<b>Additional Information</b>	Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are configured with the <b>syslog</b> statement at the <b>[edit system]</b> hierarchy level and the <b>options</b> statement at the <b>[edit routing-options]</b> hierarchy level. The trace files generated by the routing protocol process are configured with <b>traceoptions</b> statements at the <b>[edit routing-options]</b> , <b>[edit interfaces]</b> , and <b>[edit protocols protocol]</b> hierarchy levels.



**NOTE:** To monitor a log file within a logical system, issue the **monitor start *logical-system-name/filename*** command.

<b>Required Privilege Level</b>	trace
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">monitor list on page 873</a></li> <li><a href="#">monitor stop on page 876</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">monitor start on page 875</a>
<b>Output Fields</b>	<a href="#">Table 104 on page 874</a> describes the output fields for the <b>monitor start</b> command. Output fields are listed in the approximate order in which they appear.

**Table 104: monitor start Output Fields**

Field Name	Field Description
<b>***<i>filename</i>***</b>	Name of the file from which entries are being displayed. This line is displayed initially and when the command switches between log files.
<b><i>Date and time</i></b>	Timestamp for the log entry.

## Sample Output

### monitor start

```
user@host> monitor start system-log
*** system-log***
Jul 20 15:07:34 hang sshd[5845]: log: Generating 768 bit RSA key.
Jul 20 15:07:35 hang sshd[5845]: log: RSA key generation complete.
Jul 20 15:07:35 hang sshd[5845]: log: Connection from 204.69.248.180 port 912
Jul 20 15:07:37 hang sshd[5845]: log: RSA authentication for root accepted.
Jul 20 15:07:37 hang sshd[5845]: log: ROOT LOGIN as 'root' from trip.jcmax.com
Jul 20 15:07:37 hang sshd[5845]: log: Closing connection to 204.69.248.180
```

## monitor stop

---

<b>Syntax</b>	<code>monitor stop <i>filename</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Stop displaying the system log or trace file.
<b>Options</b>	<i>filename</i> —Specific log or trace file.
<b>Additional Information</b>	Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are those configured with the <b>syslog</b> statement at the <b>[edit system]</b> hierarchy level and the <b>options</b> statement at the <b>[edit routing-options]</b> hierarchy level. The trace files generated by the routing protocol process are those configured with <b>traceoptions</b> statements at the <b>[edit routing-options]</b> , <b>[edit interfaces]</b> , and <b>[edit protocols <i>protocol</i>]</b> hierarchy levels.
<b>Required Privilege Level</b>	trace
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">monitor list on page 873</a></li><li>• <a href="#">monitor start on page 874</a></li></ul>
<b>List of Sample Output</b>	<a href="#">monitor stop on page 876</a>
<b>Output Fields</b>	This command produces no output.

## Sample Output

### monitor stop

```
user@host> monitor stop
```

## request chassis cb

<b>List of Syntax</b>	<a href="#">Syntax on page 877</a> <a href="#">Syntax (TX Matrix Router) on page 877</a> <a href="#">Syntax (TX Matrix Plus Router) on page 877</a> <a href="#">Syntax (QFabric System) on page 877</a>
<b>Syntax</b>	<code>request chassis cb (offline   online) slot <i>slot-number</i></code>
<b>Syntax (TX Matrix Router)</b>	<code>request chassis cb (offline   online) &lt;slot <i>slot-number</i>   lcc <i>number</i> slot <i>cb-slot-number</i>   scc <i>number</i> slot <i>cb-slot-number</i>&gt;</code>
<b>Syntax (TX Matrix Plus Router)</b>	<code>request chassis cb (offline   online) &lt;slot <i>slot-number</i>   lcc <i>number</i> slot <i>cb-slot-number</i>   sfc <i>number</i> slot <i>cb-slot-number</i>&gt;</code>
<b>Syntax (QFabric System)</b>	<code>request chassis cb (offline   online) interconnect-device <i>name</i> slot <i>slot-number</i> &lt;interconnect-device <i>name</i> slot <i>slot-number</i> (offline   online)&gt;</code>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS 9.4 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS 11.3 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
<b>Description</b>	(M120, M320, and MX Series routers and T Series routers, QFabric systems, and EX8200 switches only) Control the operation of the Control Board (CB). For information about the meaning of “CBs” on the switches, see <a href="#">“EX Series Switches Hardware and CLI Terminology Mapping” on page 315</a> .
<b>Options</b>	<b>offline</b> —Take the Control Board offline.



**NOTE:** On a QFabric system, to bring the backup Control Board on a QFX3008-I Interconnect device offline, issue the `request chassis cb slot backup-slot-number offline` command.



**NOTE:** Only backup Control Board can be turned offline or online. To turn a Control Board offline or to bring it back online, the Routing Engine should be turned offline first.

**online**—Bring the Control Board online.

**interconnect-device *name***—(QFabric systems only) (Optional) Bring the QFX3008-I Interconnect device Control Board either offline or online:

**slot slot-number**—Control Board slot number:

- (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, if you specify the number of the T640 router by using the **lcc number** option (the recommended method), replace **cb-slot-number** with a value from 0 through 1.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 or T4000 router by using the **lcc number** option (the recommended method), replace **cb-slot-number** with a value from 0 through 1.

- M320 router—Replace **slot-number** with a value from 0 through 1.
- MX480/MX240 routers—Replace **slot-number** with a value from 0 through 1.
- MX960 router—Replace **slot-number** with a value from 0 through 2.
- MX2020 and MX2010 routers—Replace **slot-number** with 0 or 1.
- EX8208 switch—Replace **slot-number** with a value from 0 through 2.
- EX8216 switch—Replace **slot-number** with a value from 0 through 1.
- QFabric System—Replace **slot-number** with a value from 0 through 1.

**lcc number**—(TX Matrix, TX Matrix Plus routers only) (Optional) Line-card chassis number.

Replace **number** with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**sfc number**—(TX Matrix Plus routers only) (Optional) Change the CB status for the TX Matrix Plus router (switch-fabric chassis). Replace **number** with 0.

**Required Privilege Level** maintenance

**Related Documentation**

- [show chassis environment cb on page 970](#)
- *Understanding Switching Control Board Redundancy*
- *Routing Engine and Switching Control Board Redundancy Configuration Statements*

**List of Sample Output**

- [request chassis cb on page 879](#)
- [request chassis cb interconnect-device \(QFabric System\) on page 879](#)
- [request chassis cb \(MX2020 Router\) on page 879](#)
- [request chassis cb \(MX2010 Router\) on page 879](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request chassis cb

```
user@host> request chassis cb offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

### request chassis cb interconnect-device (QFabric System)

```
user@switch> request chassis cb interconnect-device interconnect1 offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

### request chassis cb (MX2020 Router)

```
user@host> request chassis cb offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

### request chassis cb (MX2010 Router)

```
user@host> request chassis cb offline slot 1
Backup CB 1 cannot be set offline, backup RE is online
```

## request chassis fabric plane

---

<b>Syntax</b>	<code>request chassis fabric plane <i>plane-number</i> (offline   online)</code>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
<b>Description</b>	<p>(M120 and MX Series routers and EX8200 switches only) Control the operation of the specified fabric plane.</p> <p>On an MX480 or MX240 series router, you can configure the active control board for redundancy mode or increased bandwidth mode. When running in increased bandwidth mode, MX series routers with Trio chips and the MPC3E will use all eight active fabric planes.</p> <p>To take both plane 0 and plane 1 offline on a MX480 and MX240 series routers with one or more MPC4E MICs installed, a X86 Media Service Blade, and/or 100G PFE, and where redundancy-mode is configured for "increased-bandwidth", Juniper recommends taking plane 1 offline before plane 0. Likewise, when the router is configured for increased-bandwidth mode, taking fabric planes 0, 2, 4, and 6 offline can cause the chassis to run in a reduced fabric bandwidth mode. Plane 7 may remain in a "spare" state (as seen in the "show chassis fabric summary" command output) until plane 3 is taken offline and then brought back up.</p>
<b>Options</b>	<p><b>offline</b>—Take the fabric plane offline. Use the <code>request chassis fabric plane <i>plane-number</i> offline</code> command to clear a <b>FAULT</b> state on a fabric plane. To bring the fabric plane back online, use the <code>request chassis fabric plane <i>plane-number</i> online</code> command.</p> <p><b>online</b>—Bring the fabric plane online.</p> <p><b>plane <i>plane-number</i></b>—Fabric plane number.</p> <ul style="list-style-type: none"><li>• For the M120 router, replace <i>plane-number</i> with a value from 0 through 3.</li><li>• For the MX480 and MX240 routers, replace <i>plane-number</i> with a value from 0 through 7.</li><li>• For the MX2020 and MX2010 routers, replace <i>plane-number</i> with a value from 0 through 7.</li><li>• For the MX960 router, replace <i>plane-number</i> with a value from 0 through 5.</li><li>• For the EX8208 switch, replace <i>plane-number</i> with a value from 0 through 11.</li><li>• For the EX8216 switch, replace <i>plane-number</i> with a value from 0 through 7.</li></ul>
<b>Required Privilege Level</b>	maintenance



Related Documentation	<ul style="list-style-type: none"> <li>• <a href="#">show chassis fabric plane on page 1108</a></li> <li>• <a href="#">show chassis fabric plane-location on page 1150</a></li> <li>• <a href="#">show chassis fabric summary on page 1155</a></li> <li>• <i>Fabric Management Overview</i></li> </ul>
List of Sample Output	<a href="#">request chassis fabric plane 0 online on page 881</a> <a href="#">request chassis fabric plane 0 offline on page 881</a> <a href="#">request chassis fabric plane 0 online (EX8200 switch) on page 881</a> <a href="#">request chassis fabric plane (MX2020 Router) on page 881</a> <a href="#">request chassis fabric plane (MX2010 Router) on page 881</a>
Output Fields	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### [request chassis fabric plane 0 online](#)

```
user@host> request chassis fabric plane 0 online
Online initiated, use "show chassis fabric plane" to verify
```

### [request chassis fabric plane 0 offline](#)

```
user@host> request chassis fabric plane 0 offline
Offline initiated, use "show chassis fabric plane" to verify
```

### [request chassis fabric plane 0 online \(EX8200 switch\)](#)

```
user@host> request chassis fabric plane 0 online

Plane 0 is already active
```

### [request chassis fabric plane \(MX2020 Router\)](#)

```
user@host> request chassis fabric plane 2 online
Plane 2 is already active
```

### [request chassis fabric plane \(MX2010 Router\)](#)

```
user@host> request chassis fabric plane 4 online
Plane 4 is already active
```

## request chassis fpc

---

<b>List of Syntax</b>	<a href="#">Syntax on page 882</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 882</a> <a href="#">Syntax (MX Series Routers) on page 882</a> <a href="#">Syntax (MX2020 3D Universal Edge Routers) on page 882</a> <a href="#">Syntax (MX2010 3D Universal Edge Routers) on page 882</a> <a href="#">Syntax (QFabric System) on page 882</a> <a href="#">Syntax (PTX Series Packet Transport Routers) on page 882</a>
<b>Syntax</b>	<code>request chassis fpc (offline   online   restart) slot <i>slot-number</i></code>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	<code>request chassis fpc (offline   online   restart) slot <i>slot-number</i> &lt;fcc <i>number</i>&gt;</code>
<b>Syntax (MX Series Routers)</b>	<code>request chassis fpc (offline   online   restart) slot <i>slot-number</i> &lt;all-members&gt;</code> <code>&lt;local&gt;</code> <code>&lt;member <i>member-id</i>&gt;</code>
<b>Syntax (MX2020 3D Universal Edge Routers)</b>	<code>request chassis fpc (offline   online   restart) slot <i>slot-number</i></code>
<b>Syntax (MX2010 3D Universal Edge Routers)</b>	<code>request chassis fpc (offline   online   restart) slot <i>slot-number</i></code>
<b>Syntax (QFabric System)</b>	<code>request chassis fpc</code> <code>&lt;interconnect-device <i>name</i> slot <i>slot-number</i> (offline   online)&gt;</code> <code>&lt;(offline   online) interconnect-device <i>name</i> slot <i>slot-number</i>&gt;</code> <code>&lt;slot <i>slot-number</i> interconnect-device <i>name</i> (offline   online)&gt;</code>
<b>Syntax (PTX Series Packet Transport Routers)</b>	<code>request chassis fpc (offline   online   restart) slot <i>slot-number</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS 11.3 for QFX Series. Command introduced in Junos OS 12.1x48 for PTX Series Packet Transport Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
<b>Description</b>	(M20, M40, M40e, M120, M160, M320, MX Series, and T Series routers, QFabric systems, EX Series switches, and PTX Series Packet Transport Routers only) Control the operation of the Flexible PIC Concentrator (FPC). For information about the meaning of “FPCs” on the switches, see <a href="#">“EX Series Switches Hardware and CLI Terminology Mapping” on page 315</a> .



**NOTE:** Beginning in Junos OS 12.3, it is possible that FPCs brought offline using the `request chassis fpc slot fpc-slot offline` operational-mode CLI command can come online during a configuration commit or power-supply replacement procedure. As an alternative, use the `set fpc fpc-slot power off` configuration-mode command at the `[edit chassis]` hierarchy level to ensure that the FPCs remain offline.

**Options**    **offline**—Take the FPC offline.

**online**—Bring the FPC online.

**interconnect-device *name***—(QFabric systems only) Bring the Flexible Port Concentrator (FPC) on the QFX3008-I Interconnect device either offline or online:

- (QFabric System) On a QFabric system, specify the name of the QFX3008-I Interconnect device containing the Flexible Port Concentrator (FPC) you want to bring either offline or online.

**restart**—Restart the FPC.

**slot *slot-number***—FPC slot number:

- M20 router—0 through 3.
- M120 router—0 through 5.
- MX240 router—0 through 2. On the MX240 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX480 router—0 through 5. On the MX480 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX960 router—0 through 11. On the MX960 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX2020 router—0 through 19.
- MX2010 router—0 through 9.
- TX Matrix and TX Matrix Plus routers only—On the TX Matrix router, if you specify the number of the T640 router by using the ***lcc number*** option (the recommended method), replace ***slot-number*** with a value from 0 through 7. Otherwise, replace ***slot-number*** with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 or T4000 router by using the ***lcc number*** option (the recommended method), replace ***slot-number*** with a value from 0 through 7. Otherwise, replace ***slot-number*** with a value from 0 through 31. In case of TX Matrix Plus router with 3D SIBs, replace

*slot-number* with a value from 0 through 63. For example, the following commands have the same result:

```
user@host> request chassis fpc lcc 1 slot 1 offline
user@host> request chassis fpc slot 9 offline
```

- Other routers—0 through 7.
- QFabric System—Replace *slot-number* with a value from 0 through 2.
- EX Series switches:
  - EX4200 switches in a Virtual Chassis configuration—Replace *slot-number* with a value from 0 through 9.
  - EX6210 switches—Replace *slot-number* with a value from 0 through 9.



**NOTE:** These commands are not supported for slots 4 and 5 when a Switch Fabric and Routing Engine (SRE) module is installed in those slots. These commands are supported for slots 4 and 5 only if a line card is installed in them.

---

- EX8208 switches—Replace *slot-number* with a value from 0 through 7.
- EX8216 switches—Replace *slot-number* with a value from 0 through 15.
- PTX5000 Packet Transport Router—Replace *slot-number* with a value from 0 through 7.

**all-members**—(MX Series routers only) (Optional) Change FPC status of all members of the Virtual Chassis configuration.

**local**—(MX Series routers only) (Optional) Change FPC status of the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Change FPC status of the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**Required Privilege Level** maintenance

**Related Documentation**

- [show chassis fpc on page 1160](#)
- *show chassis fpc-feb-connectivity*
- [show chassis fabric fpcs on page 1063](#)
- *Configuring the Junos OS to Make a Flexible PIC Concentrator Stay Offline*
- *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
- *MX960 Flexible PIC Concentrator Description*

**List of Sample Output** [request chassis fpc on page 885](#)  
[request chassis fpc \(MX Series Routers with Media Services Blade \[MSB\]\) on page 885](#)  
[request chassis fpc \(MX2020 Router\) on page 885](#)  
[request chassis fpc \(MX2010 Router\) on page 885](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### [request chassis fpc](#)

```
user@host> request chassis fpc online slot 0
FPC 0 already online
```

### [request chassis fpc \(MX Series Routers with Media Services Blade \[MSB\]\)](#)

```
user@host> request chassis fpc slot 0
Possible completions:
offline           Take FPC offline
online            Bring FPC online
restart           Restart FPC
```

### [request chassis fpc \(MX2020 Router\)](#)

```
user@host >request chassis fpc online slot 2
FPC 2 already online
```

### [request chassis fpc \(MX2010 Router\)](#)

```
user@host >request chassis fpc offline slot 5
Offline initiated, use "show chassis fpc" to verify
```

## request system configuration rescue delete

---

**Syntax** request system configuration rescue delete

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Delete an existing rescue configuration.



**NOTE:** The [edit system configuration] hierarchy is not available on QFabric systems.

**Options** This command has no options.

**Required Privilege Level** maintenance

**Related Documentation**

- [request system configuration rescue save on page 666](#)
- [request system software rollback on page 756](#)
- [show system commit on page 675](#)

**List of Sample Output** [request system configuration rescue delete on page 886](#)


**Output Fields** This command produces no output.

### Sample Output

request system configuration rescue delete

```
user@host> request system configuration rescue delete
```

## request system configuration rescue save

<b>Syntax</b>	request system configuration rescue save
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Save the most recently committed configuration as the rescue configuration so that you can return to it at any time by using the <b>rollback</b> command.
<div>  <b>NOTE:</b> The [edit system configuration] hierarchy is not available on QFabric systems. </div>	
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request system software delete on page 752</a></li> <li>• <a href="#">request system software rollback on page 756</a></li> <li>• <a href="#">show system commit on page 675</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system configuration rescue save on page 887</a>
<b>Output Fields</b>	This command produces no output.

### Sample Output

#### request system configuration rescue save

```
user@host> request system configuration rescue save
```

## request system scripts refresh-from commit

---

<b>Syntax</b>	<code>request system scripts refresh-from commit file <i>file-name</i> url <i>url-path</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.1 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Automatically download the initial Junos OS configuration and a set of standard commit scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.</p> <p>The Junos XML management protocol equivalent for this operational mode command is:</p> <pre>&lt;request-script-refresh-from&gt;   &lt;type&gt;commit&lt;/type&gt;   &lt;file&gt;file-name&lt;/file&gt;   &lt;URL&gt;URL&lt;/URL&gt; &lt;/request-script-refresh-from&gt;</pre>
<b>Options</b>	<p><b>file <i>file-name</i></b>—Name of the file to be downloaded.</p> <p><b>url <i>url-path</i></b>—URL of the file to be downloaded.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Automatic Refreshing of Scripts on EX Series Switches on page 613</a></li><li>• <a href="#">Junos OS Junos XML Management Protocol Guide</a></li><li>• <a href="#">Junos OS NETCONF XML Management Protocol Guide</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request system scripts refresh-from commit file config.txt url http://host1.juniper.net on page 888</a>

### Sample Output

`request system scripts refresh-from commit file config.txt url http://host1.juniper.net`

```
user@switch> request system scripts refresh-from commit file config.txt url
http://host1.juniper.net
user@switch>
```



## request system scripts refresh-from event

<b>Syntax</b>	<code>request system scripts refresh-from event file <i>file-name</i> url <i>url-path</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.1 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Automatically download the initial Junos OS configuration and a set of standard event scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.</p> <p>The Junos XML management protocol equivalent for this operational mode command is:</p> <pre>&lt;request-script-refresh-from&gt;   &lt;type&gt;event&lt;/type&gt;   &lt;file&gt;file-name&lt;/file&gt;   &lt;URL&gt;URL&lt;/URL&gt; &lt;/request-script-refresh-from&gt;</pre>
<b>Options</b>	<p><b>file <i>file-name</i></b>—Name of the file to be downloaded.</p> <p><b>url <i>url-path</i></b>—URL of the file to be downloaded.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Automatic Refreshing of Scripts on EX Series Switches on page 613</a></li> <li>• <a href="#">Junos OS Junos XML Management Protocol Guide</a></li> <li>• <a href="#">Junos OS NETCONF XML Management Protocol Guide</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system scripts refresh-from event file config.txt url http://host1.juniper.net on page 889</a>

### Sample Output

`request system scripts refresh-from event file config.txt url http://host1.juniper.net`

```
user@switch> request system scripts refresh-from event file config.txt url http://host1.juniper.net
user@switch>
```

## request system scripts refresh-from op

---

<b>Syntax</b>	<code>request system scripts refresh-from op file <i>file-name</i> url <i>url-path</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.1 for EX Series switches.
<b>Description</b>	<p>Automatically download the initial Junos OS configuration and a set of standard op scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.</p> <p>The Junos XML management protocol equivalent for this operational mode command is:</p> <pre>&lt;request-script-refresh-from&gt;   &lt;type&gt;op&lt;/type&gt;   &lt;file&gt;file-name&lt;/file&gt;   &lt;URL&gt;URL&lt;/URL&gt; &lt;/request-script-refresh-from&gt;</pre>
<b>Options</b>	<p><b>file <i>file-name</i></b>—Name of the file to be downloaded.</p> <p><b>url <i>url-path</i></b>—URL of the file to be downloaded.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Automatic Refreshing of Scripts on EX Series Switches on page 613</a></li><li>• <a href="#">Junos OS Junos XML Management Protocol Guide</a></li><li>• <a href="#">Junos OS NETCONF XML Management Protocol Guide</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request system scripts refresh-from op file config.txt url http://host1.juniper.net on page 890</a>

### Sample Output

`request system scripts refresh-from op file config.txt url http://host1.juniper.net`

```
user@switch> request system scripts refresh-from op file config.txt url http://host1.juniper.net
user@switch>
```

## show chassis alarms

<b>List of Syntax</b>	<a href="#">Syntax on page 891</a> <a href="#">Syntax (TX Matrix Routers) on page 891</a> <a href="#">Syntax (TX Matrix Plus Routers) on page 891</a> <a href="#">Syntax (MX Series Routers) on page 891</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 891</a> <a href="#">Syntax (QFX Series) on page 891</a> <a href="#">Syntax (PTX Series Packet Transport Routers) on page 891</a> <a href="#">Syntax (ACX Series Universal Access Routers) on page 891</a>
<b>Syntax</b>	show chassis alarms
<b>Syntax (TX Matrix Routers)</b>	show chassis alarms <lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Routers)</b>	show chassis alarms <lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Routers)</b>	show chassis alarms <all-members> <local> <member <i>member-id</i> >
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)</b>	show chassis alarms
<b>Syntax (QFX Series)</b>	show chassis alarms <interconnect-device <i>name</i> > <node-device <i>name</i> >
<b>Syntax (PTX Series Packet Transport Routers)</b>	show chassis alarms
<b>Syntax (ACX Series Universal Access Routers)</b>	show chassis alarms
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option for the TX Matrix Plus router introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>

Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.

**Description** Display information about the conditions that have been configured to trigger alarms.

**Options** **none**—Display information about the conditions that have been configured to trigger alarms.

**all-members**—(MX Series routers only) (Optional) Display information about alarm conditions for all the member routers of the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display information about alarm conditions for the Interconnect device.

**lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display information about alarm conditions for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display information about alarm conditions for the specified member of the Virtual Chassis configuration. Replace *member-id* variable with a value of 0 or 1.

**node-device *name***—(QFabric systems only) (Optional) Display information about alarm conditions for the Node device.

**scc**—(TX Matrix router only) (Optional) Show information about the TX Matrix router (switch-card chassis).

**sfc *number***—(TX Matrix Plus router only) (Optional) Show information about the respective TX Matrix Plus router, which is the switch-fabric chassis. Replace *number* variable with 0.

**Additional Information** You cannot clear the alarms for chassis components. Instead, you must remedy the cause of the alarm. When a chassis alarm LED is lit, it indicates that you are running the router or switch in a manner that we do not recommend.

On routers, you can manually silence external devices connected to the alarm relay contacts by pressing the alarm cutoff button, located on the craft interface. Silencing the device does not remove the alarm messages from the display (if present on the

router) or extinguish the alarm LEDs. In addition, new alarms that occur after you silence an external device reactivate the external device.

In Junos OS release 11.1 and later, alarms for fans also show the slot number of the fans in the CLI output.

In Junos OS Release 11.2 and later, the command output on EX8200 switches shows the detailed location (**Plane/FPC/PFE**) for link errors in the chassis.

In Junos OS Release 10.2 and later, an alarm is shown on T Series routers for a standby sonic clock generator (SCG) that is offline or absent.

You may often see the following error messages, in which only the error code is shown and no other information is provided:

```
Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code:
257
Apr 12 08:04:19 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code:
559
```

To understand what CM\_ALARM error codes mean, you need to first identify the structure of the CM Alarm codes. A CM\_ALARM code has the following structure:

Bits:	Error type:
1-31	Major (1)
0	Minor (0)

According to the table above, the LSB (bit 0) identifies the **Error Type** (major alarm, if the bit is set and minor alarm if the bit is unset). The rest of the bits (1 - 31) identify the actual error code.

Take an example of the following error code, which was logged on a T1600:

```
Apr 12 08:04:10 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code:
559
```

First, you have to convert 559 to binary; that is **100010111**. The LSB in this case is 1, which means that this is a major alarm. After removing the LSB, you are left with **10001011**, which is equal to 279 in decimal. This is the actual error code, its meaning can be found from the following list:

Chip Type: L Chip	Code
CMALARM_LCHIP_LOUT_DESRD_PARITY_ERR	1
CMALARM_LCHIP_LOUT_DESRD_UNINIT_ERR	2
CMALARM_LCHIP_LOUT_DESRD_ILLEGALLINK_ERR	3
CMALARM_LCHIP_LOUT_DESRD_ILLEGALSIZE_ERR	4

CMALARM_LCHIP_LOUT_HDRF_TOERR_ERR	5
CMALARM_LCHIP_LOUT_HDRF_PARITY_ERR	6
CMALARM_LCHIP_LOUT_HDRF_UCERR_ERR	7
CMALARM_LCHIP_LOUT_NLIF_CRCDROP_ERR	8
CMALARM_LCHIP_LOUT_NLIF_CRCERR_ERR	9
CMALARM_LCHIP_UCODE_TIMEOUT_ERR	10
CMALARM_LCHIP_LIN_SRCTL_ACCT_DROP_ERR	11
CMALARM_LCHIP_LIN_SRCTL_ACCT_ADDR_SIZE_ERR	12
CMALARM_LCHIP_SRAM_PARITY_ERR	13
CMALARM_LCHIP_UCODE_OVFLW_ERR	14
CMALARM_LCHIP_LOUT_HDRF_MTU_ERR	15

Chip Type: M Chip	Code
CMALARM_MCHIP_ECC_UNCORRECT_ERR	128

Chip Type: N Chip	Code
CMALARM_NCHIP_RDDMA_JBUS_TIMEOUT_ERR	256
CMALARM_NCHIP_RDDMA_FIFO_OVFLW_ERR	257
CMALARM_NCHIP_RDDMA_FIFO_UNFLW_ERR	258
CMALARM_NCHIP_RDDMA_SIZE_ERR	259
CMALARM_NCHIP_RDDMA_JBUS_CRC_ERR	260
CMALARM_NCHIP_WRDMA_PKTR_ERR	261
CMALARM_NCHIP_WRDMA_PKT_CRC_ERR	262
CMALARM_NCHIP_WRDMA_JBUS_TIMEOUT_ERR	263
CMALARM_NCHIP_WRDMA_FIFO_OVFLW_ERR	264
CMALARM_NCHIP_WRDMA_FIFO_UNFLW_ERR	265
CMALARM_NCHIP_WRDMA_PKT_LEN_ERR	266

CMALARM_NCHIP_WRDMA_JBUS_CRC_ERR	267
CMALARM_NCHIP_PKTR_DMA_AGE_ERR	268
CMALARM_NCHIP_PKTR_ICELLSIG_ERR	269
CMALARM_NCHIP_PKTR_FTTL_ERR	270
CMALARM_NCHIP_RODR_OFFSET_OVFLW_ERR	271
CMALARM_NCHIP_PKTR_TMO_CELL_ERR	272
CMALARM_NCHIP_PKTR_TMO_OUTRANGE_ERR	273
CMALARM_NCHIP_PKTR_MD_REQUEST_Q_OVFLW_ERR	274
CMALARM_NCHIP_PKTR_DMA_BUFFER_OVFLW_ERR	275
CMALARM_NCHIP_PKTR_GRT_OVFLW_ERR	276
CMALARM_NCHIP_FRQ_ERR	277
CMALARM_NCHIP_RODR_IN_Q_OVFLW_ERR	278
CMALARM_NCHIP_DBUF_CRC_ERR	279

Chip Type: R Chip	Code
CMALARM_RCHIP_SRAM_PARITY_ERR	512

Chip Type: R Chip	Code
CMALARM_ICHIP_WO_DESRD_ID_ERR	601
CMALARM_ICHIP_WO_DESRD_DATA_ERR	602
CMALARM_ICHIP_WO_DESRD_OFLOW_ERR	603
CMALARM_ICHIP_WO_HDRF_UCERR_ERR	604
CMALARM_ICHIP_WO_HDRF_MTUERR_ERR	605
CMALARM_ICHIP_WO_HDRF_PARITY_ERR	606
CMALARM_ICHIP_WO_HDRF_TOERR_ERR	607
CMALARM_ICHIP_WO_IP_CRC_ERR	608
CMALARM_ICHIP_WO_IP_INTER_ERR	609

CMALARM_ICHIP_WI_WAN_TIMEOUT_ERR	625
CMALARM_ICHIP_WI_FAB_TIMEOUT_ERR	626
CMALARM_ICHIP_RLDRAM_BIST_ERR	630
CMALARM_ICHIP_SDRAM_BIST_ERR	631
CMALARM_ICHIP_RLDRAM_PARITY_ERR	632
CMALARM_ICHIP_SDRAM_UNCORRECT_ERR	633
CMALARM_ICHIP_SDRAM_CORRECT_ERR	634
CMALARM_ICHIP_FUSE_DONE_ERR	635

According to the table above, the **279** error code corresponds to **CMALARM\_NCHIP\_DBUF\_CRC\_ERR**; this means that new CRC errors were seen on the NCHIP of this particular FPC, which is FPC as per the logs.

If you do not want to convert decimal to binary and vice versa, you may use the following shortcut:

For major alarms, the **Actual Error Code = (Error Code - 1)/2**, where **Error Code** is the code that you get in the log message. For example, if you get the following log:

Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code: 257

Actual Error Code =  $(257-1)/2 = 128$ . Similarly, for minor alarms, Actual Error Code =  $(\text{Error Code})/2$

**Required Privilege Level** view

**Related Documentation**

- *Configuring an RMON Alarm Entry and Its Attributes*
- *Chassis Conditions That Trigger Alarms*

**List of Sample Output**

- [show chassis alarms \(Alarms Active\) on page 897](#)
- [show chassis alarms \(No Alarms Active\) on page 898](#)
- [show chassis alarms \(Fan Tray\) on page 898](#)
- [show chassis alarms \(MX104 Router\) on page 898](#)
- [show chassis alarms \(MX2010 Router\) on page 898](#)
- [show chassis alarms \(MX2020 Router\) on page 898](#)
- [show chassis alarms \(MX960, MX480, and MX240 Routers showing Major CB Failure\) on page 898](#)
- [show chassis alarms \(T4000 Router\) on page 899](#)
- [show chassis alarms \(Unreachable Destinations Present on a T Series Router\) on page 899](#)



[show chassis alarms \(FPC Offline Due to Unreachable Destinations on a T Series Router\) on page 899](#)  
[show chassis alarms \(SCG Absent on a T Series Router\) on page 899](#)  
[show chassis alarms \(Alarms Active on a TX Matrix Router\) on page 899](#)  
[show chassis alarms \(TX Matrix Plus router with 3D SIBs\) on page 900](#)  
[show chassis alarms \(Alarms on a T4000 Router After the enhanced-mode Statement is Enabled\) on page 902](#)  
[show chassis alarms \(Backup Routing Engine\) on page 902](#)  
[show chassis alarms \(EX Series Switch\) on page 902](#)  
[show chassis alarms \(Alarms Active on the QFX Series\) on page 902](#)  
[show chassis alarms node-device \(Alarms Active on the QFabric System\) on page 902](#)  
[show chassis alarms \(Alarms Active on the QFabric System\) on page 903](#)  
[show chassis alarms \(Alarms Active on an EX8200 Switch\) on page 903](#)  
[show chassis alarms \(Alarms Active on a PTX5000 Packet Transport Router\) on page 903](#)  
[show chassis alarms \(Mix of PDUs Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 904](#)  
[show chassis alarms \(PDU Converter Failed Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 904](#)  
[show chassis alarms \(No Power for System Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 904](#)  
[show chassis alarms \(Alarms Active on an ACX2000 Universal Access Router\) on page 904](#)  
[show chassis alarms \(Active Alarm to Indicate Status of the Bad SCB Clock on MX Series\) on page 905](#)

**Output Fields** [Table 105 on page 897](#) lists the output fields for the **show chassis alarms** command. Output fields are listed in the approximate order in which they appear.

**Table 105: show chassis alarms Output Fields**

Field Name	Field Description
Alarm time	Date and time the alarm was first recorded.
Class	Severity class for this alarm: <b>Minor</b> or <b>Major</b> .
Description	Information about the alarm.

## Sample Output

**show chassis alarms (Alarms Active)**

```

user@host> show chassis alarms
3 alarms are currently active
Alarm time           Class  Description
2000-02-07 10:12:22 UTC Major fxp0: ethernet link down
2000-02-07 10:11:54 UTC Minor YELLOW ALARM - PEM 1 Removed
2000-02-07 10:11:03 UTC Minor YELLOW ALARM - Lower Fan Tray Removed

```

### show chassis alarms (No Alarms Active)

```
user@host> show chassis alarms
No alarms are currently active
```

### show chassis alarms (Fan Tray)

```
user@host> show chassis alarms
4 alarms currently active
Alarm time          Class Description
2010-11-11 20:27:38 UTC Major Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC Minor Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC Major Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC Major Side Fan Tray 0 Failure
```

### show chassis alarms (MX104 Router)

```
user@host >show chassis alarms
1 alarms currently active
Alarm time          Class Description
2013-06-05 14:43:31 IST Minor Backup RE Active
```

### show chassis alarms (MX2010 Router)

```
user@host> show chassis alarms
7 alarms currently active
Alarm time          Class Description
2012-08-07 00:46:06 PDT Major Fan Tray 2 Failure
2012-08-06 18:24:36 PDT Minor Redundant feed missing for PSM 6
2012-08-06 07:41:04 PDT Minor Redundant feed missing for PSM 8
2012-08-04 02:42:06 PDT Minor Redundant feed missing for PSM 5
2012-08-03 21:14:24 PDT Minor Loss of communication with Backup RE
2012-08-03 12:26:03 PDT Minor Redundant feed missing for PSM 4
2012-08-03 10:40:18 PDT Minor Redundant feed missing for PSM 7
```

### show chassis alarms (MX2020 Router)

```
user@host> show chassis alarms
1 alarms currently active
Alarm time Class Description
2012-10-03 12:14:59 PDT Minor Plane 0 not online
```

### show chassis alarms (MX960, MX480, and MX240 Routers showing Major CB Failure)

A Major CB 0 Failure alarm occurs in the event of a bad CB (unknown or mismatched CBs do not trigger this alarm in Junos Release 12.3R9 and later). Following GRES/recovery, if the hardware issue persists, the traffic moves to the good CB and continues. If the alarm was triggered by something transient like a power zone budget on GRES, bringing the CB back online can clear the alarm. Otherwise, replace the bad CB. Note that fabric link speed is not impacted by an offline SCB. The alarm may be raised on CB0, CB1, and CB2.

```
user@host> show chassis alarms
6 alarms currently active
Alarm time          Class Description
2014-10-31 16:49:41 EDT Major PEM 3 Not OK
2014-10-31 16:49:41 EDT Major PEM 2 Not OK
2014-10-31 16:49:31 EDT Major CB 0 Failure
2014-10-31 16:49:31 EDT Minor CB 0 Fabric Chip 0 Not Online
2014-10-31 16:49:31 EDT Minor CB 0 Fabric Chip 1 Not Online
2014-10-31 16:49:31 EDT Minor Backup RE Active
```

**show chassis alarms (T4000 Router)**

```

user@host> show chassis alarms
9 alarms currently active
Alarm time          Class Description
2007-06-02 01:41:10 UTC Minor RE 0 Not Supported
2007-06-02 01:41:10 UTC Minor CB 0 Not Supported
2007-06-02 01:41:10 UTC Minor Mixed Master and Backup RE types
2007-05-30 19:37:33 UTC Major SPMB 1 not online
2007-05-30 19:37:29 UTC Minor Front Bottom Fan Tray Absent
2007-05-30 19:37:13 UTC Major PEM 1 Input Failure
2007-05-30 19:37:13 UTC Major PEM 0 Not OK
2007-05-30 19:37:03 UTC Major PEM 0 Improper for Platform
2007-05-30 19:37:03 UTC Minor Backup RE Active

```

**show chassis alarms (Unreachable Destinations Present on a T Series Router)**

```

user@host> show chassis alarms
10 alarms currently active
Alarm time          Class Description
2011-08-30 18:43:53 PDT Major FPC 7 has unreachable destinations
2011-08-30 18:43:53 PDT Major FPC 5 has unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 3 has unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 2 has unreachable destinations
2011-08-30 18:43:52 PDT Minor SIB 0 Not Online
2011-08-30 18:43:33 PDT Minor SIB 4 Not Online
2011-08-30 18:43:28 PDT Minor SIB 3 Not Online
2011-08-30 18:43:05 PDT Minor SIB 2 Not Online
2011-08-30 18:43:28 PDT Minor SIB 1 Not Online
2011-08-30 18:43:05 PDT Major PEM 1 Not Ok

```

**show chassis alarms (FPC Offline Due to Unreachable Destinations on a T Series Router)**

```

user@host> show chassis alarms
10 alarms currently active
Alarm time          Class Description
2011-08-30 18:43:53 PDT Major FPC 7 offline due to unreachable destinations
2011-08-30 18:43:53 PDT Major FPC 5 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 3 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 2 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Minor SIB 0 Not Online
2011-08-30 18:43:33 PDT Minor SIB 4 Not Online
2011-08-30 18:43:28 PDT Minor SIB 3 Not Online
2011-08-30 18:43:05 PDT Minor SIB 2 Not Online
2011-08-30 18:43:28 PDT Minor SIB 1 Not Online
2011-08-30 18:43:05 PDT Major PEM 1 Not Ok

```

**show chassis alarms (SCG Absent on a T Series Router)**

```

user@host> show chassis alarms
4 alarms currently active
Alarm time          Class Description
2011-01-23 21:42:46 PST Major SCG 0 NO EXT CLK MEAS-BKUP SCG ABS

```

**show chassis alarms (Alarms Active on a TX Matrix Router)**

```

user@host> show chassis alarms
scc-re0:
-----
8 alarms currently active
Alarm time          Class Description

```

```

2004-08-05 18:43:53 PDT Minor LCC 0 Minor Errors
2004-08-05 18:43:53 PDT Minor SIB 3 Not Online
2004-08-05 18:43:52 PDT Major SIB 2 Absent
2004-08-05 18:43:52 PDT Major SIB 1 Absent
2004-08-05 18:43:52 PDT Major SIB 0 Absent
2004-08-05 18:43:33 PDT Major LCC 2 Major Errors
2004-08-05 18:43:28 PDT Major LCC 0 Major Errors
2004-08-05 18:43:05 PDT Minor LCC 2 Minor Errors
lcc0-re0:

```

```
-----
5 alarms currently active

```

Alarm time	Class	Description
2004-08-05 18:43:53 PDT	Minor	SIB 3 Not Online
2004-08-05 18:43:49 PDT	Major	SIB 2 Absent
2004-08-05 18:43:49 PDT	Major	SIB 1 Absent
2004-08-05 18:43:49 PDT	Major	SIB 0 Absent
2004-08-05 18:43:28 PDT	Major	PEM 0 Not OK

```
lcc2-re0:

```

```
-----
5 alarms currently active

```

Alarm time	Class	Description
2004-08-05 18:43:35 PDT	Minor	SIB 3 Not Online
2004-08-05 18:43:33 PDT	Major	SIB 2 Absent
2004-08-05 18:43:33 PDT	Major	SIB 1 Absent
2004-08-05 18:43:33 PDT	Major	SIB 0 Absent
2004-08-05 18:43:05 PDT	Minor	PEM 1 Absent

#### show chassis alarms (TX Matrix Plus router with 3D SIBs)

```

user@host> show chassis alarms
sfc0-re0:

```

```

-----
Alarm time          Class  Description

2014-04-08 14:35:13 IST Minor  FPM 0 SFC Config Size Changed
2014-04-08 14:32:58 IST Major  Fan Tray Failure
2014-04-08 14:31:53 IST Major  SIB F13 6 Fault
2014-04-08 14:31:43 IST Major  SIB F13 11 Fault
2014-04-08 14:31:08 IST Minor  Check SIB F13 12 CXP 14 Fbr Cbl
2014-04-08 14:31:08 IST Minor  Check SIB F13 12 CXP 8 Fbr Cbl
2014-04-08 14:31:08 IST Minor  Check SIB F13 12 CXP 3 Fbr Cbl
2014-04-08 14:31:08 IST Major  SIB F13 12 CXP 15 fault
2014-04-08 14:31:08 IST Minor  SIB F13 12 CXP 14 LOL
2014-04-08 14:31:08 IST Minor  Check SIB F13 12 CXP 14
2014-04-08 14:31:08 IST Major  SIB F13 12 CXP 10 fault
2014-04-08 14:31:08 IST Minor  SIB F13 12 CXP 8 LOL
2014-04-08 14:31:08 IST Minor  Check SIB F13 12 CXP 8
2014-04-08 14:31:08 IST Major  SIB F13 12 CXP 7 fault
2014-04-08 14:31:08 IST Major  SIB F13 12 CXP 4 fault
2014-04-08 14:31:08 IST Minor  SIB F13 12 CXP 3 LOL
2014-04-08 14:31:08 IST Minor  Check SIB F13 12 CXP 3
2014-04-08 14:31:08 IST Minor  Check SIB F13 6 CXP 14 Fbr Cbl
2014-04-08 14:31:08 IST Minor  Check SIB F13 6 CXP 12 Fbr Cbl
2014-04-08 14:31:08 IST Minor  Check SIB F13 6 CXP 8 Fbr Cbl
2014-04-08 14:31:08 IST Minor  Check SIB F13 6 CXP 6 Fbr Cbl
2014-04-08 14:31:08 IST Minor  Check SIB F13 6 CXP 4 Fbr Cbl
2014-04-08 14:31:08 IST Minor  Check SIB F13 6 CXP 2 Fbr Cbl
2014-04-08 14:31:08 IST Minor  Check SIB F13 6 CXP 0 Fbr Cbl
2014-04-08 14:31:08 IST Minor  SIB F13 6 CXP 14 LOL
2014-04-08 14:31:08 IST Minor  Check SIB F13 6 CXP 14
2014-04-08 14:31:08 IST Minor  SIB F13 6 CXP 12 LOL

```

```

2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 12
2014-04-08 14:31:08 IST Major SIB F13 6 CXP 10 fault
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 8 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 8
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 6 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 6
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 4 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 4
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 2 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 2
2014-04-08 14:31:08 IST Minor SIB F13 6 CXP 0 LOL
2014-04-08 14:31:08 IST Minor Check SIB F13 6 CXP 0
2014-04-08 14:31:08 IST Minor SIB F13 12 CXP 14 XC HSL Link Error
2014-04-08 14:29:27 IST Minor LCC 0 Minor Errors
2014-04-08 14:28:37 IST Major LCC 0 Major Errors
2014-04-08 14:28:37 IST Major LCC 2 Major Errors
2014-04-08 14:28:37 IST Minor LCC 2 Minor Errors
2014-04-08 14:28:24 IST Major SIB F2S 4/6 Absent
2014-04-08 14:28:24 IST Major SIB F2S 4/4 Absent
2014-04-08 14:28:24 IST Major SIB F2S 4/2 Absent
2014-04-08 14:28:24 IST Major SIB F2S 4/0 Absent
2014-04-08 14:28:24 IST Major SIB F2S 3/6 Absent
2014-04-08 14:28:24 IST Major SIB F2S 3/4 Absent
2014-04-08 14:28:24 IST Major SIB F2S 3/2 Absent
2014-04-08 14:28:24 IST Major SIB F2S 3/0 Absent
2014-04-08 14:28:24 IST Major SIB F13 9 Absent
2014-04-08 14:28:24 IST Major SIB F13 8 Absent
2014-04-08 14:28:24 IST Major SIB F13 7 Absent
2014-04-08 14:28:24 IST Major SIB F13 4 Absent
2014-04-08 14:28:24 IST Major SIB F13 1 Absent
2014-04-08 14:28:22 IST Major PEM 0 Input Failure
2014-04-08 14:28:22 IST Major PEM 0 Not OK

```

lcc0-re0:

-----

12 alarms currently active

Alarm time	Class	Description
2014-04-08 14:36:08 IST	Minor	CB 1 M/S Switch Changed
2014-04-08 14:36:08 IST	Minor	CB 1 CHASSIS ID Changed
2014-04-08 14:35:43 IST	Minor	CB 0 M/S Switch Changed
2014-04-08 14:35:43 IST	Minor	CB 0 CHASSIS ID Changed
2014-04-08 14:29:30 IST	Minor	SIB 4 Not Online
2014-04-08 14:29:30 IST	Minor	SIB 3 Not Online
2014-04-08 14:29:30 IST	Minor	SIB 2 Not Online
2014-04-08 14:29:24 IST	Major	Rear Fan Tray Failure
2014-04-08 14:29:24 IST	Major	Front Bottom Fan Tray Improper for Platform
2014-04-08 14:29:24 IST	Major	Front Top Fan Tray Improper for Platform
2014-04-08 14:28:37 IST	Major	SIB 4 Absent
2014-04-08 14:28:37 IST	Major	SIB 3 Absent

lcc2-re0:

-----

12 alarms currently active

Alarm time	Class	Description
2014-04-08 14:36:02 IST	Minor	CB 1 M/S Switch Changed
2014-04-08 14:36:02 IST	Minor	CB 1 CHASSIS ID Changed
2014-04-08 14:35:42 IST	Minor	CB 0 M/S Switch Changed
2014-04-08 14:34:42 IST	Minor	CB 0 CHASSIS ID Changed
2014-04-08 14:29:29 IST	Minor	SIB 0 CXP 7 Unsupported Optics
2014-04-08 14:29:27 IST	Major	Front Bottom Fan Tray Improper for Platform
2014-04-08 14:29:27 IST	Major	Front Top Fan Tray Improper for Platform

```
2014-04-08 14:29:25 IST Minor SIB 4 Not Online
2014-04-08 14:29:25 IST Minor SIB 3 Not Online
2014-04-08 14:28:47 IST Major PEM 0 Not OK
2014-04-08 14:28:36 IST Major SIB 2 Absent
2014-04-08 14:28:36 IST Minor Host 0 Boot from alternate media
```

```
lcc6-re0:
```

```
-----
2 alarms currently active
```

Alarm time	Class	Description
2013-11-06 04:03:56 PST	Minor	SIB 1 CXP 0 XC HSL Link Error
2013-11-06 03:49:32 PST	Major	PEM 1 Not OK

### show chassis alarms (Alarms on a T4000 Router After the enhanced-mode Statement is Enabled)

To enable improved virtual private LAN service (VPLS) MAC address learning on T4000 routers, you must include the **enhanced-mode** statement at the **[edit chassis network-services]** hierarchy level and reboot the router. When router reboots, only the T4000 Type 5 FPCs are required to be present on the router. If there are any other FPCs (apart from T4000 Type 5 FPCs) on the T4000 router, such FPCs become offline, and FPC misconfiguration alarms are generated. The **show chassis alarm** command output displays FPC misconfiguration (**FPC *fpc-slot* misconfig**) as the reason for the generation of the alarms.

```
user@host> show chassis alarms
2 alarms currently active
Alarm time          Class  Description
2011-10-22 10:10:47 PDT Major  FPC 1 misconfig
2011-10-22 10:10:46 PDT Major  FPC 0 misconfig
```

### show chassis alarms (Backup Routing Engine)

```
user@host> show chassis alarms
2 alarms are currently active
Alarm time          Class  Description
2005-04-07 10:12:22 PDT Minor  Host 1 Boot from alternate media
2005-04-07 10:11:54 PDT Major  Host 1 compact-flash missing in Boot List
```

### show chassis alarms (EX Series Switch)

```
user@switch> show chassis alarms
4 alarms currently active
Alarm time          Class  Description
2014-03-12 15:36:09 UTC Minor  Require a Fan Tray upgrade
2014-03-12 15:00:02 UTC Major  PEM 0 Input Failure
2014-03-12 15:00:02 UTC Major  PEM 0 Not OK
2014-03-12 14:59:51 UTC Minor  Host 1 Boot from alternate media
```

### show chassis alarms (Alarms Active on the QFX Series)

```
user@switch> show chassis alarms
1 alarms currently active
Alarm time          Class  Description
2012-03-05 2:10:24 UTC Major  FPC 0 PEM 0 Airflow not matching Chassis Airflow
```

### show chassis alarms node-device (Alarms Active on the QFabric System)

```
user@switch> show chassis alarms node-device ED3691
node-device ED3694
3 alarms currently active
```

Alarm time	Class	Description
2011-08-24 16:04:15 UTC	Major	ED3694:fte-0/1/2: Link down
2011-08-24 16:04:14 UTC	Major	ED3694:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC	Major	ED3694 PEM 0 is not supported/powered

### show chassis alarms (Alarms Active on the QFabric System)

```
user@switch> show chassis alarms
IC-A0001:
```

```
-----
1 alarms currently active
Alarm time      Class  Description
2011-08-24 16:04:15 UTC  Minor  Backup RE Active
```

```
ED3694:
```

```
-----
3 alarms currently active
Alarm time      Class  Description
2011-08-24 16:04:15 UTC  Major  ED3694:fte-0/1/2: Link down
2011-08-24 16:04:14 UTC  Major  ED3694:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC  Major  ED3694 PEM 0 is not supported/powered
```

```
SNG-0:
```

```
NW-NG-0:
```

```
-----
1 alarms currently active
Alarm time      Class  Description
2011-08-24 15:49:27 UTC  Major  ED3691 PEM 0 is not supported/powered
```

### show chassis alarms (Alarms Active on an EX8200 Switch)

```
user@switch> show chassis alarms
```

```
6 alarms currently active
Alarm time      Class  Description
2010-12-02 19:15:22 UTC  Major  Fan Tray Failure
2010-12-02 19:15:22 UTC  Major  Fan Tray Failure
2010-12-02 19:15:14 UTC  Minor  Check CB 0 Fabric Chip 1 on Plane/FPC/PFE: 1/5/0,
1/5/1, 1/5/2, 1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:15:14 UTC  Minor  Check CB 0 Fabric Chip 0 on Plane/FPC/PFE: 1/5/0,
1/5/1, 1/5/2, 1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:14:18 UTC  Major  PSU 1 Output Failure
2010-12-02 19:14:18 UTC  Minor  Loss of communication with Backup RE
```

### show chassis alarms (Alarms Active on a PTX5000 Packet Transport Router)

```
user@host> show chassis alarms
```

```
23 alarms currently active
Alarm time      Class  Description
2011-07-12 16:22:05 PDT  Minor  No Redundant Power for Rear Chassis
2011-07-12 16:22:05 PDT  Major  PDU 0 PSM 1 Not OK
2011-07-12 16:21:57 PDT  Minor  No Redundant Power for Fan 0-2
2011-07-12 16:21:57 PDT  Major  PDU 0 PSM 0 Not OK
2011-07-12 15:56:06 PDT  Major  PDU 1 PSM 2 Not OK
2011-07-12 15:56:06 PDT  Minor  No Redundant Power for FPC 0-7
2011-07-12 15:56:06 PDT  Major  PDU 0 PSM 3 Not OK
2011-07-12 15:28:20 PDT  Major  PDU 0 PSM 2 Not OK
2011-07-12 15:19:14 PDT  Minor  Backup RE Active
```

### show chassis alarms (Mix of PDUs Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-PIA)

All PDUs installed on a PTX5000 router must be of the same type. The **Mix of PDUs or Power Manager Non Operational** alarm is raised when different types of PDUs are installed on a PTX5000 router.

```
user@host> show chassis alarms
15 alarms currently active
Alarm time          Class Description
2013-03-19 23:03:53 PDT Minor No Redundant Power
2013-03-19 23:03:48 PDT Minor Mix of PDUs
2013-03-19 23:03:47 PDT Minor PDU 1 PSM 3 Absent
2013-03-19 23:03:47 PDT Minor PDU 1 PSM 2 Absent
2013-03-19 23:03:47 PDT Minor PDU 1 PSM 1 Absent
2013-03-19 23:03:47 PDT Minor PDU 1 PSM 0 Absent
2013-03-19 23:03:46 PDT Major No CG Online
```

### show chassis alarms (PDU Converter Failed Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-PIA)

The **PDU Converter Failed** alarm is raised when one or more 36 V booster converter of a DC PDU fails. If two or more 36 V booster converter fails, fan trays fail and the router might get over heated. Therefore, when this alarm is raised, check the PDU and replace it, if required.

```
user@host> show chassis alarms
11 alarms currently active
Alarm time          Class Description
2013-12-11 22:14:13 PST Minor No Redundant Power for System
2013-12-11 22:14:10 PST Major PDU 0 PSM 7 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 6 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 5 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 4 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 3 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 2 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 1 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 0 Not OK
2013-12-11 22:14:10 PST Major PDU 0 Not OK
2013-12-11 22:14:01 PST Major PDU 0 Converter Failed
```

### show chassis alarms (No Power for System Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-PIA)

```
user@host> show chassis alarms
8 alarms currently active
Alarm time          Class Description
2013-11-19 01:58:41 PST Major No Power for System
2013-11-19 01:58:37 PST Major PDU 0 PSM 1 Not OK
2013-11-19 01:56:46 PST Major PDU 0 PSM 2 Not OK
2013-11-19 01:54:26 PST Major PDU 0 PSM 3 Not OK
2013-11-19 01:53:30 PST Major PDU 1 PSM 3 Not OK
2013-11-19 01:53:29 PST Major PDU 1 PSM 2 Not OK
2013-11-19 01:53:29 PST Major PDU 1 PSM 1 Not OK
2013-11-19 01:53:29 PST Major PDU 1 PSM 0 Not OK
```

### show chassis alarms (Alarms Active on an ACX2000 Universal Access Router)

```
user@host> show chassis alarms
7 alarms currently active
Alarm time          Class Description
2012-05-22 11:19:09 UTC Major xe-0/3/1: Link down
2012-05-22 11:19:09 UTC Major xe-0/3/0: Link down
```



```
2012-05-22 11:19:09 UTC Major ge-0/1/7: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/6: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/3: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/2: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/1: Link down
```

#### show chassis alarms (Active Alarm to Indicate Status of the Bad SCB Clock on MX Series)

```
user@host> show chassis alarms
1 alarm currently active
Alarm time      Class  Description
2013-08-06 07:48:35 PDT Major  CB 0 19.44 MHz clock failure
```

## show chassis environment

---

- List of Syntax**
- Syntax on page 906
  - Syntax (T320, T640, T1600, and T4000 Routers) on page 906
  - Syntax (TX Matrix Routers) on page 906
  - Syntax (TX Matrix Plus Routers) on page 906
  - Syntax (MX Series Routers) on page 906
  - Syntax (MX104 3D Universal Edge Routers) on page 906
  - Syntax (MX2010 and MX2020 3D Universal Edge Routers) on page 907
  - Syntax (EX8200 Switches) on page 907
  - Syntax (EX Series Switches except EX8200) on page 907
  - Syntax (QFX Series) on page 907
  - Syntax (PTX Series Packet Transport Routers) on page 907
  - Syntax (ACX Series Universal Access Routers) on page 907

**Syntax** show chassis environment

**Syntax (T320, T640, T1600, and T4000 Routers)**

```
show chassis environment
<cb cb-slot-number>
<fpc fpc-slot-number>
<fpm>
<pem pem-slot-number>
<routing-engine re-slot-number>
<scg scg-slot-number>
<sib sib-slot-number>
```

**Syntax (TX Matrix Routers)**

```
show chassis environment
<lcc number | scc>
```

**Syntax (TX Matrix Plus Routers)**

```
show chassis environment
<cb cb-slot-number>
<cip cip-slot-number>
<fpc fpc-slot-number>
<fpm>
<lcc number>
<pem pem-slot-number>
<routing-engine re-slot-number>
<scg scg-slot-number>
<sfc number>
<sib sib-slot-number>
```

**Syntax (MX Series Routers)**

```
show chassis environment
<all-members>
<local>
<member member-id>
```

**Syntax (MX104 3D Universal Edge Routers)**

```
show chassis environment
<cb>
<pem pem-slot-number>
<routing-engine re-slot-number>
```

Syntax (MX2010 and MX2020 3D Universal Edge Routers)	<pre> show chassis environment &lt;adc <i>adc-slot-number</i>&gt; &lt;cb <i>cb-slot-number</i>&gt; &lt;fpc <i>fpc-slot-number</i>&gt; &lt;fpm&gt; &lt;monitored&gt; &lt;psm <i>psm-slot-number</i>&gt; &lt;routing-engine <i>re-slot-number</i>&gt; &lt;sfb <i>sfb-slot-number</i>&gt; </pre>
Syntax (EX8200 Switches)	<pre> show chassis environment &lt;all-members&gt; &lt;cb <i>cb-slot-number</i>&gt; &lt;fpc <i>fpc-slot-number</i>&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;psu <i>psu-slot-number</i>&gt; &lt;routing-engine <i>re-slot-number</i>&gt; </pre>
Syntax (EX Series Switches except EX8200)	<pre> show chassis environment &lt;all-members&gt; &lt;fpc <i>fpc-slot-number</i>&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;power-supply-unit&gt; &lt;routing-engine&gt; </pre>
Syntax (QFX Series)	<pre> show chassis environment &lt;cb <i>slot-number</i> &lt;interconnect-device name&gt;&gt; &lt;fpc <i>slot-number</i> &lt;interconnect-device name&gt;&gt; &lt;interconnect-device name &lt;slot-number&gt; &lt;node-device name&gt; &lt;pem <i>slot-number</i> (interconnect-device name <i>slot-number</i>)   (node-device name)&gt; &lt;routing-engine name &lt;interconnect-device name <i>slot-number</i>&gt;&gt; </pre>
Syntax (PTX Series Packet Transport Routers)	<pre> show chassis environment &lt;cb <i>cb-slot-number</i>&gt; &lt;ccg <i>ccg-slot-number</i>&gt; &lt;fpc <i>fpc-slot-number</i>&gt; &lt;fpm&gt; &lt;monitored&gt; &lt;pdu <i>pdu-slot-number</i>&gt; &lt;routing-engine <i>re-slot-number</i>&gt; &lt;sib <i>sib-slot-number</i>&gt; </pre>
Syntax (ACX Series Universal Access Routers)	<pre> show chassis environment &lt;cb <i>cb-slot-number</i>&gt; &lt;pem <i>pem-slot-number</i>&gt; &lt;routing-engine <i>re-slot-number</i>&gt; </pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p>

Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

**monitored** option added in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

Command introduced in Junos OS Release 12.1 for T4000 Core Routers.

Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.

Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.

**pem** option introduced in Junos OS Release 12.3 for ACX4000 Universal Access Routers.

Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.

**Description** Display environmental information about the router or switch chassis, including the temperature and information about the fans, power supplies, and Routing Engine.

In addition, on ACX4000 routers, display temperature information about the different channels of a Modular Interface Card (MIC). The number of channels displayed depends on the type of MIC installed.

Starting with Junos OS Release 14.1, the **show chassis environment cb cb-slot-number | ccg ccg-slot-number | fpc fpc-slot-number | fpm | monitored | pdu pdu-slot-number | routing-engine re-slot-number | sib sib-slot-number** operational mode command output displays environmental information for the the new DC power supply module (PSM) and power distribution unit (PDU) that are added to provide power to the high-density FPC (FPC2-PTX-PIA) and other components in a PTX5000 Packet Transport Router.

**Options** **none**—Display environmental information about the router or switch chassis. On a TX Matrix router, display environmental information about the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about the TX Matrix Plus router and its attached routers.

**all-members**—(MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for all the members of the Virtual Chassis configuration.

**adc adc-slot-number**—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the adapter cards. For MX2020 routers, replace **adc-slot-number** with a value from 0 through 19. For MX2010 routers, replace **adc-slot-number** with a value from 0 through 9.

**cb cb-slot-number**—(ACX Series Universal Access Routers, EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2020 routers, MX2010 routers, PTX Series Packet Transport Routers, QFX Series, and T Series routers, and TX Matrix Plus routers only) (Optional) Display chassis environmental information for the Control Board. On devices other than EX Series switches, replace **cb-slot** with 0 or 1. For the EX Series switches, see [“EX Series Switches Hardware and CLI Terminology Mapping” on page 315](#) for information on CB slot numbering.

**cip cip-slot-number**—(TX Matrix Plus routers only) (Optional) Display chassis environmental information for the Connection Interface Panel (CIP). Replace the **cip-slot-number** variable with a value of 0 or 1.

**cb interconnect-device *name***—(QFabric systems only) (Optional) Display chassis environmental information for the Control Board on an Interconnect device.

**cgc *ccg-slot-number***—(PTX Series only) (Optional) Display chassis environmental information for the Centralized Clock Generator. Replace ***cb-slot*** with a value of 0 or 1.

**fpc *fpc-slot***—(EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2010 routers, MX2020 routers, PTX Series Packet Transport Routers, QFX Series, QFX3500 switches, QFabric systems, T Series routers, and TX Matrix Plus routers) (Optional) Display chassis environmental information for a specified Flexible PIC Concentrator. For MX2010 routers, replace ***fpc-slot*** with a value from 0 through 9. For MX2020 routers, replace ***fpc-slot*** with a value from 0 through 19. For information about FPC numbering, see [show chassis environment fpc](#). On a QFabric system, display chassis environmental information for a specified Flexible PIC Concentrator on an Interconnect device. On an EX Series switch, display chassis environmental information for a specified Flexible PIC Concentrator; see [“EX Series Switches Hardware and CLI Terminology Mapping” on page 315](#) for information on FPC numbering. On a TX Matrix Plus router with 3D SIBs replace ***fpc-slot*** with a value from 0 through 63.

**fpm**—(M120, M320, and M40e routers, MX2010 routers, MX2020 routers, PTX Series, Packet Transport Routers, T Series routers, and TX Matrix Plus routers only) (Optional) Display chassis environmental information for the craft interface (FPM).

**interconnect-device *name***—(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.

**monitored**—(MX2020 routers and PTX Series Packet Transport Routers only) (Optional) Display chassis environmental information for monitored temperatures only. Temperatures that are not included in temperature alarm computations are not displayed.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers and EX Series switches) (Optional) Display chassis environmental information for the local Virtual Chassis member.

**member *member-id***—(MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for the specified member of the Virtual Chassis configuration. On MX Series routers, replace *member-id* variable with a value of 0 or 1. For EX Series switches, see [member](#) for member ID values.

**node-device *name***—(QFabric systems only) (Optional) Display chassis environmental information for the Node device.

**pdu *pdu-slot-number***—(PTX Series only) (Optional) Display chassis environmental information for the specified power distribution unit.

**pem**—(QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Interconnect device or Node device.

**pem *pem-slot-number***—(ACX Series Universal Access Routers, M120, M320, and M40e routers, MX Series routers, MX104 routers, QFX Series, and T Series routers only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Power Entry Module. For information about the options, see *show chassis environment pem*.

**psm *psm-slot-number***—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the power supply module. For MX2020 routers, replace *psm-slot-number* with a value from 0 through 17. For MX2010 routers, replace *psm-slot-number* with a value from 0 through 8.

**psu *psu-slot-number***—(EX Series switches only) (Optional) Display chassis environmental information for a specified power supply. See [“EX Series Switches Hardware and CLI Terminology Mapping” on page 315](#) for detailed information.

**routing-engine**—(QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Routing Engine on the specified Interconnect device.

**routing-engine *re-slot-number***—(Optional) Display chassis environmental information for the specified Routing Engine. For information about the options, see [show chassis environment routing-engine](#).

**scg**—(T Series routers only) (Optional) Display chassis environmental information about the SONET Clock Generator.

**scc**—(TX Matrix routers only) (Optional) Display chassis environmental information about the TX Matrix router (switch-card chassis).

**sfb *sfb-slot-number***—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the power supply module. Replace *sfb-slot-number* with a value from 0 through 7.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display chassis environmental information about the respective TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with 0.

**sib *sib-slot-number***—(M320 routers, PTX Series Packet Transport Routers, and T Series routers only) (Optional) Display chassis environmental information about the specified switch interface board. For information about the options, see *show chassis environment sib*.

**Required Privilege Level**

view

**Related Documentation**

- *show chassis environment adc*
- [show chassis environment cb on page 970](#)
- *show chassis environment ccg*
- *show chassis environment cip*
- [show chassis environment fpc on page 988](#)
- *show chassis environment fpm*
- *show chassis environment lcc*
- *show chassis environment mcs*
- *show chassis environment monitored*
- *show chassis environment pcg*
- *show chassis environment pdu*
- *show chassis environment pem*
- *show chassis environment psm*
- *show chassis environment psu*
- [show chassis environment routing-engine on page 1014](#)
- *show chassis environment scg*
- *show chassis environment sfb*
- *show chassis environment sib*
- *show chassis environment sfc*

**List of Sample Output**

[show chassis environment \(J2300 Router\) on page 914](#)  
[show chassis environment \(J4300 or J6300 Router\) on page 914](#)  
[show chassis environment \(M5 Router\) on page 914](#)  
[show chassis environment \(M7i Router\) on page 915](#)  
[show chassis environment \(M10 Router\) on page 915](#)  
[show chassis environment \(M10i Router\) on page 915](#)  
[show chassis environment \(M20 Router\) on page 916](#)  
[show chassis environment \(M40 Router\) on page 916](#)  
[show chassis environment \(M40e Router\) on page 916](#)  
[show chassis environment \(M120 Router\) on page 917](#)  
[show chassis environment \(M160 Router\) on page 918](#)  
[show chassis environment \(M320 Router\) on page 918](#)

[show chassis environment \(MX104 Router\) on page 919](#)  
[show chassis environment \(MX240 Router\) on page 920](#)  
[show chassis environment \(MX240 Router with SCBE\) on page 921](#)  
[show chassis environment \(MX480 Router\) on page 921](#)  
[show chassis environment \(MX480 Router with SCBE\) on page 922](#)  
[show chassis environment \(MX960 Router\) on page 923](#)  
[show chassis environment \(MX960 Router with SCBE\) on page 924](#)  
[show chassis environment \(MX960 Router with MPC5EQ\) on page 927](#)  
[show chassis environment \(MX2020 Router\) on page 931](#)  
[show chassis environment \(MX2020 Router with MPC5EQ and MPC6E\) on page 940](#)  
[show chassis environment \(MX2010 Router\) on page 944](#)  
[show chassis environment \(T320 Router\) on page 949](#)  
[show chassis environment \(T640 Router\) on page 950](#)  
[show chassis environment \(T4000 Router\) on page 951](#)  
[show chassis environment \(TX Matrix Router\) on page 953](#)  
[show chassis environment \(T1600 Router\) on page 954](#)  
[show chassis environment \(TX Matrix Plus Router\) on page 955](#)  
[show chassis environment \(TX Matrix Plus router with 3D SIBs\) on page 957](#)  
[show chassis environment \(EX4200 Standalone Switch\) on page 960](#)  
[show chassis environment \(EX8216 Switch\) on page 961](#)  
[show chassis environment \(EX9200 Switch\) on page 961](#)  
[show chassis environment \(QFX Series\) on page 962](#)  
[show chassis environment interconnect-device \(QFabric System\) on page 962](#)  
[show chassis environment node-device \(QFabric System\) on page 964](#)  
[show chassis environment pem node-device \(QFabric System\) on page 964](#)  
[show chassis environment \(PTX5000 Packet Transport Router\) on page 965](#)  
[show chassis environment \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 967](#)  
[show chassis environment \(ACX2000 Universal Access Router\) on page 968](#)  
[show chassis environment \(ACX4000 Universal Access Router\) on page 968](#)

**Output Fields** [Table 106 on page 913](#) lists the output fields for the **show chassis environment** command. Output fields are listed in the approximate order in which they appear.



Table 106: show chassis environment Output Fields

Field Name	Field Description
<b>Class</b>	<p>Information about the category or class of chassis component:</p> <ul style="list-style-type: none"> <li>• <b>Power:</b> Power information: <ul style="list-style-type: none"> <li>• (M5, M10, M20, and M40 routers and EX Series switches only) Power supply status: <b>OK</b>, <b>Testing</b>, (during initial power-on), <b>Failed</b>, or <b>Absent</b>.</li> <li>• (M7i, M10i, M40e, M120, M160, M320, and T Series routers and EX Series switches only) Power Entry Modules status: <b>OK</b>, <b>Testing</b>, (during initial power-on), <b>Check</b>, <b>Failed</b>, or <b>Absent</b>.</li> <li>• (PTX Series only) Power information is reported in PDU or PSM combinations. The status is: <b>OK</b>, <b>Testing</b>, (during initial power-on), <b>Check</b>, <b>Failed</b>, or <b>Absent</b>.</li> </ul> </li> <li>• <b>Temp:</b> Temperature of air flowing through the chassis in degrees Celsius (C) and Fahrenheit (F). <ul style="list-style-type: none"> <li>• On PTX Series Packet Transport Routers and MX2010 and MX2020 Routers, multiple cooling zones are supported. FRU temperatures in each zone are coordinated with the fan speed of fan trays in those zones.</li> <li>• EX2200 switches have a side-to-rear cooling system. The <b>Local Intake</b> temperature is measured by the sensor on the right side of the chassis, and the <b>Remote Intake</b> temperature is measured by the sensor on the left side of the chassis.</li> </ul> </li> <li>• <b>Pic:</b> On ACX4000 Routers, multiple temperature channels on a MIC. The status is: <b>OK</b> and the <b>Measurement</b> is in degrees Celsius (C) and Fahrenheit (F).</li> <li>• <b>Fan:</b> Fan status: <b>OK</b>, <b>Testing</b> (during initial power-on), <b>Failed</b>, or <b>Absent</b>. On PTX Series Packet Transport Routers and MX2010 and MX2020 Routers, multiple fan trays are supported. Fan status is reported in Fan Tray or Fan combinations. <b>Measurement</b> indicates actual fan RPM (PTX and MX2010 and MX2020 Routers only).</li> <li>• <b>Misc:</b> Information about other components of the chassis. <ul style="list-style-type: none"> <li>• On some routers, this field indicates the status of one or more additional components.</li> <li>• On the M40e, M160, and M320 router, <b>Misc</b> includes <b>CIP</b> (Connector Interface Panel). <b>OK</b> indicates that the CIP is present. <b>Absent</b> indicates that the CIP is not present.</li> <li>• On T Series routers, <b>Misc</b> includes <b>CIP</b> and <b>SPMB</b> (Switch Processor Mezzanine Board). <b>OK</b> indicates that the <b>CIP</b> or <b>SPMB</b> is present. <b>Absent</b> indicates that the <b>CIP</b> or <b>SPMB</b> is not present.</li> <li>• On PTX Series Packet Transport Routers, <b>Misc</b> includes the <b>SPMB</b> (Switch Processor Mezzanine Board). The SPMB is located on the control boards. <b>OK</b> indicates that the control board is present. <b>Absent</b> indicates that the control board is not present.</li> </ul> </li> </ul>
<b>Item</b>	<p>(MX2010 and MX2020 Routers) Information about the chassis component: Routing Engines, Controls Boards (CBs), Switch Fabric Boards (SFBs), PICs, Flexible PIC Concentrators (FPCs), and Adapter Cards (ADCs).</p> <p>(MX104 Routers) Information about the chassis components: Routing Engines, Control Board (CB), Power Entry Module (PEM), and Compact Forwarding Engine Board (AFEB).</p> <p>(QFabric Systems) Information about the chassis component: Control Boards, Routing Engines, Flexible PIC Concentrators (FPCs), and Power Entry Modules (PEMs), Node Devices, and Interconnect Devices.</p> <p>(QFX Series) Information about the chassis component: Flexible PIC Concentrators (FPCs), and Power Entry Modules (PEMs).</p>

Table 106: show chassis environment Output Fields (*continued*)

Field Name	Field Description
<b>Status</b>	<p>(MX104, MX2010, and MX2020 Routers) Status of the specified chassis component. For example, if the Class is Fan, the fan status can be:</p> <ul style="list-style-type: none"> <li>• <b>OK:</b> The fans are operational.</li> <li>• <b>Testing:</b> The fans are being tested during initial power-on.</li> <li>• <b>Failed:</b> The fans have failed or the fans are not spinning.</li> <li>• <b>Absent:</b> The fan tray is not installed.</li> </ul> <p>If the Class is Power, the power supply status can be:</p> <ul style="list-style-type: none"> <li>• <b>OK:</b> The power component is operational.</li> <li>• <b>Testing:</b> The power component is being tested during initial power-on.</li> <li>• <b>Check:</b> There is insufficient power---that is, fewer than the minimum required feeds are connected.</li> <li>• <b>Failed:</b> The inputs leads have failed.</li> <li>• <b>Absent:</b> The power component is not installed.</li> </ul>
<b>Measurement</b>	<p>(MX104, MX2010, and MX2020 Routers) Dependant on the Class. For example, if the Class is Temp, indicates the temperature in degree Celsius and degrees Fahrenheit. If the Class is Fan, indicates actual fan RPM.</p>

## Sample Output

### show chassis environment (J2300 Router)

```

user@host> show chassis environment
Class Item           Status Measurement
Temp  Routing Engine    OK      40 degrees C / 104 degrees F
Fan   Fan              OK

```

### show chassis environment (J4300 or J6300 Router)

```

user@host> show chassis environment
Class Item           Status Measurement
Temp  Routing Engine    OK      41 degrees C / 105 degrees F
Fan   Fan 0             OK
      Fan 1             OK

```

### show chassis environment (M5 Router)

```

user@host> show chassis environment
Class Item           Status Measurement
Power Power Supply A    OK
      Power Supply B    Absent
Temp  FPC 0             OK      30 degrees C / 86 degrees F
      FEB              OK      33 degrees C / 91 degrees F
      PS Intake         OK      27 degrees C / 80 degrees F
      PS Exhaust        OK      27 degrees C / 80 degrees F
      Routing Engine    OK      34 degrees C / 93 degrees F
Fans  Left Fan 1        OK      Spinning at normal speed
      Left Fan 2        OK      Spinning at normal speed
      Left Fan 3        OK      Spinning at normal speed
      Left Fan 4        OK      Spinning at normal speed
Misc  Craft Interface    OK

```

## show chassis environment (M7i Router)

```

user@host> show chassis environment
Class Item                Status      Measurement
Power Power Supply 0        OK
      Power Supply 1      Absent
Temp  Intake               OK          22 degrees C / 71 degrees F
      FPC 0                OK          23 degrees C / 73 degrees F
      Power Supplies       OK          23 degrees C / 73 degrees F
      CFEB Intake          OK          24 degrees C / 75 degrees F
      CFEB Exhaust         OK          29 degrees C / 84 degrees F
      Routing Engine       OK          26 degrees C / 78 degrees F
Fans  Fan 1                 OK          Spinning at normal speed
      Fan 2                 OK          Spinning at normal speed
      Fan 3                 OK          Spinning at normal speed
      Fan 4                 OK          Spinning at normal speed

```

## show chassis environment (M10 Router)

```

user@host> show chassis environment
Class Item                Status      Measurement
Power Power Supply A        OK
      Power Supply B      Failed
Temp  FPC 0                 OK          36 degrees C / 96 degrees F
      FPC 1                 OK          35 degrees C / 95 degrees F
      FEB                   OK          34 degrees C / 93 degrees F
      PS Intake             OK          31 degrees C / 87 degrees F
      PS Exhaust            OK          34 degrees C / 93 degrees F
      Routing Engine        OK          35 degrees C / 95 degrees F
Fans  Left Fan 1            OK          Spinning at normal speed
      Left Fan 2            OK          Spinning at normal speed
      Left Fan 3            OK          Spinning at normal speed
      Left Fan 4            OK          Spinning at normal speed
Misc  Craft Interface       OK

```

## show chassis environment (M10i Router)

```

user@host> show chassis environment
Class Item                Status      Measurement
Power Power Supply 0        OK
      Power Supply 1      OK
      Power Supply 2      Absent
      Power Supply 3      Absent
Temp  Intake               OK          26 degrees C / 78 degrees F
      FPC 0                OK          27 degrees C / 80 degrees F
      FPC 1                OK          28 degrees C / 82 degrees F
      Lower Power Supplies OK          29 degrees C / 84 degrees F
      Upper Power Supplies OK          28 degrees C / 82 degrees F
      CFEB Intake          OK          27 degrees C / 80 degrees F
      CFEB Exhaust         OK          36 degrees C / 96 degrees F
      Routing Engine 0     OK          31 degrees C / 87 degrees F
      Routing Engine 1     OK          27 degrees C / 80 degrees F
Fans  Fan Tray 0 Fan 1     OK          Spinning at normal speed
      Fan Tray 0 Fan 2     OK          Spinning at normal speed
      Fan Tray 0 Fan 3     OK          Spinning at normal speed
      Fan Tray 0 Fan 4     OK          Spinning at normal speed
      Fan Tray 0 Fan 5     OK          Spinning at normal speed
      Fan Tray 0 Fan 6     OK          Spinning at normal speed
      Fan Tray 0 Fan 7     OK          Spinning at normal speed

```

Fan Tray 0 Fan 8	OK	Spinning at normal speed
Fan Tray 1 Fan 1	Absent	
Fan Tray 1 Fan 2	Absent	
Fan Tray 1 Fan 3	Absent	
Fan Tray 1 Fan 4	Absent	
Fan Tray 1 Fan 5	Absent	
Fan Tray 1 Fan 6	Absent	
Fan Tray 1 Fan 7	Absent	
Fan Tray 1 Fan 8	Absent	

**show chassis environment (M20 Router)**

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply A	OK	
	Power Supply B	Absent	
Temp	FPC 0	OK	28 degrees C / 82 degrees F
	FPC 1	OK	27 degrees C / 80 degrees F
	Power Supply A	OK	22 degrees C / 71 degrees F
	Power Supply B	Absent	
	SSB 0	OK	30 degrees C / 86 degrees F
	Backplane	OK	22 degrees C / 71 degrees F
	Routing Engine 0	OK	26 degrees C / 78 degrees F
	Routing Engine 1	Testing	
Fans	Rear Fan	OK	Spinning at normal speed
	Front Upper Fan	OK	Spinning at normal speed
	Front Middle Fan	OK	Spinning at normal speed
	Front Bottom Fan	OK	Spinning at normal speed
Misc	Craft Interface	OK	

**show chassis environment (M40 Router)**

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply A	OK	
	Power Supply B	Absent	
Temp	FPC 3	OK	24 degrees C / 75 degrees F
	FPC 6	OK	26 degrees C / 78 degrees F
	SCB	OK	26 degrees C / 78 degrees F
	Backplane @ A1	OK	28 degrees C / 82 degrees F
	Backplane @ A2	OK	23 degrees C / 73 degrees F
	Routing Engine	OK	26 degrees C / 78 degrees F
Fans	Top Impeller	OK	Spinning at normal speed
	Bottom impeller	OK	Spinning at normal speed
	Rear Left Fan	OK	Spinning at normal speed
	Rear Center Fan	OK	Spinning at normal speed
	Rear Right Fan	OK	Spinning at normal speed
Misc	Craft Interface	OK	

**show chassis environment (M40e Router)**

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	
	PEM 1	Absent	
Temp	PCG 0	OK	44 degrees C / 111 degrees F
	PCG 1	OK	47 degrees C / 116 degrees F
	Routing Engine 0	OK	40 degrees C / 104 degrees F
	Routing Engine 1	OK	37 degrees C / 98 degrees F

	MCS 0	OK	45 degrees C / 113 degrees F
	MCS 1	OK	42 degrees C / 107 degrees F
	SFM 0 SPP	OK	40 degrees C / 104 degrees F
	SFM 0 SPR	OK	44 degrees C / 111 degrees F
	SFM 1 SPP	OK	43 degrees C / 109 degrees F
	SFM 1 SPR	OK	45 degrees C / 113 degrees F
	FPC 0	OK	38 degrees C / 100 degrees F
	FPC 1	OK	40 degrees C / 104 degrees F
	FPC 2	OK	38 degrees C / 100 degrees F
	FPC 4	OK	34 degrees C / 93 degrees F
	FPC 5	OK	43 degrees C / 109 degrees F
	FPC 6	OK	41 degrees C / 105 degrees F
	FPC 7	OK	43 degrees C / 109 degrees F
	FPM CMB	OK	28 degrees C / 82 degrees F
	FPM Display	OK	28 degrees C / 82 degrees F
Fans	Rear Bottom Blower	OK	Spinning at normal speed
	Rear Top Blower	OK	Spinning at normal speed
	Front Top Blower	OK	Spinning at normal speed
	Fan Tray Rear Left	OK	Spinning at normal speed
	Fan Tray Rear Right	OK	Spinning at normal speed
	Fan Tray Front Left	OK	Spinning at normal speed
	Fan Tray Front Right	OK	Spinning at normal speed
Misc	CIP	OK	

### show chassis environment (M120 Router)

user@host> show chassis environment

Class	Item	Status	Measurement
Temp	PEM 0	OK	
	PEM 1	OK	
	Routing Engine 0	OK	43 degrees C / 109 degrees F
	Routing Engine 1	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	33 degrees C / 91 degrees F
	CB 0 Exhaust A	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust B	OK	35 degrees C / 95 degrees F
	CB 1 Intake	OK	34 degrees C / 93 degrees F
	CB 1 Exhaust A	OK	38 degrees C / 100 degrees F
	CB 1 Exhaust B	OK	35 degrees C / 95 degrees F
	FEB 3 Intake	OK	35 degrees C / 95 degrees F
	FEB 3 Exhaust A	OK	37 degrees C / 98 degrees F
	FEB 3 Exhaust B	OK	39 degrees C / 102 degrees F
	FEB 4 Intake	OK	33 degrees C / 91 degrees F
	FEB 4 Exhaust A	OK	39 degrees C / 102 degrees F
	FEB 4 Exhaust B	OK	36 degrees C / 96 degrees F
	FPC 2 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 2 Exhaust B	OK	31 degrees C / 87 degrees F
	FPC 3 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 3 Exhaust B	OK	33 degrees C / 91 degrees F
	FPC 4 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 4 Exhaust B	OK	30 degrees C / 86 degrees F
Fans	Front Top Tray Fan 1	OK	Spinning at normal speed
	Front Top Tray Fan 2	OK	Spinning at normal speed
	Front Top Tray Fan 3	OK	Spinning at normal speed
	Front Top Tray Fan 4	OK	Spinning at normal speed
	Front Top Tray Fan 5	OK	Spinning at normal speed
	Front Top Tray Fan 6	OK	Spinning at normal speed
	Front Top Tray Fan 7	OK	Spinning at normal speed
	Front Top Tray Fan 8	OK	Spinning at normal speed
	Front Bottom Tray Fan 1	OK	Spinning at normal speed
	Front Bottom Tray Fan 2	OK	Spinning at normal speed

Front Bottom Tray Fan 3	OK	Spinning at normal speed
Front Bottom Tray Fan 4	OK	Spinning at normal speed
Front Bottom Tray Fan 5	OK	Spinning at normal speed
Front Bottom Tray Fan 6	OK	Spinning at normal speed
Front Bottom Tray Fan 7	OK	Spinning at normal speed
Front Bottom Tray Fan 8	OK	Spinning at normal speed
Rear Top Tray Fan 1	OK	Spinning at normal speed
Rear Top Tray Fan 2	OK	Spinning at normal speed
Rear Top Tray Fan 3	OK	Spinning at normal speed
Rear Top Tray Fan 4	OK	Spinning at normal speed
Rear Top Tray Fan 5	OK	Spinning at normal speed
Rear Top Tray Fan 6	OK	Spinning at normal speed
Rear Top Tray Fan 7	OK	Spinning at normal speed
Rear Top Tray Fan 8	OK	Spinning at normal speed
Rear Bottom Tray Fan 1	OK	Spinning at normal speed
Rear Bottom Tray Fan 2	OK	Spinning at normal speed
Rear Bottom Tray Fan 3	OK	Spinning at normal speed
Rear Bottom Tray Fan 4	OK	Spinning at normal speed
Rear Bottom Tray Fan 5	OK	Spinning at normal speed
Rear Bottom Tray Fan 6	OK	Spinning at normal speed
Rear Bottom Tray Fan 7	OK	Spinning at normal speed
Rear Bottom Tray Fan 8	OK	Spinning at normal speed

#### show chassis environment (M160 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	PEM 1
Temp	PCG 0	OK	45 degrees C / 113 degrees F
	PCG 1	Absent	
	Routing Engine 0	OK	35 degrees C / 95 degrees F
	Routing Engine 1	Absent	
	MCS 0	OK	50 degrees C / 122 degrees F
	SFM 0 SPP	OK	47 degrees C / 116 degrees F
	SFM 0 SPR	OK	49 degrees C / 120 degrees F
	SFM 1 SPP	OK	50 degrees C / 122 degrees F
	SFM 1 SPR	OK	50 degrees C / 122 degrees F
	SFM 2 SPP	OK	51 degrees C / 123 degrees F
	SFM 2 SPR	OK	52 degrees C / 125 degrees F
	SFM 3 SPP	OK	52 degrees C / 125 degrees F
	SFM 3 SPR	OK	48 degrees C / 118 degrees F
	FPC 0	OK	45 degrees C / 113 degrees F
	FPC 6	OK	43 degrees C / 109 degrees F
	FPM CMB	OK	31 degrees C / 87 degrees F
	FPM Display	OK	33 degrees C / 91 degrees F
Fans	Rear Bottom Blower	OK	Spinning at normal speed
	Rear Top Blower	OK	Spinning at normal speed
	Front Top Blower	OK	Spinning at normal speed
	Fan Tray Rear Left	OK	Spinning at normal speed
	Fan Tray Rear Right	OK	Spinning at normal speed
	Fan Tray Front Left	OK	Spinning at normal speed
	Fan Tray Front Right	OK	Spinning at normal speed
Misc	CIP	OK	

#### show chassis environment (M320 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	Absent	

	PEM 2	OK	
	PEM 3	OK	
	Routing Engine 0	OK	33 degrees C / 91 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	CB 0	OK	36 degrees C / 96 degrees F
	CB 1	OK	36 degrees C / 96 degrees F
	SIB 0	OK	38 degrees C / 100 degrees F
	SIB 1	OK	29 degrees C / 84 degrees F
	SIB 2	OK	38 degrees C / 100 degrees F
	SIB 3	OK	41 degrees C / 105 degrees F
	FPC 0 Intake	OK	28 degrees C / 82 degrees F
	FPC 0 Exhaust	OK	40 degrees C / 104 degrees F
	FPC 1 Intake	OK	29 degrees C / 84 degrees F
	FPC 1 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 2 Intake	OK	28 degrees C / 82 degrees F
	FPC 2 Exhaust	OK	38 degrees C / 100 degrees F
	FPC 3 Intake	OK	28 degrees C / 82 degrees F
	FPC 3 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 6 Intake	OK	27 degrees C / 80 degrees F
	FPC 6 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 7 Intake	OK	27 degrees C / 80 degrees F
	FPC 7 Exhaust	OK	42 degrees C / 107 degrees F
	FPM GBUS	OK	30 degrees C / 86 degrees F
Fan	Top Left Front fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Rear Fan 1 (TOP)	OK	Spinning at normal speed
	Rear Fan 2	OK	Spinning at normal speed
	Rear Fan 3	OK	Spinning at normal speed
	Rear Fan 4	OK	Spinning at normal speed
	Rear Fan 5	OK	Spinning at normal speed
	Rear Fan 6	OK	Spinning at normal speed
	Rear Fan 7 (Bottom)	OK	Spinning at normal speed
Misc	CIP	OK	

### show chassis environment (MX104 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	34 degrees C / 93 degrees F
	PEM 1	Absent	
	ABB 0 Intake	OK	33 degrees C / 91 degrees F
	ABB 0 Exhaust A	OK	42 degrees C / 107 degrees F
	ABB 0 Exhaust B	OK	43 degrees C / 109 degrees F
	ABB 1 Intake	Absent	
	ABB 1 Exhaust A	Absent	
	ABB 1 Exhaust B	Absent	
	Routing Engine 0	OK	34 degrees C / 93 degrees F
	Routing Engine 0 CPU	OK	46 degrees C / 114 degrees F
Fans	Routing Engine 1	Absent	
	Routing Engine 1 CPU	Absent	
	AFEB 0 AFEB Processor	OK	33 degrees C / 91 degrees F
	Fan 1	OK	Spinning at normal speed
	Fan 2	OK	Spinning at normal speed
	Fan 3	OK	Spinning at normal speed

Fan 4	OK	Spinning at normal speed
Fan 5	OK	Spinning at normal speed

**show chassis environment (MX240 Router)**

user@host&gt; show chassis environment

Class	Item	Status	Measurement
Temp	PEM 0	OK	40 degrees C / 104 degrees F
	PEM 1	OK	45 degrees C / 113 degrees F
	PEM 2	Absent	
	PEM 3	Absent	
	Routing Engine 0	OK	39 degrees C / 102 degrees F
	Routing Engine 1	OK	37 degrees C / 98 degrees F
	CB 0 Intake	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK	34 degrees C / 93 degrees F
	CB 0 Exhaust B	OK	38 degrees C / 100 degrees F
	CB 0 ACBC	OK	37 degrees C / 98 degrees F
	CB 0 SF A	OK	49 degrees C / 120 degrees F
	CB 0 SF B	OK	41 degrees C / 105 degrees F
	CB 1 Intake	OK	37 degrees C / 98 degrees F
	CB 1 Exhaust A	OK	34 degrees C / 93 degrees F
	CB 1 Exhaust B	OK	39 degrees C / 102 degrees F
	CB 1 ACBC	OK	38 degrees C / 100 degrees F
	CB 1 SF A	OK	47 degrees C / 116 degrees F
	CB 1 SF B	OK	41 degrees C / 105 degrees F
	FPC 1 Intake	OK	33 degrees C / 91 degrees F
	FPC 1 Exhaust A	OK	38 degrees C / 100 degrees F
	FPC 1 Exhaust B	OK	53 degrees C / 127 degrees F
	FPC 1 I3 0 TSensor	OK	50 degrees C / 122 degrees F
	FPC 1 I3 0 Chip	OK	53 degrees C / 127 degrees F
	FPC 1 I3 1 TSensor	OK	49 degrees C / 120 degrees F
	FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 1 I3 2 TSensor	OK	47 degrees C / 116 degrees F
	FPC 1 I3 2 Chip	OK	49 degrees C / 120 degrees F
	FPC 1 I3 3 TSensor	OK	44 degrees C / 111 degrees F
	FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
	FPC 1 IA 0 TSensor	OK	45 degrees C / 113 degrees F
	FPC 1 IA 0 Chip	OK	44 degrees C / 111 degrees F
	FPC 1 IA 1 TSensor	OK	44 degrees C / 111 degrees F
	FPC 1 IA 1 Chip	OK	48 degrees C / 118 degrees F
	FPC 2 Intake	OK	32 degrees C / 89 degrees F
	FPC 2 Exhaust A	OK	40 degrees C / 104 degrees F
	FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
	FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
	FPC 2 I3 0 Chip	OK	56 degrees C / 132 degrees F
	FPC 2 I3 1 TSensor	OK	52 degrees C / 125 degrees F
	FPC 2 I3 1 Chip	OK	55 degrees C / 131 degrees F
	FPC 2 I3 2 TSensor	OK	49 degrees C / 120 degrees F
	FPC 2 I3 2 Chip	OK	52 degrees C / 125 degrees F
	FPC 2 I3 3 TSensor	OK	44 degrees C / 111 degrees F
	FPC 2 I3 3 Chip	OK	48 degrees C / 118 degrees F
	FPC 2 IA 0 TSensor	OK	50 degrees C / 122 degrees F
	FPC 2 IA 0 Chip	OK	48 degrees C / 118 degrees F
	FPC 2 IA 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 2 IA 1 Chip	OK	53 degrees C / 127 degrees F
Fans	Front Fan	OK	Spinning at normal speed
	Middle Fan	OK	Spinning at normal speed
	Rear Fan	OK	Spinning at normal speed



## show chassis environment (MX240 Router with SCBE)

```

user@host> show chassis environment
Class Item                               Status      Measurement
Temp PEM 0                             OK          40 degrees C / 104 degrees F
      PEM 1                             OK          45 degrees C / 113 degrees F
      PEM 2                             Absent
      PEM 3                             Absent
      Routing Engine 0                   OK          39 degrees C / 102 degrees F
      Routing Engine 1                   OK          37 degrees C / 98 degrees F
      CB 0 Intake                        OK          36 degrees C / 96 degrees F
      CB 0 Exhaust A                     OK          34 degrees C / 93 degrees F
      CB 0 Exhaust B                     OK          38 degrees C / 100 degrees F
      CB 0 ACBC                          OK          37 degrees C / 98 degrees F
      CB 0 XF A                           OK          49 degrees C / 120 degrees F
      CB 0 XF B                           OK          41 degrees C / 105 degrees F
      CB 1 Intake                        OK          37 degrees C / 98 degrees F
      CB 1 Exhaust A                     OK          34 degrees C / 93 degrees F
      CB 1 Exhaust B                     OK          39 degrees C / 102 degrees F
      CB 1 ACBC                          OK          38 degrees C / 100 degrees F
      CB 1 XF A                           OK          47 degrees C / 116 degrees F
      CB 1 XF B                           OK          41 degrees C / 105 degrees F
      FPC 1 Intake                       OK          33 degrees C / 91 degrees F
      FPC 1 Exhaust A                     OK          38 degrees C / 100 degrees F
      FPC 1 Exhaust B                     OK          53 degrees C / 127 degrees F
      FPC 1 I3 0 TSensor                  OK          50 degrees C / 122 degrees F
      FPC 1 I3 0 Chip                     OK          53 degrees C / 127 degrees F
      FPC 1 I3 1 TSensor                  OK          49 degrees C / 120 degrees F
      FPC 1 I3 1 Chip                     OK          52 degrees C / 125 degrees F
      FPC 1 I3 2 TSensor                  OK          47 degrees C / 116 degrees F
      FPC 1 I3 2 Chip                     OK          49 degrees C / 120 degrees F
      FPC 1 I3 3 TSensor                  OK          44 degrees C / 111 degrees F
      FPC 1 I3 3 Chip                     OK          46 degrees C / 114 degrees F
      FPC 1 IA 0 TSensor                  OK          45 degrees C / 113 degrees F
      FPC 1 IA 0 Chip                     OK          44 degrees C / 111 degrees F
      FPC 1 IA 1 TSensor                  OK          44 degrees C / 111 degrees F
      FPC 1 IA 1 Chip                     OK          48 degrees C / 118 degrees F
      FPC 2 Intake                       OK          32 degrees C / 89 degrees F
      FPC 2 Exhaust A                     OK          40 degrees C / 104 degrees F
      FPC 2 Exhaust B                     OK          52 degrees C / 125 degrees F
      FPC 2 I3 0 TSensor                  OK          52 degrees C / 125 degrees F
      FPC 2 I3 0 Chip                     OK          56 degrees C / 132 degrees F
      FPC 2 I3 1 TSensor                  OK          52 degrees C / 125 degrees F
      FPC 2 I3 1 Chip                     OK          55 degrees C / 131 degrees F
      FPC 2 I3 2 TSensor                  OK          49 degrees C / 120 degrees F
      FPC 2 I3 2 Chip                     OK          52 degrees C / 125 degrees F
      FPC 2 I3 3 TSensor                  OK          44 degrees C / 111 degrees F
      FPC 2 I3 3 Chip                     OK          48 degrees C / 118 degrees F
      FPC 2 IA 0 TSensor                  OK          50 degrees C / 122 degrees F
      FPC 2 IA 0 Chip                     OK          48 degrees C / 118 degrees F
      FPC 2 IA 1 TSensor                  OK          47 degrees C / 116 degrees F
      FPC 2 IA 1 Chip                     OK          53 degrees C / 127 degrees F
Fans  Front Fan                         OK          Spinning at normal speed
      Middle Fan                         OK          Spinning at normal speed
      Rear Fan                           OK          Spinning at normal speed

```

## show chassis environment (MX480 Router)

```

user@host> show chassis environment
Class Item                               Status      Measurement
Temp PEM 0                             OK          35 degrees C / 95 degrees F

```

PEM 1	OK	40 degrees C / 104 degrees F
PEM 2	Absent	
PEM 3	Absent	
Routing Engine 0	OK	44 degrees C / 111 degrees F
Routing Engine 1	OK	45 degrees C / 113 degrees F
CB 0 Intake	OK	36 degrees C / 96 degrees F
CB 0 Exhaust A	OK	38 degrees C / 100 degrees F
CB 0 Exhaust B	OK	39 degrees C / 102 degrees F
CB 0 ACBC	OK	37 degrees C / 98 degrees F
CB 0 SF A	OK	51 degrees C / 123 degrees F
CB 0 SF B	OK	44 degrees C / 111 degrees F
CB 1 Intake	OK	36 degrees C / 96 degrees F
CB 1 Exhaust A	OK	39 degrees C / 102 degrees F
CB 1 Exhaust B	OK	40 degrees C / 104 degrees F
CB 1 ACBC	OK	37 degrees C / 98 degrees F
CB 1 SF A	OK	50 degrees C / 122 degrees F
CB 1 SF B	OK	43 degrees C / 109 degrees F
FPC 0 Intake	OK	36 degrees C / 96 degrees F
FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 0 Exhaust B	OK	51 degrees C / 123 degrees F
FPC 0 I3 0 TSensor	OK	49 degrees C / 120 degrees F
FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 0 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 0 I3 2 TSensor	OK	46 degrees C / 114 degrees F
FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
FPC 0 I3 3 TSensor	OK	42 degrees C / 107 degrees F
FPC 0 I3 3 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 1 Intake	OK	37 degrees C / 98 degrees F
FPC 1 Exhaust A	OK	41 degrees C / 105 degrees F
FPC 1 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 1 I3 0 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 I3 0 Chip	OK	57 degrees C / 134 degrees F
FPC 1 I3 1 TSensor	OK	48 degrees C / 118 degrees F
FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 1 I3 2 TSensor	OK	46 degrees C / 114 degrees F
FPC 1 I3 2 Chip	OK	50 degrees C / 122 degrees F
FPC 1 I3 3 TSensor	OK	42 degrees C / 107 degrees F
FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
FPC 1 IA 0 TSensor	OK	49 degrees C / 120 degrees F
FPC 1 IA 0 Chip	OK	48 degrees C / 118 degrees F
FPC 1 IA 1 TSensor	OK	46 degrees C / 114 degrees F
FPC 1 IA 1 Chip	OK	50 degrees C / 122 degrees F
Fans Top Rear Fan	OK	Spinning at normal speed
Bottom Rear Fan	OK	Spinning at normal speed
Top Middle Fan	OK	Spinning at normal speed
Bottom Middle Fan	OK	Spinning at normal speed
Top Front Fan	OK	Spinning at normal speed
Bottom Front Fan	OK	Spinning at normal speed

#### show chassis environment (MX480 Router with SCBE)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	OK	35 degrees C / 95 degrees F
	PEM 1	OK	40 degrees C / 104 degrees F
	PEM 2	Absent	

PEM 3	Absent	
Routing Engine 0	OK	44 degrees C / 111 degrees F
Routing Engine 1	OK	45 degrees C / 113 degrees F
CB 0 Intake	OK	36 degrees C / 96 degrees F
CB 0 Exhaust A	OK	38 degrees C / 100 degrees F
CB 0 Exhaust B	OK	39 degrees C / 102 degrees F
CB 0 ACBC	OK	37 degrees C / 98 degrees F
CB 0 XF A	OK	51 degrees C / 123 degrees F
CB 0 XF B	OK	44 degrees C / 111 degrees F
CB 1 Intake	OK	36 degrees C / 96 degrees F
CB 1 Exhaust A	OK	39 degrees C / 102 degrees F
CB 1 Exhaust B	OK	40 degrees C / 104 degrees F
CB 1 ACBC	OK	37 degrees C / 98 degrees F
CB 1 XF A	OK	50 degrees C / 122 degrees F
CB 1 XF B	OK	43 degrees C / 109 degrees F
FPC 0 Intake	OK	36 degrees C / 96 degrees F
FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 0 Exhaust B	OK	51 degrees C / 123 degrees F
FPC 0 I3 0 TSensor	OK	49 degrees C / 120 degrees F
FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 0 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 0 I3 2 TSensor	OK	46 degrees C / 114 degrees F
FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
FPC 0 I3 3 TSensor	OK	42 degrees C / 107 degrees F
FPC 0 I3 3 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 1 Intake	OK	37 degrees C / 98 degrees F
FPC 1 Exhaust A	OK	41 degrees C / 105 degrees F
FPC 1 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 1 I3 0 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 I3 0 Chip	OK	57 degrees C / 134 degrees F
FPC 1 I3 1 TSensor	OK	48 degrees C / 118 degrees F
FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 1 I3 2 TSensor	OK	46 degrees C / 114 degrees F
FPC 1 I3 2 Chip	OK	50 degrees C / 122 degrees F
FPC 1 I3 3 TSensor	OK	42 degrees C / 107 degrees F
FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
FPC 1 IA 0 TSensor	OK	49 degrees C / 120 degrees F
FPC 1 IA 0 Chip	OK	48 degrees C / 118 degrees F
FPC 1 IA 1 TSensor	OK	46 degrees C / 114 degrees F
FPC 1 IA 1 Chip	OK	50 degrees C / 122 degrees F
Fans		
Top Rear Fan	OK	Spinning at normal speed
Bottom Rear Fan	OK	Spinning at normal speed
Top Middle Fan	OK	Spinning at normal speed
Bottom Middle Fan	OK	Spinning at normal speed
Top Front Fan	OK	Spinning at normal speed
Bottom Front Fan	OK	Spinning at normal speed

### show chassis environment (MX960 Router)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	Absent	
	PEM 2	Check	
	PEM 3	OK	35 degrees C / 95 degrees F
	Routing Engine 0	OK	37 degrees C / 98 degrees F

	Routing Engine 1	Absent	
	CB 0 Intake	OK	24 degrees C / 75 degrees F
	CB 0 Exhaust A	OK	30 degrees C / 86 degrees F
	CB 0 Exhaust B	OK	27 degrees C / 80 degrees F
	CB 1 Intake	Absent	
	CB 1 Exhaust A	Absent	
	CB 1 Exhaust B	Absent	
	CB 1 ACBC	Absent	
	CB 1 SF A	Absent	
	CB 1 SF B	Absent	
	CB 2 Intake	Absent	
	CB 2 Exhaust A	Absent	
	CB 2 Exhaust B	Absent	
	CB 2 ACBC	Absent	
	CB 2 SF A	Absent	
	CB 2 SF B	Absent	
	FPC 4 Intake	OK	24 degrees C / 75 degrees F
	FPC 4 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 4 Exhaust B	OK	38 degrees C / 100 degrees F
	FPC 7 Intake	OK	24 degrees C / 75 degrees F
	FPC 7 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 7 Exhaust B	OK	42 degrees C / 107 degrees F
Fans	Top Fan Tray Temp	Failed	
	Top Tray Fan 1	OK	Spinning at normal speed
	Top Tray Fan 2	OK	Spinning at normal speed
	Top Tray Fan 3	OK	Spinning at normal speed
	Top Tray Fan 4	OK	Spinning at normal speed
	Top Tray Fan 5	OK	Spinning at normal speed
	Top Tray Fan 6	OK	Spinning at normal speed
	Bottom Fan Tray Temp	Failed	
	Bottom Tray Fan 1	OK	Spinning at normal speed
	Bottom Tray Fan 2	OK	Spinning at normal speed
	Bottom Tray Fan 3	OK	Spinning at normal speed
	Bottom Tray Fan 4	OK	Spinning at normal speed
	Bottom Tray Fan 5	OK	Spinning at normal speed
	Bottom Tray Fan 6	OK	Spinning at normal speed

### show chassis environment (MX960 Router with SCBE)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	50 degrees C / 122 degrees F
	PEM 2	OK	50 degrees C / 122 degrees F
	PEM 3	OK	50 degrees C / 122 degrees F
	Routing Engine 0	OK	42 degrees C / 107 degrees F
	Routing Engine 0 CPU	OK	51 degrees C / 123 degrees F
	Routing Engine 1	OK	39 degrees C / 102 degrees F
	Routing Engine 1 CPU	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	35 degrees C / 95 degrees F
	CB 0 Exhaust A	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust B	OK	43 degrees C / 109 degrees F
	CB 0 ACBC	OK	38 degrees C / 100 degrees F
	CB 0 XF A	OK	53 degrees C / 127 degrees F
	CB 0 XF B	OK	47 degrees C / 116 degrees F
	CB 1 Intake	OK	35 degrees C / 95 degrees F
	CB 1 Exhaust A	OK	35 degrees C / 95 degrees F
	CB 1 Exhaust B	OK	41 degrees C / 105 degrees F
	CB 1 ACBC	OK	38 degrees C / 100 degrees F
	CB 1 XF A	OK	52 degrees C / 125 degrees F
	CB 1 XF B	OK	47 degrees C / 116 degrees F

CB 2 Intake	OK	32 degrees C / 89 degrees F
CB 2 Exhaust A	OK	30 degrees C / 86 degrees F
CB 2 Exhaust B	OK	35 degrees C / 95 degrees F
CB 2 ACBC	OK	33 degrees C / 91 degrees F
CB 2 XF A	OK	51 degrees C / 123 degrees F
CB 2 XF B	OK	50 degrees C / 122 degrees F
FPC 0 Intake	OK	35 degrees C / 95 degrees F
FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 0 Exhaust B	OK	50 degrees C / 122 degrees F
FPC 0 I3 0 TSensor	OK	50 degrees C / 122 degrees F
FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 0 I3 1 Chip	OK	50 degrees C / 122 degrees F
FPC 0 I3 2 TSensor	OK	45 degrees C / 113 degrees F
FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
FPC 0 I3 3 TSensor	OK	41 degrees C / 105 degrees F
FPC 0 I3 3 Chip	OK	44 degrees C / 111 degrees F
FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 1 Intake	OK	36 degrees C / 96 degrees F
FPC 1 Exhaust A	OK	47 degrees C / 116 degrees F
FPC 1 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 1 LU 0 TCAM TSensor	OK	53 degrees C / 127 degrees F
FPC 1 LU 0 TCAM Chip	OK	57 degrees C / 134 degrees F
FPC 1 LU 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 1 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 1 MQ 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 1 MQ 0 Chip	OK	56 degrees C / 132 degrees F
FPC 1 LU 1 TCAM TSensor	OK	51 degrees C / 123 degrees F
FPC 1 LU 1 TCAM Chip	OK	52 degrees C / 125 degrees F
FPC 1 LU 1 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 LU 1 Chip	OK	53 degrees C / 127 degrees F
FPC 1 MQ 1 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 MQ 1 Chip	OK	58 degrees C / 136 degrees F
FPC 2 Intake	OK	35 degrees C / 95 degrees F
FPC 2 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 2 Exhaust B	OK	54 degrees C / 129 degrees F
FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 Chip	OK	59 degrees C / 138 degrees F
FPC 2 I3 1 TSensor	OK	48 degrees C / 118 degrees F
FPC 2 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 2 I3 2 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 I3 2 Chip	OK	49 degrees C / 120 degrees F
FPC 2 I3 3 TSensor	OK	41 degrees C / 105 degrees F
FPC 2 I3 3 Chip	OK	44 degrees C / 111 degrees F
FPC 2 IA 0 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 IA 0 Chip	OK	46 degrees C / 114 degrees F
FPC 2 IA 1 TSensor	OK	45 degrees C / 113 degrees F
FPC 2 IA 1 Chip	OK	49 degrees C / 120 degrees F
FPC 3 Intake	OK	34 degrees C / 93 degrees F
FPC 3 Exhaust A	OK	34 degrees C / 93 degrees F
FPC 3 Exhaust B	OK	47 degrees C / 116 degrees F
FPC 3 I3 0 TSensor	OK	48 degrees C / 118 degrees F
FPC 3 I3 0 Chip	OK	52 degrees C / 125 degrees F
FPC 3 I3 1 TSensor	OK	46 degrees C / 114 degrees F
FPC 3 I3 1 Chip	OK	48 degrees C / 118 degrees F
FPC 3 IA 0 TSensor	OK	41 degrees C / 105 degrees F
FPC 3 IA 0 Chip	OK	40 degrees C / 104 degrees F
FPC 5 Intake	OK	42 degrees C / 107 degrees F

	FPC 5 Exhaust A	OK	42 degrees C / 107 degrees F
	FPC 5 Exhaust B	OK	53 degrees C / 127 degrees F
	FPC 5 LU 0 TSensor	OK	53 degrees C / 127 degrees F
	FPC 5 LU 0 Chip	OK	54 degrees C / 129 degrees F
	FPC 5 LU 1 TSensor	OK	53 degrees C / 127 degrees F
	FPC 5 LU 1 Chip	OK	61 degrees C / 141 degrees F
	FPC 5 LU 2 TSensor	OK	53 degrees C / 127 degrees F
	FPC 5 LU 2 Chip	OK	51 degrees C / 123 degrees F
	FPC 5 LU 3 TSensor	OK	53 degrees C / 127 degrees F
	FPC 5 LU 3 Chip	OK	53 degrees C / 127 degrees F
	FPC 5 MQ 0 TSensor	OK	47 degrees C / 116 degrees F
	FPC 5 MQ 0 Chip	OK	52 degrees C / 125 degrees F
	FPC 5 MQ 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 5 MQ 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 5 MQ 2 TSensor	OK	47 degrees C / 116 degrees F
	FPC 5 MQ 2 Chip	OK	46 degrees C / 114 degrees F
	FPC 5 MQ 3 TSensor	OK	47 degrees C / 116 degrees F
	FPC 5 MQ 3 Chip	OK	45 degrees C / 113 degrees F
	FPC 7 Intake	OK	36 degrees C / 96 degrees F
	FPC 7 Exhaust A	OK	35 degrees C / 95 degrees F
	FPC 7 Exhaust B	OK	33 degrees C / 91 degrees F
	FPC 7 QX 0 TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 QX 0 Chip	OK	47 degrees C / 116 degrees F
	FPC 7 LU 0 TCAM TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 LU 0 TCAM Chip	OK	44 degrees C / 111 degrees F
	FPC 7 LU 0 TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 LU 0 Chip	OK	46 degrees C / 114 degrees F
	FPC 7 MQ 0 TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 MQ 0 Chip	OK	45 degrees C / 113 degrees F
	FPC 8 Intake	OK	33 degrees C / 91 degrees F
	FPC 8 Exhaust A	OK	33 degrees C / 91 degrees F
	FPC 8 Exhaust B	OK	36 degrees C / 96 degrees F
	FPC 8 I3 0 TSensor	OK	38 degrees C / 100 degrees F
	FPC 8 I3 0 Chip	OK	43 degrees C / 109 degrees F
	FPC 8 BDS 0 TSensor	OK	37 degrees C / 98 degrees F
	FPC 8 BDS 0 Chip	OK	36 degrees C / 96 degrees F
	FPC 8 IA 0 TSensor	OK	37 degrees C / 98 degrees F
	FPC 8 IA 0 Chip	OK	37 degrees C / 98 degrees F
	FPC 10 Intake	OK	38 degrees C / 100 degrees F
	FPC 10 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 10 Exhaust B	OK	41 degrees C / 105 degrees F
	FPC 10 I3 0 TSensor	OK	40 degrees C / 104 degrees F
	FPC 10 I3 0 Chip	OK	42 degrees C / 107 degrees F
	FPC 10 I3 1 TSensor	OK	40 degrees C / 104 degrees F
	FPC 10 I3 1 Chip	OK	44 degrees C / 111 degrees F
	FPC 10 I3 2 TSensor	OK	42 degrees C / 107 degrees F
	FPC 10 I3 2 Chip	OK	43 degrees C / 109 degrees F
	FPC 10 I3 3 TSensor	OK	39 degrees C / 102 degrees F
	FPC 10 I3 3 Chip	OK	44 degrees C / 111 degrees F
	FPC 10 IA 0 TSensor	OK	36 degrees C / 96 degrees F
	FPC 10 IA 0 Chip	OK	36 degrees C / 96 degrees F
	FPC 10 IA 1 TSensor	OK	43 degrees C / 109 degrees F
	FPC 10 IA 1 Chip	OK	42 degrees C / 107 degrees F
Fans	Top Fan Tray Temp	OK	37 degrees C / 98 degrees F
	Top Tray Fan 1	OK	Spinning at normal speed
	Top Tray Fan 2	OK	Spinning at normal speed
	Top Tray Fan 3	OK	Spinning at normal speed
	Top Tray Fan 4	OK	Spinning at normal speed
	Top Tray Fan 5	OK	Spinning at normal speed
	Top Tray Fan 6	OK	Spinning at normal speed
	Bottom Fan Tray Temp	OK	28 degrees C / 82 degrees F

Bottom Tray Fan 1	OK	Spinning at normal speed
Bottom Tray Fan 2	OK	Spinning at normal speed
Bottom Tray Fan 3	OK	Spinning at normal speed
Bottom Tray Fan 4	OK	Spinning at normal speed
Bottom Tray Fan 5	OK	Spinning at normal speed
Bottom Tray Fan 6	OK	Spinning at normal speed

#### show chassis environment (MX960 Router with MPC5EQ)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	50 degrees C / 122 degrees F
	PEM 1	OK	45 degrees C / 113 degrees F
	PEM 2	OK	45 degrees C / 113 degrees F
	PEM 3	Absent	
	Routing Engine 0	OK	31 degrees C / 87 degrees F
	Routing Engine 0 CPU	OK	30 degrees C / 86 degrees F
	Routing Engine 1	Present	
	Routing Engine 1 CPU	Present	
	CB 0 Intake	OK	29 degrees C / 84 degrees F
	CB 0 Exhaust A	OK	29 degrees C / 84 degrees F
	CB 0 Exhaust B	OK	34 degrees C / 93 degrees F
	CB 0 ACBC	OK	32 degrees C / 89 degrees F
	CB 0 XF A	OK	49 degrees C / 120 degrees F
	CB 0 XF B	OK	45 degrees C / 113 degrees F
	CB 1 Intake	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust A	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust B	OK	27 degrees C / 80 degrees F
	CB 1 ACBC	OK	26 degrees C / 78 degrees F
	CB 1 XF A	OK	32 degrees C / 89 degrees F
	CB 1 XF B	OK	32 degrees C / 89 degrees F
	CB 2 Intake	OK	28 degrees C / 82 degrees F
	CB 2 Exhaust A	OK	27 degrees C / 80 degrees F
	CB 2 Exhaust B	OK	33 degrees C / 91 degrees F
	CB 2 ACBC	OK	30 degrees C / 86 degrees F
	CB 2 XF A	OK	48 degrees C / 118 degrees F
	CB 2 XF B	OK	46 degrees C / 114 degrees F
	FPC 0 Intake	OK	38 degrees C / 100 degrees F
	FPC 0 Exhaust A	OK	48 degrees C / 118 degrees F
	FPC 0 Exhaust B	OK	49 degrees C / 120 degrees F
	FPC 0 XL TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XL Chip	OK	50 degrees C / 122 degrees F
	FPC 0 XL_XR0 TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XL_XR0 Chip	OK	53 degrees C / 127 degrees F
	FPC 0 XL_XR1 TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XL_XR1 Chip	OK	54 degrees C / 129 degrees F
	FPC 0 XQ TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XQ Chip	OK	52 degrees C / 125 degrees F
	FPC 0 XQ_XR0 TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XQ_XR0 Chip	OK	62 degrees C / 143 degrees F
	FPC 0 XQ_XR1 TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XQ_XR1 Chip	OK	62 degrees C / 143 degrees F
	FPC 0 XM 0 TSen	OK	53 degrees C / 127 degrees F
	FPC 0 XM 0 Chip	OK	63 degrees C / 145 degrees F
	FPC 0 XM 1 TSen	OK	53 degrees C / 127 degrees F
	FPC 0 XM 1 Chip	OK	46 degrees C / 114 degrees F
	FPC 0 PLX PCIe Switch TSe	OK	53 degrees C / 127 degrees F
	FPC 0 PLX PCIe Switch Chi	OK	66 degrees C / 150 degrees F
	FPC 1 Intake	OK	31 degrees C / 87 degrees F
	FPC 1 Exhaust A	OK	38 degrees C / 100 degrees F
	FPC 1 Exhaust B	OK	49 degrees C / 120 degrees F

FPC 1 LU 0 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 0 Chip	OK	47 degrees C / 116 degrees F
FPC 1 LU 1 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 1 Chip	OK	42 degrees C / 107 degrees F
FPC 1 LU 2 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 2 Chip	OK	46 degrees C / 114 degrees F
FPC 1 LU 3 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 1 XM 0 TSen	OK	41 degrees C / 105 degrees F
FPC 1 XM 0 Chip	OK	49 degrees C / 120 degrees F
FPC 1 XF 0 TSen	OK	41 degrees C / 105 degrees F
FPC 1 XF 0 Chip	OK	63 degrees C / 145 degrees F
FPC 1 PLX Switch TSen	OK	41 degrees C / 105 degrees F
FPC 1 PLX Switch Chip	OK	43 degrees C / 109 degrees F
FPC 3 Intake	OK	31 degrees C / 87 degrees F
FPC 3 Exhaust A	OK	37 degrees C / 98 degrees F
FPC 3 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 3 LU 0 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 0 Chip	OK	43 degrees C / 109 degrees F
FPC 3 LU 1 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 1 Chip	OK	46 degrees C / 114 degrees F
FPC 3 LU 2 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 2 Chip	OK	40 degrees C / 104 degrees F
FPC 3 LU 3 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 3 Chip	OK	41 degrees C / 105 degrees F
FPC 3 MQ 0 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 0 Chip	OK	37 degrees C / 98 degrees F
FPC 3 MQ 1 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 1 Chip	OK	40 degrees C / 104 degrees F
FPC 3 MQ 2 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 2 Chip	OK	36 degrees C / 96 degrees F
FPC 3 MQ 3 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 3 Chip	OK	38 degrees C / 100 degrees F
FPC 4 Intake	OK	34 degrees C / 93 degrees F
FPC 4 Exhaust A	OK	45 degrees C / 113 degrees F
FPC 4 Exhaust B	OK	47 degrees C / 116 degrees F
FPC 4 XL TSen	OK	44 degrees C / 111 degrees F
FPC 4 XL Chip	OK	47 degrees C / 116 degrees F
FPC 4 XL_XR0 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XL_XR0 Chip	OK	48 degrees C / 118 degrees F
FPC 4 XL_XR1 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XL_XR1 Chip	OK	47 degrees C / 116 degrees F
FPC 4 XQ TSen	OK	44 degrees C / 111 degrees F
FPC 4 XQ Chip	OK	47 degrees C / 116 degrees F
FPC 4 XQ_XR0 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XQ_XR0 Chip	OK	57 degrees C / 134 degrees F
FPC 4 XQ_XR1 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XQ_XR1 Chip	OK	58 degrees C / 136 degrees F
FPC 4 XM 0 TSen	OK	51 degrees C / 123 degrees F
FPC 4 XM 0 Chip	OK	61 degrees C / 141 degrees F
FPC 4 XM 1 TSen	OK	51 degrees C / 123 degrees F
FPC 4 XM 1 Chip	OK	47 degrees C / 116 degrees F
FPC 4 PLX PCIe Switch TSe	OK	51 degrees C / 123 degrees F
FPC 4 PLX PCIe Switch Chi	OK	60 degrees C / 140 degrees F
FPC 5 Intake	OK	34 degrees C / 93 degrees F
FPC 5 Exhaust A	OK	45 degrees C / 113 degrees F
FPC 5 Exhaust B	OK	47 degrees C / 116 degrees F
FPC 5 XL TSen	OK	45 degrees C / 113 degrees F
FPC 5 XL Chip	OK	47 degrees C / 116 degrees F
FPC 5 XL_XR0 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XL_XR0 Chip	OK	49 degrees C / 120 degrees F



FPC 5 XL_XR1 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XL_XR1 Chip	OK	49 degrees C / 120 degrees F
FPC 5 XQ TSen	OK	45 degrees C / 113 degrees F
FPC 5 XQ Chip	OK	48 degrees C / 118 degrees F
FPC 5 XQ_XR0 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XQ_XR0 Chip	OK	60 degrees C / 140 degrees F
FPC 5 XQ_XR1 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XQ_XR1 Chip	OK	58 degrees C / 136 degrees F
FPC 5 XM 0 TSen	OK	50 degrees C / 122 degrees F
FPC 5 XM 0 Chip	OK	48 degrees C / 118 degrees F
FPC 5 XM 1 TSen	OK	50 degrees C / 122 degrees F
FPC 5 XM 1 Chip	OK	47 degrees C / 116 degrees F
FPC 5 PLX PCIe Switch TSe	OK	50 degrees C / 122 degrees F
FPC 5 PLX PCIe Switch Chi	OK	59 degrees C / 138 degrees F
FPC 7 Intake	OK	32 degrees C / 89 degrees F
FPC 7 Exhaust A	OK	32 degrees C / 89 degrees F
FPC 7 Exhaust B	OK	33 degrees C / 91 degrees F
FPC 7 LU 0 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 0 Chip	OK	44 degrees C / 111 degrees F
FPC 7 LU 1 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 1 Chip	OK	47 degrees C / 116 degrees F
FPC 7 LU 2 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 2 Chip	OK	39 degrees C / 102 degrees F
FPC 7 LU 3 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 3 Chip	OK	43 degrees C / 109 degrees F
FPC 7 XM 0 TSen	OK	49 degrees C / 120 degrees F
FPC 7 XM 0 Chip	OK	57 degrees C / 134 degrees F
FPC 7 XM 1 TSen	OK	49 degrees C / 120 degrees F
FPC 7 XM 1 Chip	OK	48 degrees C / 118 degrees F
FPC 7 PLX Switch TSen	OK	49 degrees C / 120 degrees F
FPC 7 PLX Switch Chip	OK	45 degrees C / 113 degrees F
FPC 8 Intake	OK	36 degrees C / 96 degrees F
FPC 8 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 8 Exhaust B	OK	46 degrees C / 114 degrees F
FPC 8 XL TSen	OK	46 degrees C / 114 degrees F
FPC 8 XL Chip	OK	47 degrees C / 116 degrees F
FPC 8 XL_XR0 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XL_XR0 Chip	OK	53 degrees C / 127 degrees F
FPC 8 XL_XR1 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XL_XR1 Chip	OK	52 degrees C / 125 degrees F
FPC 8 XQ TSen	OK	46 degrees C / 114 degrees F
FPC 8 XQ Chip	OK	46 degrees C / 114 degrees F
FPC 8 XQ_XR0 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XQ_XR0 Chip	OK	59 degrees C / 138 degrees F
FPC 8 XQ_XR1 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XQ_XR1 Chip	OK	57 degrees C / 134 degrees F
FPC 8 XM 0 TSen	OK	52 degrees C / 125 degrees F
FPC 8 XM 0 Chip	OK	61 degrees C / 141 degrees F
FPC 8 XM 1 TSen	OK	52 degrees C / 125 degrees F
FPC 8 XM 1 Chip	OK	47 degrees C / 116 degrees F
FPC 8 PLX PCIe Switch TSe	OK	52 degrees C / 125 degrees F
FPC 8 PLX PCIe Switch Chi	OK	63 degrees C / 145 degrees F
FPC 9 Intake	OK	31 degrees C / 87 degrees F
FPC 9 Exhaust A	OK	34 degrees C / 93 degrees F
FPC 9 Exhaust B	OK	35 degrees C / 95 degrees F
FPC 9 QX 0 TSen	OK	42 degrees C / 107 degrees F
FPC 9 QX 0 Chip	OK	45 degrees C / 113 degrees F
FPC 9 LU 0 TCAM TSen	OK	42 degrees C / 107 degrees F
FPC 9 LU 0 TCAM Chip	OK	41 degrees C / 105 degrees F
FPC 9 LU 0 TSen	OK	42 degrees C / 107 degrees F
FPC 9 LU 0 Chip	OK	43 degrees C / 109 degrees F

FPC 9 MQ 0 TSen	OK	42 degrees C / 107 degrees F
FPC 9 MQ 0 Chip	OK	43 degrees C / 109 degrees F
FPC 9 QX 1 TSen	OK	38 degrees C / 100 degrees F
FPC 9 QX 1 Chip	OK	40 degrees C / 104 degrees F
FPC 9 LU 1 TCAM TSen	OK	38 degrees C / 100 degrees F
FPC 9 LU 1 TCAM Chip	OK	38 degrees C / 100 degrees F
FPC 9 LU 1 TSen	OK	38 degrees C / 100 degrees F
FPC 9 LU 1 Chip	OK	41 degrees C / 105 degrees F
FPC 9 MQ 1 TSen	OK	38 degrees C / 100 degrees F
FPC 9 MQ 1 Chip	OK	41 degrees C / 105 degrees F
FPC 10 Intake	OK	35 degrees C / 95 degrees F
FPC 10 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 10 Exhaust B	OK	46 degrees C / 114 degrees F
FPC 10 XL TSen	OK	42 degrees C / 107 degrees F
FPC 10 XL Chip	OK	44 degrees C / 111 degrees F
FPC 10 XL_XR0 TSen	OK	42 degrees C / 107 degrees F
FPC 10 XL_XR0 Chip	OK	47 degrees C / 116 degrees F
FPC 10 XL_XR1 TSen	OK	42 degrees C / 107 degrees F
FPC 10 XL_XR1 Chip	OK	48 degrees C / 118 degrees F
FPC 10 XQ TSen	OK	42 degrees C / 107 degrees F
FPC 10 XQ Chip	OK	46 degrees C / 114 degrees F
FPC 10 XQ_XR0 TSen	OK	42 degrees C / 107 degrees F
FPC 10 XQ_XR0 Chip	OK	57 degrees C / 134 degrees F
FPC 10 XQ_XR1 TSen	OK	42 degrees C / 107 degrees F
FPC 10 XQ_XR1 Chip	OK	53 degrees C / 127 degrees F
FPC 10 XM 0 TSen	OK	51 degrees C / 123 degrees F
FPC 10 XM 0 Chip	OK	61 degrees C / 141 degrees F
FPC 10 XM 1 TSen	OK	51 degrees C / 123 degrees F
FPC 10 XM 1 Chip	OK	49 degrees C / 120 degrees F
FPC 10 PLX PCIe Switch TSe	OK	51 degrees C / 123 degrees F
FPC 10 PLX PCIe Switch Chi	OK	61 degrees C / 141 degrees F
FPC 11 Intake	OK	33 degrees C / 91 degrees F
FPC 11 Exhaust A	OK	33 degrees C / 91 degrees F
FPC 11 Exhaust B	OK	34 degrees C / 93 degrees F
FPC 11 LU 0 TSen	OK	50 degrees C / 122 degrees F
FPC 11 LU 0 Chip	OK	48 degrees C / 118 degrees F
FPC 11 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 11 LU 1 Chip	OK	50 degrees C / 122 degrees F
FPC 11 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 11 LU 2 Chip	OK	41 degrees C / 105 degrees F
FPC 11 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 11 LU 3 Chip	OK	48 degrees C / 118 degrees F
FPC 11 XM 0 TSen	OK	50 degrees C / 122 degrees F
FPC 11 XM 0 Chip	OK	57 degrees C / 134 degrees F
FPC 11 XM 1 TSen	OK	50 degrees C / 122 degrees F
FPC 11 XM 1 Chip	OK	52 degrees C / 125 degrees F
FPC 11 PLX Switch TSen	OK	50 degrees C / 122 degrees F
FPC 11 PLX Switch Chip	OK	45 degrees C / 113 degrees F
Fans Top Fan Tray Temp	OK	42 degrees C / 107 degrees F
Top Tray Fan 1	OK	Spinning at high speed
Top Tray Fan 2	OK	Spinning at high speed
Top Tray Fan 3	OK	Spinning at high speed
Top Tray Fan 4	OK	Spinning at high speed
Top Tray Fan 5	OK	Spinning at high speed
Top Tray Fan 6	OK	Spinning at high speed
Top Tray Fan 7	OK	Spinning at high speed
Top Tray Fan 8	OK	Spinning at high speed
Top Tray Fan 9	OK	Spinning at high speed
Top Tray Fan 10	OK	Spinning at high speed
Top Tray Fan 11	OK	Spinning at high speed
Top Tray Fan 12	OK	Spinning at high speed

Bottom Fan Tray Temp	OK	33 degrees C / 91 degrees F
Bottom Tray Fan 1	OK	Spinning at high speed
Bottom Tray Fan 2	OK	Spinning at high speed
Bottom Tray Fan 3	OK	Spinning at high speed
Bottom Tray Fan 4	OK	Spinning at high speed
Bottom Tray Fan 5	OK	Spinning at high speed
Bottom Tray Fan 6	OK	Spinning at high speed
Bottom Tray Fan 7	OK	Spinning at high speed
Bottom Tray Fan 8	OK	Spinning at high speed
Bottom Tray Fan 9	OK	Spinning at high speed
Bottom Tray Fan 10	OK	Spinning at high speed
Bottom Tray Fan 11	OK	Spinning at high speed
Bottom Tray Fan 12	OK	Spinning at high speed

### show chassis environment (MX2020 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PSM 0	Absent	
	PSM 1	Absent	
	PSM 2	OK	41 degrees C / 105 degrees F
	PSM 3	OK	39 degrees C / 102 degrees F
	PSM 4	OK	39 degrees C / 102 degrees F
	PSM 5	OK	38 degrees C / 100 degrees F
	PSM 6	OK	38 degrees C / 100 degrees F
	PSM 7	OK	38 degrees C / 100 degrees F
	PSM 8	OK	37 degrees C / 98 degrees F
	PSM 9	Absent	
	PSM 10	Absent	
	PSM 11	OK	47 degrees C / 116 degrees F
	PSM 12	OK	45 degrees C / 113 degrees F
	PSM 13	OK	44 degrees C / 111 degrees F
	PSM 14	OK	44 degrees C / 111 degrees F
	PSM 15	OK	43 degrees C / 109 degrees F
	PSM 16	OK	42 degrees C / 107 degrees F
	PSM 17	OK	41 degrees C / 105 degrees F
	PDM 0	OK	
	PDM 1	Absent	
	PDM 2	Absent	
	PDM 3	OK	
	CB 0 IntakeA-Zone0	OK	45 degrees C / 113 degrees F
	CB 0 IntakeB-Zone1	OK	34 degrees C / 93 degrees F
	CB 0 IntakeC-Zone0	OK	48 degrees C / 118 degrees F
	CB 0 ExhaustA-Zone0	OK	45 degrees C / 113 degrees F
	CB 0 ExhaustB-Zone1	OK	37 degrees C / 98 degrees F
	CB 0 TCBC-Zone0	OK	41 degrees C / 105 degrees F
	CB 1 IntakeA-Zone0	OK	46 degrees C / 114 degrees F
	CB 1 IntakeB-Zone1	OK	42 degrees C / 107 degrees F
	CB 1 IntakeC-Zone0	OK	49 degrees C / 120 degrees F
	CB 1 ExhaustA-Zone0	OK	46 degrees C / 114 degrees F
	CB 1 ExhaustB-Zone1	OK	41 degrees C / 105 degrees F
	CB 1 TCBC-Zone0	OK	46 degrees C / 114 degrees F
	SPMB 0 Intake	OK	33 degrees C / 91 degrees F
	SPMB 1 Intake	OK	42 degrees C / 107 degrees F
	Routing Engine 0	OK	35 degrees C / 95 degrees F
	Routing Engine 0 CPU	OK	34 degrees C / 93 degrees F
	Routing Engine 1	OK	44 degrees C / 111 degrees F
	Routing Engine 1 CPU	OK	42 degrees C / 107 degrees F
	SFB 0 Intake-Zone0	OK	55 degrees C / 131 degrees F
	SFB 0 Exhaust-Zone1	OK	48 degrees C / 118 degrees F
	SFB 0 IntakeA-Zone0	OK	50 degrees C / 122 degrees F

SFB 0 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 0 Exhaust-Zone0	OK	52 degrees C / 125 degrees F
SFB 0 SFB-XF2-Zone1	OK	61 degrees C / 141 degrees F
SFB 0 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 0 SFB-XF0-Zone0	OK	68 degrees C / 154 degrees F
SFB 1 Intake-Zone0	OK	56 degrees C / 132 degrees F
SFB 1 Exhaust-Zone1	OK	47 degrees C / 116 degrees F
SFB 1 IntakeA-Zone0	OK	51 degrees C / 123 degrees F
SFB 1 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 1 Exhaust-Zone0	OK	51 degrees C / 123 degrees F
SFB 1 SFB-XF2-Zone1	OK	62 degrees C / 143 degrees F
SFB 1 SFB-XF1-Zone0	OK	67 degrees C / 152 degrees F
SFB 1 SFB-XF0-Zone0	OK	69 degrees C / 156 degrees F
SFB 2 Intake-Zone0	OK	56 degrees C / 132 degrees F
SFB 2 Exhaust-Zone1	OK	47 degrees C / 116 degrees F
SFB 2 IntakeA-Zone0	OK	51 degrees C / 123 degrees F
SFB 2 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 2 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 2 SFB-XF2-Zone1	OK	65 degrees C / 149 degrees F
SFB 2 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 2 SFB-XF0-Zone0	OK	70 degrees C / 158 degrees F
SFB 3 Intake-Zone0	OK	57 degrees C / 134 degrees F
SFB 3 Exhaust-Zone1	OK	48 degrees C / 118 degrees F
SFB 3 IntakeA-Zone0	OK	52 degrees C / 125 degrees F
SFB 3 IntakeB-Zone1	OK	41 degrees C / 105 degrees F
SFB 3 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 3 SFB-XF2-Zone1	OK	66 degrees C / 150 degrees F
SFB 3 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 3 SFB-XF0-Zone0	OK	71 degrees C / 159 degrees F
SFB 4 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 4 Exhaust-Zone1	OK	49 degrees C / 120 degrees F
SFB 4 IntakeA-Zone0	OK	54 degrees C / 129 degrees F
SFB 4 IntakeB-Zone1	OK	42 degrees C / 107 degrees F
SFB 4 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 4 SFB-XF2-Zone1	OK	64 degrees C / 147 degrees F
SFB 4 SFB-XF1-Zone0	OK	68 degrees C / 154 degrees F
SFB 4 SFB-XF0-Zone0	OK	71 degrees C / 159 degrees F
SFB 5 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 5 Exhaust-Zone1	OK	50 degrees C / 122 degrees F
SFB 5 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 5 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 5 Exhaust-Zone0	OK	54 degrees C / 129 degrees F
SFB 5 SFB-XF2-Zone1	OK	66 degrees C / 150 degrees F
SFB 5 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 5 SFB-XF0-Zone0	OK	74 degrees C / 165 degrees F
SFB 6 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 6 Exhaust-Zone1	OK	49 degrees C / 120 degrees F
SFB 6 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 6 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 6 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 6 SFB-XF2-Zone1	OK	65 degrees C / 149 degrees F
SFB 6 SFB-XF1-Zone0	OK	68 degrees C / 154 degrees F
SFB 6 SFB-XF0-Zone0	OK	72 degrees C / 161 degrees F
SFB 7 Intake-Zone0	OK	57 degrees C / 134 degrees F
SFB 7 Exhaust-Zone1	OK	50 degrees C / 122 degrees F
SFB 7 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 7 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 7 Exhaust-Zone0	OK	54 degrees C / 129 degrees F
SFB 7 SFB-XF2-Zone1	OK	68 degrees C / 154 degrees F
SFB 7 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 7 SFB-XF0-Zone0	OK	73 degrees C / 163 degrees F

FPC 0 Intake	OK	41 degrees C / 105 degrees F
FPC 0 Exhaust A	OK	48 degrees C / 118 degrees F
FPC 0 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 0 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 0 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 1 Chip	OK	64 degrees C / 147 degrees F
FPC 0 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 0 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 0 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 0 Chip	OK	49 degrees C / 120 degrees F
FPC 0 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 1 Chip	OK	51 degrees C / 123 degrees F
FPC 0 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 0 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 3 Chip	OK	45 degrees C / 113 degrees F
FPC 1 Intake	OK	40 degrees C / 104 degrees F
FPC 1 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 1 Exhaust B	OK	58 degrees C / 136 degrees F
FPC 1 LU 0 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 0 Chip	OK	56 degrees C / 132 degrees F
FPC 1 LU 1 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 1 Chip	OK	58 degrees C / 136 degrees F
FPC 1 LU 2 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 2 Chip	OK	49 degrees C / 120 degrees F
FPC 1 LU 3 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 1 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 0 Chip	OK	48 degrees C / 118 degrees F
FPC 1 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 1 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 1 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 3 Chip	OK	44 degrees C / 111 degrees F
FPC 2 Intake	OK	39 degrees C / 102 degrees F
FPC 2 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 2 Exhaust B	OK	61 degrees C / 141 degrees F
FPC 2 LU 0 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 2 LU 1 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 2 LU 2 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 2 LU 3 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 2 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 2 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 2 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 2 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 3 Intake	OK	40 degrees C / 104 degrees F
FPC 3 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 3 Exhaust B	OK	61 degrees C / 141 degrees F
FPC 3 LU 0 TSen	OK	58 degrees C / 136 degrees F

FPC 3 LU 0 Chip	OK	61 degrees C / 141 degrees F
FPC 3 LU 1 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 3 LU 2 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 3 LU 3 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 3 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 3 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 3 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 3 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 3 Chip	OK	48 degrees C / 118 degrees F
FPC 4 Intake	OK	40 degrees C / 104 degrees F
FPC 4 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 4 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 4 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 4 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 4 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 4 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 4 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 4 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 4 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 4 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 5 Intake	OK	41 degrees C / 105 degrees F
FPC 5 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 5 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 5 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 0 Chip	OK	63 degrees C / 145 degrees F
FPC 5 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 1 Chip	OK	66 degrees C / 150 degrees F
FPC 5 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 2 Chip	OK	56 degrees C / 132 degrees F
FPC 5 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 5 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 5 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 5 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 2 Chip	OK	48 degrees C / 118 degrees F
FPC 5 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 6 Intake	OK	42 degrees C / 107 degrees F
FPC 6 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 6 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 6 LU 0 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 0 Chip	OK	64 degrees C / 147 degrees F
FPC 6 LU 1 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 1 Chip	OK	66 degrees C / 150 degrees F
FPC 6 LU 2 TSen	OK	61 degrees C / 141 degrees F

FPC 6 LU 2 Chip	OK	56 degrees C / 132 degrees F
FPC 6 LU 3 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 3 Chip	OK	56 degrees C / 132 degrees F
FPC 6 MQ 0 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 0 Chip	OK	56 degrees C / 132 degrees F
FPC 6 MQ 1 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 1 Chip	OK	59 degrees C / 138 degrees F
FPC 6 MQ 2 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 2 Chip	OK	49 degrees C / 120 degrees F
FPC 6 MQ 3 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 7 Intake	OK	41 degrees C / 105 degrees F
FPC 7 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 7 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 7 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 0 Chip	OK	61 degrees C / 141 degrees F
FPC 7 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 7 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 2 Chip	OK	54 degrees C / 129 degrees F
FPC 7 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 7 MQ 0 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 0 Chip	OK	53 degrees C / 127 degrees F
FPC 7 MQ 1 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 7 MQ 2 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 7 MQ 3 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 8 Intake	OK	41 degrees C / 105 degrees F
FPC 8 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 8 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 8 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 8 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 1 Chip	OK	64 degrees C / 147 degrees F
FPC 8 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 2 Chip	OK	55 degrees C / 131 degrees F
FPC 8 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 8 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 0 Chip	OK	51 degrees C / 123 degrees F
FPC 8 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 8 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 8 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 9 Intake	OK	42 degrees C / 107 degrees F
FPC 9 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 9 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 9 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 0 Chip	OK	65 degrees C / 149 degrees F
FPC 9 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 1 Chip	OK	67 degrees C / 152 degrees F
FPC 9 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 2 Chip	OK	54 degrees C / 129 degrees F
FPC 9 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 9 MQ 0 TSen	OK	51 degrees C / 123 degrees F

FPC 9 MQ 0 Chip	OK	55 degrees C / 131 degrees F
FPC 9 MQ 1 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 1 Chip	OK	59 degrees C / 138 degrees F
FPC 9 MQ 2 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 2 Chip	OK	49 degrees C / 120 degrees F
FPC 9 MQ 3 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 10 Intake	OK	44 degrees C / 111 degrees F
FPC 10 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 10 Exhaust B	OK	55 degrees C / 131 degrees F
FPC 10 LU 0 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 0 Chip	OK	55 degrees C / 131 degrees F
FPC 10 LU 1 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 1 Chip	OK	59 degrees C / 138 degrees F
FPC 10 LU 2 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 2 Chip	OK	52 degrees C / 125 degrees F
FPC 10 LU 3 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 10 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 0 Chip	OK	49 degrees C / 120 degrees F
FPC 10 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 10 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 10 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 11 Intake	OK	30 degrees C / 86 degrees F
FPC 11 Exhaust A	OK	35 degrees C / 95 degrees F
FPC 11 Exhaust B	OK	30 degrees C / 86 degrees F
FPC 11 LU 0 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 0 Chip	OK	58 degrees C / 136 degrees F
FPC 11 LU 1 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 11 LU 2 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 11 LU 3 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 11 MQ 0 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 11 MQ 1 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 1 Chip	OK	57 degrees C / 134 degrees F
FPC 11 MQ 2 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 2 Chip	OK	48 degrees C / 118 degrees F
FPC 11 MQ 3 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 3 Chip	OK	52 degrees C / 125 degrees F
FPC 12 Intake	OK	40 degrees C / 104 degrees F
FPC 12 Exhaust A	OK	47 degrees C / 116 degrees F
FPC 12 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 12 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 12 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 12 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 2 Chip	OK	47 degrees C / 116 degrees F
FPC 12 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 3 Chip	OK	50 degrees C / 122 degrees F
FPC 12 MQ 0 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 12 MQ 1 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 12 MQ 2 TSen	OK	46 degrees C / 114 degrees F



FPC 12 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 12 MQ 3 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 13 Intake	OK	40 degrees C / 104 degrees F
FPC 13 Exhaust A	OK	48 degrees C / 118 degrees F
FPC 13 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 13 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 13 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 13 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 2 Chip	OK	48 degrees C / 118 degrees F
FPC 13 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 3 Chip	OK	48 degrees C / 118 degrees F
FPC 13 MQ 0 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 13 MQ 1 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 13 MQ 2 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 13 MQ 3 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 14 Intake	OK	40 degrees C / 104 degrees F
FPC 14 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 14 Exhaust B	OK	51 degrees C / 123 degrees F
FPC 14 LU 0 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 0 Chip	OK	50 degrees C / 122 degrees F
FPC 14 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 1 Chip	OK	54 degrees C / 129 degrees F
FPC 14 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 2 Chip	OK	47 degrees C / 116 degrees F
FPC 14 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 3 Chip	OK	49 degrees C / 120 degrees F
FPC 14 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 14 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 1 Chip	OK	51 degrees C / 123 degrees F
FPC 14 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 14 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 3 Chip	OK	48 degrees C / 118 degrees F
FPC 15 Intake	OK	44 degrees C / 111 degrees F
FPC 15 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 15 Exhaust B	OK	60 degrees C / 140 degrees F
FPC 15 LU 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 0 Chip	OK	56 degrees C / 132 degrees F
FPC 15 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 1 Chip	OK	50 degrees C / 122 degrees F
FPC 15 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 2 Chip	OK	58 degrees C / 136 degrees F
FPC 15 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 3 Chip	OK	63 degrees C / 145 degrees F
FPC 15 XM 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 XM 0 Chip	OK	56 degrees C / 132 degrees F
FPC 15 XF 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 XF 0 Chip	OK	68 degrees C / 154 degrees F
FPC 15 PLX Switch TSen	OK	50 degrees C / 122 degrees F
FPC 15 PLX Switch Chip	OK	56 degrees C / 132 degrees F
FPC 16 Intake	OK	42 degrees C / 107 degrees F
FPC 16 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 16 Exhaust B	OK	53 degrees C / 127 degrees F

FPC 16 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 16 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 16 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 2 Chip	OK	48 degrees C / 118 degrees F
FPC 16 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 3 Chip	OK	49 degrees C / 120 degrees F
FPC 16 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 0 Chip	OK	48 degrees C / 118 degrees F
FPC 16 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 16 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 16 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 17 Intake	OK	43 degrees C / 109 degrees F
FPC 17 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 17 Exhaust B	OK	55 degrees C / 131 degrees F
FPC 17 LU 0 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 0 Chip	OK	57 degrees C / 134 degrees F
FPC 17 LU 1 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 1 Chip	OK	60 degrees C / 140 degrees F
FPC 17 LU 2 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 17 LU 3 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 17 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 17 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 17 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 17 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 3 Chip	OK	51 degrees C / 123 degrees F
FPC 18 Intake	OK	44 degrees C / 111 degrees F
FPC 18 Exhaust A	OK	53 degrees C / 127 degrees F
FPC 18 Exhaust B	OK	57 degrees C / 134 degrees F
FPC 18 LU 0 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 0 Chip	OK	57 degrees C / 134 degrees F
FPC 18 LU 1 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 18 LU 2 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 18 LU 3 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 3 Chip	OK	55 degrees C / 131 degrees F
FPC 18 MQ 0 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 0 Chip	OK	54 degrees C / 129 degrees F
FPC 18 MQ 1 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 1 Chip	OK	58 degrees C / 136 degrees F
FPC 18 MQ 2 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 2 Chip	OK	50 degrees C / 122 degrees F
FPC 18 MQ 3 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 3 Chip	OK	53 degrees C / 127 degrees F
FPC 19 Intake	OK	48 degrees C / 118 degrees F
FPC 19 Exhaust A	OK	56 degrees C / 132 degrees F
FPC 19 Exhaust B	OK	64 degrees C / 147 degrees F
FPC 19 LU 0 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 0 Chip	OK	64 degrees C / 147 degrees F
FPC 19 LU 1 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 1 Chip	OK	70 degrees C / 158 degrees F

FPC 19 LU 2 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 2 Chip	OK	61 degrees C / 141 degrees F
FPC 19 LU 3 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 3 Chip	OK	62 degrees C / 143 degrees F
FPC 19 MQ 0 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 0 Chip	OK	60 degrees C / 140 degrees F
FPC 19 MQ 1 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 1 Chip	OK	62 degrees C / 143 degrees F
FPC 19 MQ 2 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 2 Chip	OK	56 degrees C / 132 degrees F
FPC 19 MQ 3 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 3 Chip	OK	57 degrees C / 134 degrees F
ADC 0 Intake	OK	40 degrees C / 104 degrees F
ADC 0 Exhaust	OK	52 degrees C / 125 degrees F
ADC 0 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 0 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 1 Intake	OK	38 degrees C / 100 degrees F
ADC 1 Exhaust	OK	50 degrees C / 122 degrees F
ADC 1 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 1 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 2 Intake	OK	37 degrees C / 98 degrees F
ADC 2 Exhaust	OK	52 degrees C / 125 degrees F
ADC 2 ADC-XF1	OK	53 degrees C / 127 degrees F
ADC 2 ADC-XF0	OK	61 degrees C / 141 degrees F
ADC 3 Intake	OK	40 degrees C / 104 degrees F
ADC 3 Exhaust	OK	51 degrees C / 123 degrees F
ADC 3 ADC-XF1	OK	61 degrees C / 141 degrees F
ADC 3 ADC-XF0	OK	64 degrees C / 147 degrees F
ADC 4 Intake	OK	39 degrees C / 102 degrees F
ADC 4 Exhaust	OK	51 degrees C / 123 degrees F
ADC 4 ADC-XF1	OK	60 degrees C / 140 degrees F
ADC 4 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 5 Intake	OK	38 degrees C / 100 degrees F
ADC 5 Exhaust	OK	54 degrees C / 129 degrees F
ADC 5 ADC-XF1	OK	56 degrees C / 132 degrees F
ADC 5 ADC-XF0	OK	67 degrees C / 152 degrees F
ADC 6 Intake	OK	39 degrees C / 102 degrees F
ADC 6 Exhaust	OK	52 degrees C / 125 degrees F
ADC 6 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 6 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 7 Intake	OK	39 degrees C / 102 degrees F
ADC 7 Exhaust	OK	54 degrees C / 129 degrees F
ADC 7 ADC-XF1	OK	62 degrees C / 143 degrees F
ADC 7 ADC-XF0	OK	70 degrees C / 158 degrees F
ADC 8 Intake	OK	39 degrees C / 102 degrees F
ADC 8 Exhaust	OK	52 degrees C / 125 degrees F
ADC 8 ADC-XF1	OK	61 degrees C / 141 degrees F
ADC 8 ADC-XF0	OK	65 degrees C / 149 degrees F
ADC 9 Intake	OK	41 degrees C / 105 degrees F
ADC 9 Exhaust	OK	51 degrees C / 123 degrees F
ADC 9 ADC-XF1	OK	63 degrees C / 145 degrees F
ADC 9 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 10 Intake	OK	48 degrees C / 118 degrees F
ADC 10 Exhaust	OK	53 degrees C / 127 degrees F
ADC 10 ADC-XF1	OK	67 degrees C / 152 degrees F
ADC 10 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 12 Intake	OK	49 degrees C / 120 degrees F
ADC 12 Exhaust	OK	54 degrees C / 129 degrees F
ADC 12 ADC-XF1	OK	67 degrees C / 152 degrees F
ADC 12 ADC-XF0	OK	67 degrees C / 152 degrees F
ADC 13 Intake	OK	49 degrees C / 120 degrees F

	ADC 13 Exhaust	OK	57 degrees C / 134 degrees F
	ADC 13 ADC-XF1	OK	66 degrees C / 150 degrees F
	ADC 13 ADC-XF0	OK	69 degrees C / 156 degrees F
	ADC 14 Intake	OK	51 degrees C / 123 degrees F
	ADC 14 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 14 ADC-XF1	OK	69 degrees C / 156 degrees F
	ADC 14 ADC-XF0	OK	74 degrees C / 165 degrees F
	ADC 15 Intake	OK	50 degrees C / 122 degrees F
	ADC 15 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 15 ADC-XF1	OK	68 degrees C / 154 degrees F
	ADC 15 ADC-XF0	OK	69 degrees C / 156 degrees F
	ADC 16 Intake	OK	52 degrees C / 125 degrees F
	ADC 16 Exhaust	OK	58 degrees C / 136 degrees F
	ADC 16 ADC-XF1	OK	68 degrees C / 154 degrees F
	ADC 16 ADC-XF0	OK	70 degrees C / 158 degrees F
	ADC 17 Intake	OK	52 degrees C / 125 degrees F
	ADC 17 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 17 ADC-XF1	OK	69 degrees C / 156 degrees F
	ADC 17 ADC-XF0	OK	71 degrees C / 159 degrees F
	ADC 18 Intake	OK	53 degrees C / 127 degrees F
	ADC 18 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 18 ADC-XF1	OK	68 degrees C / 154 degrees F
	ADC 18 ADC-XF0	OK	73 degrees C / 163 degrees F
	ADC 19 Intake	OK	50 degrees C / 122 degrees F
	ADC 19 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 19 ADC-XF1	OK	68 degrees C / 154 degrees F
	ADC 19 ADC-XF0	OK	72 degrees C / 161 degrees F
Fans	Fan Tray 0 Fan 1	OK	7440 RPM
	Fan Tray 0 Fan 2	OK	7200 RPM
	Fan Tray 0 Fan 3	OK	6960 RPM
	Fan Tray 0 Fan 4	OK	7200 RPM
	Fan Tray 0 Fan 5	OK	7080 RPM
	Fan Tray 0 Fan 6	OK	6840 RPM
	Fan Tray 1 Fan 1	OK	6840 RPM
	Fan Tray 1 Fan 2	OK	6960 RPM
	Fan Tray 1 Fan 3	OK	6960 RPM
	Fan Tray 1 Fan 4	OK	7080 RPM
	Fan Tray 1 Fan 5	OK	6960 RPM
	Fan Tray 1 Fan 6	OK	6960 RPM
	Fan Tray 2 Fan 1	OK	8640 RPM
	Fan Tray 2 Fan 2	OK	8640 RPM
	Fan Tray 2 Fan 3	OK	8760 RPM
	Fan Tray 2 Fan 4	OK	8760 RPM
	Fan Tray 2 Fan 5	OK	8640 RPM
	Fan Tray 2 Fan 6	OK	8640 RPM
	Fan Tray 3 Fan 1	OK	8520 RPM
	Fan Tray 3 Fan 2	OK	8520 RPM
	Fan Tray 3 Fan 3	OK	8640 RPM
	Fan Tray 3 Fan 4	OK	8640 RPM
	Fan Tray 3 Fan 5	OK	8520 RPM
	Fan Tray 3 Fan 6	OK	8520 RPM

## show chassis environment (MX2020 Router with MPC5EQ and MPC6E)

Class	Item	Status	Measurement
Temp	PSM 0	OK	32 degrees C / 89 degrees F
	PSM 1	OK	32 degrees C / 89 degrees F
	PSM 2	OK	32 degrees C / 89 degrees F
	PSM 3	OK	32 degrees C / 89 degrees F
	PSM 4	OK	32 degrees C / 89 degrees F
	PSM 5	OK	33 degrees C / 91 degrees F

PSM 6	OK	32 degrees C / 89 degrees F
PSM 7	OK	32 degrees C / 89 degrees F
PSM 8	OK	32 degrees C / 89 degrees F
PSM 9	Absent	
PSM 10	Absent	
PSM 11	Absent	
PSM 12	OK	33 degrees C / 91 degrees F
PSM 13	OK	33 degrees C / 91 degrees F
PSM 14	OK	34 degrees C / 93 degrees F
PSM 15	OK	34 degrees C / 93 degrees F
PSM 16	OK	33 degrees C / 91 degrees F
PSM 17	OK	33 degrees C / 91 degrees F
PDM 0	OK	
PDM 1	OK	
PDM 2	OK	
PDM 3	OK	
CB 0 IntakeA-Zone0	OK	34 degrees C / 93 degrees F
CB 0 IntakeB-Zone1	OK	26 degrees C / 78 degrees F
CB 0 IntakeC-Zone0	OK	38 degrees C / 100 degrees F
CB 0 ExhaustA-Zone0	OK	34 degrees C / 93 degrees F
CB 0 ExhaustB-Zone1	OK	27 degrees C / 80 degrees F
CB 0 TCBC-Zone0	OK	32 degrees C / 89 degrees F
CB 1 IntakeA-Zone0	OK	24 degrees C / 75 degrees F
CB 1 IntakeB-Zone1	OK	22 degrees C / 71 degrees F
CB 1 IntakeC-Zone0	OK	34 degrees C / 93 degrees F
CB 1 ExhaustA-Zone0	OK	31 degrees C / 87 degrees F
CB 1 ExhaustB-Zone1	OK	24 degrees C / 75 degrees F
CB 1 TCBC-Zone0	OK	27 degrees C / 80 degrees F
SPMB 0 Intake	OK	25 degrees C / 77 degrees F
SPMB 1 Intake	OK	23 degrees C / 73 degrees F
Routing Engine 0	OK	28 degrees C / 82 degrees F
Routing Engine 0 CPU	OK	25 degrees C / 77 degrees F
Routing Engine 1	OK	25 degrees C / 77 degrees F
Routing Engine 1 CPU	OK	24 degrees C / 75 degrees F
SFB 0 Intake-Zone0	OK	45 degrees C / 113 degrees F
SFB 0 Exhaust-Zone1	OK	34 degrees C / 93 degrees F
SFB 0 IntakeA-Zone0	OK	32 degrees C / 89 degrees F
SFB 0 IntakeB-Zone1	OK	28 degrees C / 82 degrees F
SFB 0 Exhaust-Zone0	OK	36 degrees C / 96 degrees F
SFB 0 SFB-XF2-Zone1	OK	46 degrees C / 114 degrees F
SFB 0 SFB-XF1-Zone0	OK	48 degrees C / 118 degrees F
SFB 0 SFB-XF0-Zone0	OK	60 degrees C / 140 degrees F
SFB 1 Intake-Zone0	OK	44 degrees C / 111 degrees F
SFB 1 Exhaust-Zone1	OK	34 degrees C / 93 degrees F
SFB 1 IntakeA-Zone0	OK	35 degrees C / 95 degrees F
SFB 1 IntakeB-Zone1	OK	27 degrees C / 80 degrees F
SFB 1 Exhaust-Zone0	OK	37 degrees C / 98 degrees F
SFB 1 SFB-XF2-Zone1	OK	47 degrees C / 116 degrees F
SFB 1 SFB-XF1-Zone0	OK	49 degrees C / 120 degrees F
SFB 1 SFB-XF0-Zone0	OK	56 degrees C / 132 degrees F
SFB 2 Intake-Zone0	OK	41 degrees C / 105 degrees F
SFB 2 Exhaust-Zone1	OK	34 degrees C / 93 degrees F
SFB 2 IntakeA-Zone0	OK	35 degrees C / 95 degrees F
SFB 2 IntakeB-Zone1	OK	28 degrees C / 82 degrees F
SFB 2 Exhaust-Zone0	OK	37 degrees C / 98 degrees F
SFB 2 SFB-XF2-Zone1	OK	47 degrees C / 116 degrees F
SFB 2 SFB-XF1-Zone0	OK	55 degrees C / 131 degrees F
SFB 2 SFB-XF0-Zone0	OK	55 degrees C / 131 degrees F
SFB 3 Intake-Zone0	OK	43 degrees C / 109 degrees F
SFB 3 Exhaust-Zone1	OK	33 degrees C / 91 degrees F
SFB 3 IntakeA-Zone0	OK	35 degrees C / 95 degrees F

SFB 3 IntakeB-Zone1	OK	27 degrees C / 80 degrees F
SFB 3 Exhaust-Zone0	OK	36 degrees C / 96 degrees F
SFB 3 SFB-XF2-Zone1	OK	46 degrees C / 114 degrees F
SFB 3 SFB-XF1-Zone0	OK	46 degrees C / 114 degrees F
SFB 3 SFB-XF0-Zone0	OK	57 degrees C / 134 degrees F
SFB 4 Intake-Zone0	OK	36 degrees C / 96 degrees F
SFB 4 Exhaust-Zone1	OK	32 degrees C / 89 degrees F
SFB 4 IntakeA-Zone0	OK	31 degrees C / 87 degrees F
SFB 4 IntakeB-Zone1	OK	26 degrees C / 78 degrees F
SFB 4 Exhaust-Zone0	OK	32 degrees C / 89 degrees F
SFB 4 SFB-XF2-Zone1	OK	44 degrees C / 111 degrees F
SFB 4 SFB-XF1-Zone0	OK	45 degrees C / 113 degrees F
SFB 4 SFB-XF0-Zone0	OK	52 degrees C / 125 degrees F
SFB 5 Intake-Zone0	OK	31 degrees C / 87 degrees F
SFB 5 Exhaust-Zone1	OK	30 degrees C / 86 degrees F
SFB 5 IntakeA-Zone0	OK	26 degrees C / 78 degrees F
SFB 5 IntakeB-Zone1	OK	24 degrees C / 75 degrees F
SFB 5 Exhaust-Zone0	OK	29 degrees C / 84 degrees F
SFB 5 SFB-XF2-Zone1	OK	43 degrees C / 109 degrees F
SFB 5 SFB-XF1-Zone0	OK	47 degrees C / 116 degrees F
SFB 5 SFB-XF0-Zone0	OK	49 degrees C / 120 degrees F
SFB 6 Intake-Zone0	OK	30 degrees C / 86 degrees F
SFB 6 Exhaust-Zone1	OK	29 degrees C / 84 degrees F
SFB 6 IntakeA-Zone0	OK	25 degrees C / 77 degrees F
SFB 6 IntakeB-Zone1	OK	24 degrees C / 75 degrees F
SFB 6 Exhaust-Zone0	OK	29 degrees C / 84 degrees F
SFB 6 SFB-XF2-Zone1	OK	43 degrees C / 109 degrees F
SFB 6 SFB-XF1-Zone0	OK	44 degrees C / 111 degrees F
SFB 6 SFB-XF0-Zone0	OK	45 degrees C / 113 degrees F
SFB 7 Intake-Zone0	OK	31 degrees C / 87 degrees F
SFB 7 Exhaust-Zone1	OK	30 degrees C / 86 degrees F
SFB 7 IntakeA-Zone0	OK	26 degrees C / 78 degrees F
SFB 7 IntakeB-Zone1	OK	24 degrees C / 75 degrees F
SFB 7 Exhaust-Zone0	OK	28 degrees C / 82 degrees F
SFB 7 SFB-XF2-Zone1	OK	50 degrees C / 122 degrees F
SFB 7 SFB-XF1-Zone0	OK	43 degrees C / 109 degrees F
SFB 7 SFB-XF0-Zone0	OK	47 degrees C / 116 degrees F
FPC 0 Intake	OK	31 degrees C / 87 degrees F
FPC 0 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 0 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 0 XL TSen	OK	42 degrees C / 107 degrees F
FPC 0 XL Chip	OK	46 degrees C / 114 degrees F
FPC 0 XL_XR0 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XL_XR0 Chip	OK	48 degrees C / 118 degrees F
FPC 0 XL_XR1 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XL_XR1 Chip	OK	48 degrees C / 118 degrees F
FPC 0 XQ TSen	OK	42 degrees C / 107 degrees F
FPC 0 XQ Chip	OK	44 degrees C / 111 degrees F
FPC 0 XQ_XR0 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XQ_XR0 Chip	OK	57 degrees C / 134 degrees F
FPC 0 XQ_XR1 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XQ_XR1 Chip	OK	55 degrees C / 131 degrees F
FPC 0 XM 0 TSen	OK	48 degrees C / 118 degrees F
FPC 0 XM 0 Chip	OK	62 degrees C / 143 degrees F
FPC 0 XM 1 TSen	OK	48 degrees C / 118 degrees F
FPC 0 XM 1 Chip	OK	44 degrees C / 111 degrees F
FPC 0 PLX PCIe Switch TSe	OK	48 degrees C / 118 degrees F
FPC 0 PLX PCIe Switch Chi	OK	57 degrees C / 134 degrees F
FPC 1 Intake	OK	29 degrees C / 84 degrees F
FPC 1 Exhaust A	OK	36 degrees C / 96 degrees F
FPC 1 Exhaust B	OK	44 degrees C / 111 degrees F

FPC 1 LU 0 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 0 Chip	OK	45 degrees C / 113 degrees F
FPC 1 LU 1 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 1 Chip	OK	38 degrees C / 100 degrees F
FPC 1 LU 2 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 2 Chip	OK	42 degrees C / 107 degrees F
FPC 1 LU 3 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 3 Chip	OK	47 degrees C / 116 degrees F
FPC 1 XM 0 TSen	OK	38 degrees C / 100 degrees F
FPC 1 XM 0 Chip	OK	44 degrees C / 111 degrees F
FPC 1 XF 0 TSen	OK	38 degrees C / 100 degrees F
FPC 1 XF 0 Chip	OK	54 degrees C / 129 degrees F
FPC 1 PLX Switch TSen	OK	38 degrees C / 100 degrees F
FPC 1 PLX Switch Chip	OK	41 degrees C / 105 degrees F
FPC 2 Intake	OK	28 degrees C / 82 degrees F
FPC 2 Exhaust A	OK	28 degrees C / 82 degrees F
FPC 2 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 2 LU 0 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 0 Chip	OK	40 degrees C / 104 degrees F
FPC 2 LU 1 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 1 Chip	OK	41 degrees C / 105 degrees F
FPC 2 LU 2 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 2 Chip	OK	34 degrees C / 93 degrees F
FPC 2 LU 3 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 3 Chip	OK	38 degrees C / 100 degrees F
FPC 2 XM 0 TSen	OK	40 degrees C / 104 degrees F
FPC 2 XM 0 Chip	OK	47 degrees C / 116 degrees F
FPC 2 XM 1 TSen	OK	40 degrees C / 104 degrees F
FPC 2 XM 1 Chip	OK	42 degrees C / 107 degrees F
FPC 2 PLX Switch TSen	OK	40 degrees C / 104 degrees F
FPC 2 PLX Switch Chip	OK	39 degrees C / 102 degrees F
FPC 3 Intake	OK	27 degrees C / 80 degrees F
FPC 3 Exhaust A	OK	38 degrees C / 100 degrees F
FPC 3 Exhaust B	OK	31 degrees C / 87 degrees F
FPC 3 QX 0 TSen	OK	38 degrees C / 100 degrees F
FPC 3 QX 0 Chip	OK	42 degrees C / 107 degrees F
FPC 3 LU 0 TCAM TSen	OK	38 degrees C / 100 degrees F
FPC 3 LU 0 TCAM Chip	OK	43 degrees C / 109 degrees F
FPC 3 LU 0 TSen	OK	38 degrees C / 100 degrees F
FPC 3 LU 0 Chip	OK	42 degrees C / 107 degrees F
FPC 3 MQ 0 TSen	OK	38 degrees C / 100 degrees F
FPC 3 MQ 0 Chip	OK	39 degrees C / 102 degrees F
FPC 3 QX 1 TSen	OK	32 degrees C / 89 degrees F
FPC 3 QX 1 Chip	OK	36 degrees C / 96 degrees F
FPC 3 LU 1 TCAM TSen	OK	32 degrees C / 89 degrees F
FPC 3 LU 1 TCAM Chip	OK	35 degrees C / 95 degrees F
FPC 3 LU 1 TSen	OK	32 degrees C / 89 degrees F
FPC 3 LU 1 Chip	OK	37 degrees C / 98 degrees F
FPC 3 MQ 1 TSen	OK	32 degrees C / 89 degrees F
FPC 3 MQ 1 Chip	OK	36 degrees C / 96 degrees F
FPC 4 Intake	OK	29 degrees C / 84 degrees F
FPC 4 Exhaust A	OK	36 degrees C / 96 degrees F
FPC 4 Exhaust B	OK	40 degrees C / 104 degrees F
FPC 4 XL TSen	OK	39 degrees C / 102 degrees F
FPC 4 XL Chip	OK	42 degrees C / 107 degrees F
FPC 4 XL_XR0 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XL_XR0 Chip	OK	45 degrees C / 113 degrees F
FPC 4 XL_XR1 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XL_XR1 Chip	OK	46 degrees C / 114 degrees F
FPC 4 XQ TSen	OK	39 degrees C / 102 degrees F
FPC 4 XQ Chip	OK	42 degrees C / 107 degrees F

FPC 4 XQ_XR0 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XQ_XR0 Chip	OK	54 degrees C / 129 degrees F
FPC 4 XQ_XR1 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XQ_XR1 Chip	OK	53 degrees C / 127 degrees F
FPC 4 XM 0 TSen	OK	45 degrees C / 113 degrees F
FPC 4 XM 0 Chip	OK	59 degrees C / 138 degrees F
FPC 4 XM 1 TSen	OK	45 degrees C / 113 degrees F
FPC 4 XM 1 Chip	OK	41 degrees C / 105 degrees F
FPC 4 PLX PCIe Switch TSe	OK	45 degrees C / 113 degrees F
FPC 4 PLX PCIe Switch Chi	OK	58 degrees C / 136 degrees F
FPC 5 Intake	OK	29 degrees C / 84 degrees F
FPC 5 Exhaust A	OK	33 degrees C / 91 degrees F
FPC 5 Exhaust B	OK	39 degrees C / 102 degrees F
FPC 5 LU 0 TSen	OK	40 degrees C / 104 degrees F
FPC 5 LU 0 Chip	OK	40 degrees C / 104 degrees F
FPC 5 LU 1 TSen	OK	40 degrees C / 104 degrees F
FPC 5 LU 1 Chip	OK	45 degrees C / 113 degrees F
FPC 5 LU 2 TSen	OK	40 degrees C / 104 degrees F
FPC 5 LU 2 Chip	OK	40 degrees C / 104 degrees F
FPC 5 LU 3 TSen	OK	40 degrees C / 104 degrees F
FPC 5 LU 3 Chip	OK	46 degrees C / 114 degrees F
FPC 5 MQ 0 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 0 Chip	OK	33 degrees C / 91 degrees F
FPC 5 MQ 1 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 1 Chip	OK	35 degrees C / 95 degrees F
FPC 5 MQ 2 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 2 Chip	OK	32 degrees C / 89 degrees F
FPC 5 MQ 3 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 3 Chip	OK	32 degrees C / 89 degrees F
FPC 9 Intake	OK	25 degrees C / 77 degrees F
FPC 9 Exhaust A	OK	37 degrees C / 98 degrees F
FPC 9 Exhaust B	OK	40 degrees C / 104 degrees F
FPC 9 XL 0 TSen	OK	40 degrees C / 104 degrees F

...

**show chassis environment (MX2010 Router)**

user@host&gt; show chassis environment

Class	Item	Status	Measurement
Temp	PSM 0	OK	7 degrees C / 44 degrees F
	PSM 1	OK	7 degrees C / 44 degrees F
	PSM 2	OK	7 degrees C / 44 degrees F
	PSM 3	OK	6 degrees C / 42 degrees F
	PSM 4	OK	6 degrees C / 42 degrees F
	PSM 5	OK	6 degrees C / 42 degrees F
	PSM 6	OK	6 degrees C / 42 degrees F
	PSM 7	OK	7 degrees C / 44 degrees F
	PSM 8	OK	7 degrees C / 44 degrees F
	PDM 0	OK	
	PDM 1	Absent	
	CB 0 IntakeA-Zone0	OK	14 degrees C / 57 degrees F
	CB 0 IntakeB-Zone1	OK	7 degrees C / 44 degrees F
	CB 0 IntakeC-Zone0	OK	22 degrees C / 71 degrees F
	CB 0 ExhaustA-Zone0	OK	14 degrees C / 57 degrees F
	CB 0 ExhaustB-Zone1	OK	9 degrees C / 48 degrees F
	CB 0 TCBC-Zone0	OK	11 degrees C / 51 degrees F
	CB 1 IntakeA-Zone0	OK	9 degrees C / 48 degrees F
	CB 1 IntakeB-Zone1	OK	5 degrees C / 41 degrees F
	CB 1 IntakeC-Zone0	OK	20 degrees C / 68 degrees F
	CB 1 ExhaustA-Zone0	OK	12 degrees C / 53 degrees F
	CB 1 ExhaustB-Zone1	OK	7 degrees C / 44 degrees F



CB 1 TCBC-Zone0	OK	10 degrees C / 50 degrees F
SPMB 0 Intake	OK	5 degrees C / 41 degrees F
SPMB 1 Intake	OK	4 degrees C / 39 degrees F
Routing Engine 0	OK	9 degrees C / 48 degrees F
Routing Engine 0 CPU	OK	9 degrees C / 48 degrees F
Routing Engine 1	OK	6 degrees C / 42 degrees F
Routing Engine 1 CPU	OK	6 degrees C / 42 degrees F
SFB 0 Intake-Zone0	OK	26 degrees C / 78 degrees F
SFB 0 Exhaust-Zone1	OK	17 degrees C / 62 degrees F
SFB 0 IntakeA-Zone0	OK	16 degrees C / 60 degrees F
SFB 0 IntakeB-Zone1	OK	11 degrees C / 51 degrees F
SFB 0 Exhaust-Zone0	OK	18 degrees C / 64 degrees F
SFB 0 SFB-XF2-Zone1	OK	25 degrees C / 77 degrees F
SFB 0 SFB-XF1-Zone0	OK	23 degrees C / 73 degrees F
SFB 0 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
SFB 1 Intake-Zone0	OK	27 degrees C / 80 degrees F
SFB 1 Exhaust-Zone1	OK	15 degrees C / 59 degrees F
SFB 1 IntakeA-Zone0	OK	20 degrees C / 68 degrees F
SFB 1 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 1 Exhaust-Zone0	OK	19 degrees C / 66 degrees F
SFB 1 SFB-XF2-Zone1	OK	26 degrees C / 78 degrees F
SFB 1 SFB-XF1-Zone0	OK	27 degrees C / 80 degrees F
SFB 1 SFB-XF0-Zone0	OK	32 degrees C / 89 degrees F
SFB 2 Intake-Zone0	OK	21 degrees C / 69 degrees F
SFB 2 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 2 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 2 IntakeB-Zone1	OK	9 degrees C / 48 degrees F
SFB 2 Exhaust-Zone0	OK	16 degrees C / 60 degrees F
SFB 2 SFB-XF2-Zone1	OK	24 degrees C / 75 degrees F
SFB 2 SFB-XF1-Zone0	OK	21 degrees C / 69 degrees F
SFB 2 SFB-XF0-Zone0	OK	26 degrees C / 78 degrees F
SFB 4 Intake-Zone0	OK	28 degrees C / 82 degrees F
SFB 4 Exhaust-Zone1	OK	16 degrees C / 60 degrees F
SFB 4 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 4 IntakeB-Zone1	OK	11 degrees C / 51 degrees F
SFB 4 Exhaust-Zone0	OK	19 degrees C / 66 degrees F
SFB 4 SFB-XF2-Zone1	OK	27 degrees C / 80 degrees F
SFB 4 SFB-XF1-Zone0	OK	27 degrees C / 80 degrees F
SFB 4 SFB-XF0-Zone0	OK	32 degrees C / 89 degrees F
SFB 5 Intake-Zone0	OK	22 degrees C / 71 degrees F
SFB 5 Exhaust-Zone1	OK	14 degrees C / 57 degrees F
SFB 5 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 5 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 5 Exhaust-Zone0	OK	17 degrees C / 62 degrees F
SFB 5 SFB-XF2-Zone1	OK	22 degrees C / 71 degrees F
SFB 5 SFB-XF1-Zone0	OK	29 degrees C / 84 degrees F
SFB 5 SFB-XF0-Zone0	OK	27 degrees C / 80 degrees F
SFB 6 Intake-Zone0	OK	27 degrees C / 80 degrees F
SFB 6 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 6 IntakeA-Zone0	OK	19 degrees C / 66 degrees F
SFB 6 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 6 Exhaust-Zone0	OK	20 degrees C / 68 degrees F
SFB 6 SFB-XF2-Zone1	OK	24 degrees C / 75 degrees F
SFB 6 SFB-XF1-Zone0	OK	32 degrees C / 89 degrees F
SFB 6 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
SFB 7 Intake-Zone0	OK	25 degrees C / 77 degrees F
SFB 7 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 7 IntakeA-Zone0	OK	14 degrees C / 57 degrees F
SFB 7 IntakeB-Zone1	OK	8 degrees C / 46 degrees F
SFB 7 Exhaust-Zone0	OK	17 degrees C / 62 degrees F
SFB 7 SFB-XF2-Zone1	OK	21 degrees C / 69 degrees F

SFB 7 SFB-XF1-Zone0	OK	21 degrees C / 69 degrees F
SFB 7 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
FPC 0 Intake	OK	13 degrees C / 55 degrees F
FPC 0 Exhaust A	OK	13 degrees C / 55 degrees F
FPC 0 Exhaust B	OK	14 degrees C / 57 degrees F
FPC 0 LU 0 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 0 Chip	OK	25 degrees C / 77 degrees F
FPC 0 LU 1 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 0 LU 2 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 0 LU 3 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 3 Chip	OK	23 degrees C / 73 degrees F
FPC 0 XM 0 TSen	OK	28 degrees C / 82 degrees F
FPC 0 XM 0 Chip	OK	33 degrees C / 91 degrees F
FPC 0 XM 1 TSen	OK	28 degrees C / 82 degrees F
FPC 0 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 0 PLX Switch TSen	OK	28 degrees C / 82 degrees F
FPC 0 PLX Switch Chip	OK	26 degrees C / 78 degrees F
FPC 1 Intake	OK	10 degrees C / 50 degrees F
FPC 1 Exhaust A	OK	24 degrees C / 75 degrees F
FPC 1 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 1 LU 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 0 Chip	OK	31 degrees C / 87 degrees F
FPC 1 LU 1 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 1 Chip	OK	21 degrees C / 69 degrees F
FPC 1 LU 2 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 2 Chip	OK	25 degrees C / 77 degrees F
FPC 1 LU 3 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 1 XM 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 XM 0 Chip	OK	30 degrees C / 86 degrees F
FPC 1 XF 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 XF 0 Chip	OK	37 degrees C / 98 degrees F
FPC 1 PLX Switch TSen	OK	22 degrees C / 71 degrees F
FPC 1 PLX Switch Chip	OK	22 degrees C / 71 degrees F
FPC 2 Intake	OK	9 degrees C / 48 degrees F
FPC 2 Exhaust A	OK	10 degrees C / 50 degrees F
FPC 2 Exhaust B	OK	10 degrees C / 50 degrees F
FPC 2 LU 0 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 0 Chip	OK	25 degrees C / 77 degrees F
FPC 2 LU 1 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 1 Chip	OK	26 degrees C / 78 degrees F
FPC 2 LU 2 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 2 Chip	OK	17 degrees C / 62 degrees F
FPC 2 LU 3 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 2 XM 0 TSen	OK	26 degrees C / 78 degrees F
FPC 2 XM 0 Chip	OK	34 degrees C / 93 degrees F
FPC 2 XM 1 TSen	OK	26 degrees C / 78 degrees F
FPC 2 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 2 PLX Switch TSen	OK	26 degrees C / 78 degrees F
FPC 2 PLX Switch Chip	OK	20 degrees C / 68 degrees F
FPC 3 Intake	OK	12 degrees C / 53 degrees F
FPC 3 Exhaust A	OK	16 degrees C / 60 degrees F
FPC 3 Exhaust B	OK	26 degrees C / 78 degrees F
FPC 3 LU 0 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 0 Chip	OK	26 degrees C / 78 degrees F
FPC 3 LU 1 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 3 LU 2 TSen	OK	23 degrees C / 73 degrees F

FPC 3 LU 2 Chip	OK	22 degrees C / 71 degrees F
FPC 3 LU 3 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 3 Chip	OK	21 degrees C / 69 degrees F
FPC 3 MQ 0 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 0 Chip	OK	18 degrees C / 64 degrees F
FPC 3 MQ 1 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 1 Chip	OK	20 degrees C / 68 degrees F
FPC 3 MQ 2 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 2 Chip	OK	17 degrees C / 62 degrees F
FPC 3 MQ 3 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 3 Chip	OK	16 degrees C / 60 degrees F
FPC 4 Intake	OK	11 degrees C / 51 degrees F
FPC 4 Exhaust A	OK	22 degrees C / 71 degrees F
FPC 4 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 4 LU 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 0 Chip	OK	33 degrees C / 91 degrees F
FPC 4 LU 1 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 1 Chip	OK	21 degrees C / 69 degrees F
FPC 4 LU 2 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 2 Chip	OK	26 degrees C / 78 degrees F
FPC 4 LU 3 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 4 XM 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 XM 0 Chip	OK	30 degrees C / 86 degrees F
FPC 4 XF 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 XF 0 Chip	OK	37 degrees C / 98 degrees F
FPC 4 PLX Switch TSen	OK	22 degrees C / 71 degrees F
FPC 4 PLX Switch Chip	OK	23 degrees C / 73 degrees F
FPC 5 Intake	OK	12 degrees C / 53 degrees F
FPC 5 Exhaust A	OK	12 degrees C / 53 degrees F
FPC 5 Exhaust B	OK	12 degrees C / 53 degrees F
FPC 5 LU 0 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 0 Chip	OK	28 degrees C / 82 degrees F
FPC 5 LU 1 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 5 LU 2 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 5 LU 3 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 5 XM 0 TSen	OK	27 degrees C / 80 degrees F
FPC 5 XM 0 Chip	OK	36 degrees C / 96 degrees F
FPC 5 XM 1 TSen	OK	27 degrees C / 80 degrees F
FPC 5 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 5 PLX Switch TSen	OK	27 degrees C / 80 degrees F
FPC 5 PLX Switch Chip	OK	24 degrees C / 75 degrees F
FPC 6 Intake	OK	12 degrees C / 53 degrees F
FPC 6 Exhaust A	OK	17 degrees C / 62 degrees F
FPC 6 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 6 LU 0 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 0 Chip	OK	29 degrees C / 84 degrees F
FPC 6 LU 1 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 1 Chip	OK	30 degrees C / 86 degrees F
FPC 6 LU 2 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 2 Chip	OK	24 degrees C / 75 degrees F
FPC 6 LU 3 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 6 MQ 0 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 0 Chip	OK	19 degrees C / 66 degrees F
FPC 6 MQ 1 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 1 Chip	OK	20 degrees C / 68 degrees F
FPC 6 MQ 2 TSen	OK	16 degrees C / 60 degrees F

FPC 6 MQ 2 Chip	OK	17 degrees C / 62 degrees F
FPC 6 MQ 3 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 3 Chip	OK	16 degrees C / 60 degrees F
FPC 7 Intake	OK	10 degrees C / 50 degrees F
FPC 7 Exhaust A	OK	10 degrees C / 50 degrees F
FPC 7 Exhaust B	OK	11 degrees C / 51 degrees F
FPC 7 LU 0 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 0 Chip	OK	26 degrees C / 78 degrees F
FPC 7 LU 1 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 1 Chip	OK	29 degrees C / 84 degrees F
FPC 7 LU 2 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 7 LU 3 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 3 Chip	OK	24 degrees C / 75 degrees F
FPC 7 XM 0 TSen	OK	26 degrees C / 78 degrees F
FPC 7 XM 0 Chip	OK	34 degrees C / 93 degrees F
FPC 7 XM 1 TSen	OK	26 degrees C / 78 degrees F
FPC 7 XM 1 Chip	OK	32 degrees C / 89 degrees F
FPC 7 PLX Switch TSen	OK	26 degrees C / 78 degrees F
FPC 7 PLX Switch Chip	OK	22 degrees C / 71 degrees F
FPC 8 Intake	OK	10 degrees C / 50 degrees F
FPC 8 Exhaust A	OK	22 degrees C / 71 degrees F
FPC 8 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 8 LU 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 0 Chip	OK	33 degrees C / 91 degrees F
FPC 8 LU 1 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 1 Chip	OK	23 degrees C / 73 degrees F
FPC 8 LU 2 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 2 Chip	OK	26 degrees C / 78 degrees F
FPC 8 LU 3 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 8 XM 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 XM 0 Chip	OK	29 degrees C / 84 degrees F
FPC 8 XF 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 XF 0 Chip	OK	38 degrees C / 100 degrees F
FPC 8 PLX Switch TSen	OK	20 degrees C / 68 degrees F
FPC 8 PLX Switch Chip	OK	24 degrees C / 75 degrees F
FPC 9 Intake	OK	11 degrees C / 51 degrees F
FPC 9 Exhaust A	OK	11 degrees C / 51 degrees F
FPC 9 Exhaust B	OK	11 degrees C / 51 degrees F
FPC 9 LU 0 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 0 Chip	OK	24 degrees C / 75 degrees F
FPC 9 LU 1 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 1 Chip	OK	26 degrees C / 78 degrees F
FPC 9 LU 2 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 2 Chip	OK	16 degrees C / 60 degrees F
FPC 9 LU 3 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 3 Chip	OK	21 degrees C / 69 degrees F
FPC 9 XM 0 TSen	OK	25 degrees C / 77 degrees F
FPC 9 XM 0 Chip	OK	32 degrees C / 89 degrees F
FPC 9 XM 1 TSen	OK	25 degrees C / 77 degrees F
FPC 9 XM 1 Chip	OK	25 degrees C / 77 degrees F
FPC 9 PLX Switch TSen	OK	25 degrees C / 77 degrees F
FPC 9 PLX Switch Chip	OK	21 degrees C / 69 degrees F
ADC 0 Intake	OK	12 degrees C / 53 degrees F
ADC 0 Exhaust	OK	20 degrees C / 68 degrees F
ADC 0 ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 0 ADC-XF0	OK	32 degrees C / 89 degrees F
ADC 1 Intake	OK	11 degrees C / 51 degrees F
ADC 1 Exhaust	OK	21 degrees C / 69 degrees F
ADC 1 ADC-XF1	OK	24 degrees C / 75 degrees F

ADC 1 ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 2 Intake	OK	14 degrees C / 57 degrees F
ADC 2 Exhaust	OK	21 degrees C / 69 degrees F
ADC 2 ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 2 ADC-XF0	OK	34 degrees C / 93 degrees F
ADC 3 Intake	OK	13 degrees C / 55 degrees F
ADC 3 Exhaust	OK	19 degrees C / 66 degrees F
ADC 3 ADC-XF1	OK	24 degrees C / 75 degrees F
ADC 3 ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 4 Intake	OK	9 degrees C / 48 degrees F
ADC 4 Exhaust	OK	22 degrees C / 71 degrees F
ADC 4 ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 4 ADC-XF0	OK	35 degrees C / 95 degrees F
ADC 5 Intake	OK	12 degrees C / 53 degrees F
ADC 5 Exhaust	OK	22 degrees C / 71 degrees F
ADC 5 ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 5 ADC-XF0	OK	34 degrees C / 93 degrees F
ADC 6 Intake	OK	11 degrees C / 51 degrees F
ADC 6 Exhaust	OK	21 degrees C / 69 degrees F
ADC 6 ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 6 ADC-XF0	OK	35 degrees C / 95 degrees F
ADC 7 Intake	OK	14 degrees C / 57 degrees F
ADC 7 Exhaust	OK	22 degrees C / 71 degrees F
ADC 7 ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 7 ADC-XF0	OK	34 degrees C / 93 degrees F
ADC 8 Intake	OK	14 degrees C / 57 degrees F
ADC 8 Exhaust	OK	21 degrees C / 69 degrees F
ADC 8 ADC-XF1	OK	24 degrees C / 75 degrees F
ADC 8 ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 9 Intake	OK	10 degrees C / 50 degrees F
ADC 9 Exhaust	OK	22 degrees C / 71 degrees F
ADC 9 ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 9 ADC-XF0	OK	36 degrees C / 96 degrees F
Fans Fan Tray 0 Fan 1	OK	3480 RPM
Fan Tray 0 Fan 2	OK	3480 RPM
Fan Tray 0 Fan 3	OK	3480 RPM
Fan Tray 0 Fan 4	OK	3360 RPM
Fan Tray 0 Fan 5	OK	3360 RPM
Fan Tray 0 Fan 6	OK	3480 RPM
Fan Tray 1 Fan 1	OK	3360 RPM
Fan Tray 1 Fan 2	OK	3360 RPM
Fan Tray 1 Fan 3	OK	3360 RPM
Fan Tray 1 Fan 4	OK	3480 RPM
Fan Tray 1 Fan 5	OK	3480 RPM
Fan Tray 1 Fan 6	OK	3480 RPM
Fan Tray 2 Fan 1	OK	3360 RPM
Fan Tray 2 Fan 2	OK	3360 RPM
Fan Tray 2 Fan 3	OK	3480 RPM
Fan Tray 2 Fan 4	OK	3480 RPM
Fan Tray 2 Fan 5	OK	3360 RPM
Fan Tray 2 Fan 6	OK	3480 RPM
Fan Tray 3 Fan 1	OK	3360 RPM
Fan Tray 3 Fan 2	OK	3360 RPM
Fan Tray 3 Fan 3	OK	3480 RPM
Fan Tray 3 Fan 4	OK	3480 RPM
Fan Tray 3 Fan 5	OK	3480 RPM
Fan Tray 3 Fan 6	OK	3360 RPM

### show chassis environment (T320 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	
	PEM 1	Absent	
Temp	SCG 0	OK	28 degrees C / 82 degrees F
	SCG 1	OK	28 degrees C / 82 degrees F
	Routing Engine 0	OK	31 degrees C / 87 degrees F
	Routing Engine 1	OK	30 degrees C / 86 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	33 degrees C / 91 degrees F
	SIB 1	OK	33 degrees C / 91 degrees F
	SIB 2	OK	34 degrees C / 93 degrees F
	FPC 0 Top	OK	38 degrees C / 100 degrees F
	FPC 0 Bottom	OK	32 degrees C / 89 degrees F
	FPC 1 Top	OK	38 degrees C / 100 degrees F
	FPC 1 Bottom	OK	33 degrees C / 91 degrees F
	FPC 2 Top	OK	36 degrees C / 96 degrees F
	FPC 2 Bottom	OK	31 degrees C / 87 degrees F
	FPM GBUS	OK	26 degrees C / 78 degrees F
	FPM Display	OK	29 degrees C / 84 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Middle fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
Misc	Rear Tray Bottom fan	OK	Spinning at normal speed
	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

### show chassis environment (T640 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	22 degrees C / 71 degrees F
	SCG 0	OK	30 degrees C / 86 degrees F
	SCG 1	OK	30 degrees C / 86 degrees F
	Routing Engine 0	Present	
	Routing Engine 1	OK	27 degrees C / 80 degrees F
	CB 0	Present	
	CB 1	OK	33 degrees C / 91 degrees F
	SIB 0	Absent	
	SIB 1	Absent	
	SIB 2	Absent	
	SIB 3	Absent	
	SIB 4	Absent	
	FPC 4 Top	Testing	
	FPC 4 Bottom	Testing	

	FPC 5 Top	Testing	
	FPC 5 Bottom	Testing	
	FPC 6 Top	Testing	
	FPC 6 Bottom	Testing	
	FPM GBUS	OK	23 degrees C / 73 degrees F
	FPM Display	Absent	
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Fourth Blower from top	OK	Spinning at normal speed
	Bottom Blower	OK	Spinning at normal speed
	Middle Blower	OK	Spinning at normal speed
	Top Blower	OK	Spinning at normal speed
	Second Blower from top	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

#### show chassis environment (T4000 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	33 degrees C / 91 degrees F
	PEM 1	Absent	
	SCG 0	OK	33 degrees C / 91 degrees F
	SCG 1	OK	33 degrees C / 91 degrees F
	Routing Engine 0	OK	33 degrees C / 91 degrees F
	Routing Engine 0 CPU	OK	50 degrees C / 122 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU	OK	46 degrees C / 114 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	33 degrees C / 91 degrees F
	SIB 0	OK	42 degrees C / 107 degrees F
	SIB 1	OK	42 degrees C / 107 degrees F
	SIB 2	OK	42 degrees C / 107 degrees F
	SIB 3	OK	43 degrees C / 109 degrees F
	SIB 4	OK	45 degrees C / 113 degrees F
	FPC 0 Fan Intake	OK	34 degrees C / 93 degrees F
	FPC 0 Fan Exhaust	OK	48 degrees C / 118 degrees F
	FPC 0 PMB	OK	47 degrees C / 116 degrees F
	FPC 0 LMB0	OK	50 degrees C / 122 degrees F
	FPC 0 LMB1	OK	41 degrees C / 105 degrees F
	FPC 0 LMB2	OK	35 degrees C / 95 degrees F
	FPC 0 PFE1 LU2	OK	46 degrees C / 114 degrees F
	FPC 0 PFE1 LU0	OK	41 degrees C / 105 degrees F
	FPC 0 PFE0 LU0	OK	57 degrees C / 134 degrees F
	FPC 0 XF1	OK	46 degrees C / 114 degrees F
	FPC 0 XF0	OK	52 degrees C / 125 degrees F
	FPC 0 XM1	OK	41 degrees C / 105 degrees F
	FPC 0 XM0	OK	50 degrees C / 122 degrees F
	FPC 0 PFE0 LU1	OK	56 degrees C / 132 degrees F

	FPC 0 PFE0 LU2	OK	45 degrees C / 113 degrees F
	FPC 0 PFE1 LU1	OK	37 degrees C / 98 degrees F
	FPC 3 Fan Intake	OK	36 degrees C / 96 degrees F
	FPC 3 Fan Exhaust	OK	51 degrees C / 123 degrees F
	FPC 3 PMB	OK	43 degrees C / 109 degrees F
	FPC 3 LMB0	OK	57 degrees C / 134 degrees F
	FPC 3 LMB1	OK	54 degrees C / 129 degrees F
	FPC 3 LMB2	OK	38 degrees C / 100 degrees F
	FPC 3 PFE1 LU2	OK	63 degrees C / 145 degrees F
	FPC 3 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 3 PFE0 LU0	OK	69 degrees C / 156 degrees F
	FPC 3 XF1	OK	62 degrees C / 143 degrees F
	FPC 3 XF0	OK	63 degrees C / 145 degrees F
	FPC 3 XM1	OK	43 degrees C / 109 degrees F
	FPC 3 XM0	OK	67 degrees C / 152 degrees F
	FPC 3 PFE0 LU1	OK	63 degrees C / 145 degrees F
	FPC 3 PFE0 LU2	OK	66 degrees C / 150 degrees F
	FPC 3 PFE1 LU1	OK	41 degrees C / 105 degrees F
	FPC 5 Top	OK	39 degrees C / 102 degrees F
	FPC 5 Bottom	OK	38 degrees C / 100 degrees F
	FPC 6 Fan Intake	OK	33 degrees C / 91 degrees F
	FPC 6 Fan Exhaust	OK	49 degrees C / 120 degrees F
	FPC 6 PMB	OK	40 degrees C / 104 degrees F
	FPC 6 LMB0	OK	60 degrees C / 140 degrees F
	FPC 6 LMB1	OK	58 degrees C / 136 degrees F
	FPC 6 LMB2	OK	40 degrees C / 104 degrees F
	FPC 6 PFE1 LU2	OK	69 degrees C / 156 degrees F
	FPC 6 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 6 PFE0 LU0	OK	71 degrees C / 159 degrees F
	FPC 6 XF1	OK	58 degrees C / 136 degrees F
	FPC 6 XF0	OK	65 degrees C / 149 degrees F
	FPC 6 XM1	OK	39 degrees C / 102 degrees F
	FPC 6 XM0	OK	66 degrees C / 150 degrees F
	FPC 6 PFE0 LU1	OK	69 degrees C / 156 degrees F
	FPC 6 PFE0 LU2	OK	69 degrees C / 156 degrees F
	FPC 6 PFE1 LU1	OK	42 degrees C / 107 degrees F
	FPM GBUS	OK	24 degrees C / 75 degrees F
	FPM Display	OK	27 degrees C / 80 degrees F
Fans	Top Left Front fan	OK	Spinning at high speed
	Top Left Middle fan	OK	Spinning at high speed
	Top Left Rear fan	OK	Spinning at high speed
	Top Right Front fan	OK	Spinning at high speed
	Top Right Middle fan	OK	Spinning at high speed
	Top Right Rear fan	OK	Spinning at high speed
	Bottom Left Front fan	OK	Spinning at high speed
	Bottom Left Middle fan	OK	Spinning at high speed
	Bottom Left Rear fan	OK	Spinning at high speed
	Bottom Right Front fan	OK	Spinning at high speed
	Bottom Right Middle fan	OK	Spinning at high speed
	Bottom Right Rear fan	OK	Spinning at high speed
	Rear Tray Top fan	OK	Spinning at high speed
	Rear Tray Second fan	OK	Spinning at high speed
	Rear Tray Third fan	OK	Spinning at high speed
	Rear Tray Fourth fan	OK	Spinning at high speed
Misc	Rear Tray Fifth fan	OK	Spinning at high speed
	Rear Tray Sixth fan	OK	Spinning at high speed
	Rear Tray Seventh fan	OK	Spinning at high speed
	Rear Tray Bottom fan	OK	Spinning at high speed
	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	



## show chassis environment (TX Matrix Router)

```
user@host> show chassis environment
scc-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	29 degrees C / 84 degrees F
	Routing Engine 0	OK	34 degrees C / 93 degrees F
	Routing Engine 1	OK	34 degrees C / 93 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	44 degrees C / 111 degrees F
	SIB 0 (B)	OK	44 degrees C / 111 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	32 degrees C / 89 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP 0	OK	
	CIP 1	OK	
	SPMB 0	OK	
	SPMB 1	OK	

```
1cc0-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	29 degrees C / 84 degrees F
	PEM 1	Absent	
	SCG 0	OK	35 degrees C / 95 degrees F
	SCG 1	Absent	
	Routing Engine 0	OK	39 degrees C / 102 degrees F
	Routing Engine 1	OK	36 degrees C / 96 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	40 degrees C / 104 degrees F
	SIB 0 (B)	OK	51 degrees C / 123 degrees F
	FPC 0 Top	OK	45 degrees C / 113 degrees F
	FPC 0 Bottom	OK	31 degrees C / 87 degrees F
	FPC 1 Top	OK	34 degrees C / 93 degrees F
	FPC 1 Bottom	OK	31 degrees C / 87 degrees F
	FPM GBUS	OK	30 degrees C / 86 degrees F
	FPM Display	OK	34 degrees C / 93 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed

	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

lcc2-re0:

Class	Item	Status	Measurement
Temp	PEM 0	OK	29 degrees C / 84 degrees F
	PEM 1	Absent	
	SCG 0	OK	32 degrees C / 89 degrees F
	SCG 1	Absent	
	Routing Engine 0	OK	31 degrees C / 87 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	CB 0	OK	30 degrees C / 86 degrees F
	SIB 0	OK	38 degrees C / 100 degrees F
	SIB 0 (B)	OK	49 degrees C / 120 degrees F
	FPC 0 Top	OK	45 degrees C / 113 degrees F
Fans	FPC 0 Bottom	OK	33 degrees C / 91 degrees F
	FPC 1 Top	OK	37 degrees C / 98 degrees F
	FPC 1 Bottom	OK	33 degrees C / 91 degrees F
	FPM GBUS	OK	30 degrees C / 86 degrees F
	FPM Display	OK	34 degrees C / 93 degrees F
	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	...		

### show chassis environment (T1600 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	27 degrees C / 80 degrees F
	PEM 1	Absent	
	SCG 0	OK	31 degrees C / 87 degrees F
	SCG 1	OK	35 degrees C / 95 degrees F
	Routing Engine 0	OK	30 degrees C / 86 degrees F
	Routing Engine 1	OK	30 degrees C / 86 degrees F
	CB 0	OK	31 degrees C / 87 degrees F
	CB 1	OK	31 degrees C / 87 degrees F
	SIB 0	OK	41 degrees C / 105 degrees F
	SIB 0 (B)	OK	34 degrees C / 93 degrees F
	SIB 1	OK	0 degrees C / 32 degrees F
	SIB 1 (B)	OK	0 degrees C / 32 degrees F

	SIB 2	OK	0 degrees C / 32 degrees F
	SIB 2 (B)	OK	0 degrees C / 32 degrees F
	SIB 3	OK	0 degrees C / 32 degrees F
	SIB 3 (B)	OK	0 degrees C / 32 degrees F
	SIB 4	OK	0 degrees C / 32 degrees F
	SIB 4 (B)	OK	0 degrees C / 32 degrees F
	FPC 0 Top	OK	49 degrees C / 120 degrees F
	FPC 0 Bottom	OK	50 degrees C / 122 degrees F
	FPC 1 Top	OK	48 degrees C / 118 degrees F
	FPC 1 Bottom	OK	49 degrees C / 120 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	30 degrees C / 86 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

### show chassis environment (TX Matrix Plus Router)

```
user@host> show chassis environment
sfc0-re0:
```

Class Item		Status	Measurement
Temp	PEM 0	OK	28 degrees C / 82 degrees F
	PEM 1	Absent	
	Routing Engine 0	OK	27 degrees C / 80 degrees F
	Routing Engine 1	OK	29 degrees C / 84 degrees F
	CB 0 Intake	OK	26 degrees C / 78 degrees F
	CB 0 Exhaust A	OK	25 degrees C / 77 degrees F
	CB 0 Exhaust B	OK	25 degrees C / 77 degrees F
	CB 1 Intake	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust A	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust B	OK	26 degrees C / 78 degrees F
	SIB F13 0	OK	47 degrees C / 116 degrees F
	SIB F13 0 (B)	OK	48 degrees C / 118 degrees F
	SIB F13 1	OK	38 degrees C / 100 degrees F
	SIB F13 1 (B)	OK	37 degrees C / 98 degrees F
	SIB F2S 0/0	OK	27 degrees C / 80 degrees F
	SIB F2S 0/2	OK	28 degrees C / 82 degrees F
	SIB F2S 0/4	OK	27 degrees C / 80 degrees F
	SIB F2S 0/6	OK	28 degrees C / 82 degrees F
	SIB F2S 1/0	OK	26 degrees C / 78 degrees F

	SIB F2S 1/2	OK	26 degrees C / 78 degrees F
	SIB F2S 1/4	OK	26 degrees C / 78 degrees F
	SIB F2S 1/6	OK	26 degrees C / 78 degrees F
	SIB F2S 2/0	OK	25 degrees C / 77 degrees F
	SIB F2S 2/2	OK	25 degrees C / 77 degrees F
	SIB F2S 2/4	OK	23 degrees C / 73 degrees F
	CIP 0 Intake	OK	23 degrees C / 73 degrees F
	CIP 0 Exhaust A	OK	24 degrees C / 75 degrees F
	CIP 0 Exhaust B	OK	24 degrees C / 75 degrees F
	CIP 1 Intake	OK	24 degrees C / 75 degrees F
	CIP 1 Exhaust A	OK	25 degrees C / 77 degrees F
	CIP 1 Exhaust B	OK	25 degrees C / 77 degrees F
Fans	Fan Tray 0 Fan 1	OK	Spinning at normal speed
	Fan Tray 0 Fan 2	OK	Spinning at normal speed
	Fan Tray 0 Fan 3	OK	Spinning at normal speed
	Fan Tray 0 Fan 4	OK	Spinning at normal speed
	Fan Tray 0 Fan 5	OK	Spinning at normal speed
	Fan Tray 0 Fan 6	OK	Spinning at normal speed
	Fan Tray 1 Fan 1	OK	Spinning at normal speed
	Fan Tray 1 Fan 2	OK	Spinning at normal speed
	Fan Tray 1 Fan 3	OK	Spinning at normal speed
	Fan Tray 1 Fan 4	OK	Spinning at normal speed
	Fan Tray 1 Fan 5	OK	Spinning at normal speed
	Fan Tray 1 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 1	OK	Spinning at normal speed
	Fan Tray 2 Fan 2	OK	Spinning at normal speed
	Fan Tray 2 Fan 3	OK	Spinning at normal speed
	Fan Tray 2 Fan 4	OK	Spinning at normal speed
	Fan Tray 2 Fan 5	OK	Spinning at normal speed
	Fan Tray 2 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 7	OK	Spinning at normal speed
	Fan Tray 2 Fan 8	OK	Spinning at normal speed
	Fan Tray 2 Fan 9	OK	Spinning at normal speed
	Fan Tray 3 Fan 1	OK	Spinning at normal speed
	Fan Tray 3 Fan 2	OK	Spinning at normal speed
	Fan Tray 3 Fan 3	OK	Spinning at normal speed
	Fan Tray 3 Fan 4	OK	Spinning at normal speed
	Fan Tray 3 Fan 5	OK	Spinning at normal speed
	Fan Tray 3 Fan 6	OK	Spinning at normal speed
	Fan Tray 3 Fan 7	OK	Spinning at normal speed
	Fan Tray 3 Fan 8	OK	Spinning at normal speed
	Fan Tray 3 Fan 9	OK	Spinning at normal speed
	Fan Tray 4 Fan 1	OK	Spinning at normal speed
	Fan Tray 4 Fan 2	OK	Spinning at normal speed
	Fan Tray 4 Fan 3	OK	Spinning at normal speed
	Fan Tray 4 Fan 4	OK	Spinning at normal speed
	Fan Tray 4 Fan 5	OK	Spinning at normal speed
	Fan Tray 4 Fan 6	OK	Spinning at normal speed
	Fan Tray 4 Fan 7	OK	Spinning at normal speed
	Fan Tray 4 Fan 8	OK	Spinning at normal speed
	Fan Tray 4 Fan 9	OK	Spinning at normal speed
	Fan Tray 5 Fan 1	OK	Spinning at normal speed
	Fan Tray 5 Fan 2	OK	Spinning at normal speed
	Fan Tray 5 Fan 3	OK	Spinning at normal speed
	Fan Tray 5 Fan 4	OK	Spinning at normal speed
	Fan Tray 5 Fan 5	OK	Spinning at normal speed
	Fan Tray 5 Fan 6	OK	Spinning at normal speed
	Fan Tray 5 Fan 7	OK	Spinning at normal speed
	Fan Tray 5 Fan 8	OK	Spinning at normal speed
	Fan Tray 5 Fan 9	OK	Spinning at normal speed
Misc	SPMB 0	OK	

```

SPMB 1                                OK

lcc0-re0:
-----
Class Item                               Status Measurement
Temp PEM 0                              OK          27 degrees C / 80 degrees F
    PEM 1                              Absent
    SCG 0                              OK          31 degrees C / 87 degrees F
    SCG 1                              OK          35 degrees C / 95 degrees F
    Routing Engine 0                    OK          30 degrees C / 86 degrees F
    Routing Engine 1                    OK          30 degrees C / 86 degrees F
    CB 0                                OK          31 degrees C / 87 degrees F
    CB 1                                OK          31 degrees C / 87 degrees F
    SIB 0                                OK          41 degrees C / 105 degrees F
    SIB 0 (B)                           OK          34 degrees C / 93 degrees F
    SIB 1                                OK          0 degrees C / 32 degrees F
    SIB 1 (B)                           OK          0 degrees C / 32 degrees F
    SIB 2                                OK          0 degrees C / 32 degrees F
    SIB 2 (B)                           OK          0 degrees C / 32 degrees F
    SIB 3                                OK          0 degrees C / 32 degrees F
    SIB 3 (B)                           OK          0 degrees C / 32 degrees F
    SIB 4                                OK          0 degrees C / 32 degrees F
    SIB 4 (B)                           OK          0 degrees C / 32 degrees F
    FPC 0 Top                           OK          49 degrees C / 120 degrees F
    FPC 0 Bottom                        OK          50 degrees C / 122 degrees F
    FPC 1 Top                           OK          48 degrees C / 118 degrees F
    FPC 1 Bottom                        OK          49 degrees C / 120 degrees F
    FPM GBUS                            OK          27 degrees C / 80 degrees F
    FPM Display                         OK          30 degrees C / 86 degrees F
Fans Top Left Front fan                 OK          Spinning at normal speed
    Top Left Middle fan                 OK          Spinning at normal speed
    Top Left Rear fan                   OK          Spinning at normal speed
    Top Right Front fan                 OK          Spinning at normal speed
    Top Right Middle fan                OK          Spinning at normal speed
    Top Right Rear fan                  OK          Spinning at normal speed
    Bottom Left Front fan                OK          Spinning at normal speed
    Bottom Left Middle fan               OK          Spinning at normal speed
    Bottom Left Rear fan                 OK          Spinning at normal speed
    Bottom Right Front fan               OK          Spinning at normal speed
    Bottom Right Middle fan              OK          Spinning at normal speed
    Bottom Right Rear fan                OK          Spinning at normal speed
    Rear Tray Top fan                   OK          Spinning at normal speed
    Rear Tray Second fan                 OK          Spinning at normal speed
    Rear Tray Third fan                  OK          Spinning at normal speed
    Rear Tray Fourth fan                 OK          Spinning at normal speed
    Rear Tray Fifth fan                  OK          Spinning at normal speed
    Rear Tray Sixth fan                  OK          Spinning at normal speed
    Rear Tray Seventh fan                OK          Spinning at normal speed
    Rear Tray Bottom fan                 OK          Spinning at normal speed
Misc CIP                               OK
    SPMB 0                              OK
    SPMB 1                              OK

```

#### show chassis environment (TX Matrix Plus router with 3D SIBs)

```

user@host> show chassis environment
sfc0-re0:
-----
Class Item                               Status Measurement
Temp PEM 0                              Check       30 degrees C / 86 degrees F
    PEM 1                              OK          33 degrees C / 91 degrees F

```

	Routing Engine 0	OK	28 degrees C / 82 degrees F
	Routing Engine 0 CPU	OK	42 degrees C / 107 degrees F
	Routing Engine 1	OK	29 degrees C / 84 degrees F
	Routing Engine 1 CPU	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	30 degrees C / 86 degrees F
	CB 0 Exhaust A	OK	28 degrees C / 82 degrees F
	CB 0 Exhaust B	OK	30 degrees C / 86 degrees F
	CB 1 Intake	OK	31 degrees C / 87 degrees F
	CB 1 Exhaust A	OK	27 degrees C / 80 degrees F
	CB 1 Exhaust B	OK	31 degrees C / 87 degrees F
	SIB F13 0 Board	OK	44 degrees C / 111 degrees F
	SIB F13 0 XF Junction	OK	62 degrees C / 143 degrees F
	SIB F13 3 Board	OK	45 degrees C / 113 degrees F
	SIB F13 3 XF Junction	OK	60 degrees C / 140 degrees F
	SIB F13 6 Board	OK	47 degrees C / 116 degrees F
	SIB F13 6 XF Junction	OK	62 degrees C / 143 degrees F
	SIB F2S 0/0 Board	OK	32 degrees C / 89 degrees F
	SIB F2S 0/0 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 0/2 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/2 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 0/4 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/4 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 0/6 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/6 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 1/0 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 1/0 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 1/2 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 1/2 XF Junction	OK	39 degrees C / 102 degrees F
	SIB F2S 1/4 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 1/4 XF Junction	OK	35 degrees C / 95 degrees F
	SIB F2S 1/6 Board	OK	30 degrees C / 86 degrees F
	SIB F2S 1/6 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 2/0 Board	OK	30 degrees C / 86 degrees F
	SIB F2S 2/0 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 2/2 Board	OK	28 degrees C / 82 degrees F
	SIB F2S 2/2 XF Junction	OK	39 degrees C / 102 degrees F
	SIB F2S 2/4 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 2/4 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 2/6 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 2/6 XF Junction	OK	41 degrees C / 105 degrees F
	CIP 0 Intake	OK	25 degrees C / 77 degrees F
	CIP 0 Exhaust A	OK	26 degrees C / 78 degrees F
	CIP 0 Exhaust B	OK	26 degrees C / 78 degrees F
	CIP 1 Intake	OK	26 degrees C / 78 degrees F
	CIP 1 Exhaust A	OK	27 degrees C / 80 degrees F
	CIP 1 Exhaust B	OK	27 degrees C / 80 degrees F
Fans	Fan Tray 0 Fan 1	OK	Spinning at normal speed
	Fan Tray 0 Fan 2	OK	Spinning at normal speed
	Fan Tray 0 Fan 3	OK	Spinning at normal speed
	Fan Tray 0 Fan 4	OK	Spinning at normal speed
	Fan Tray 0 Fan 5	OK	Spinning at normal speed
	Fan Tray 0 Fan 6	OK	Spinning at normal speed
	Fan Tray 1 Fan 1	OK	Spinning at normal speed
	Fan Tray 1 Fan 2	OK	Spinning at normal speed
	Fan Tray 1 Fan 3	OK	Spinning at normal speed
	Fan Tray 1 Fan 4	OK	Spinning at normal speed
	Fan Tray 1 Fan 5	OK	Spinning at normal speed
	Fan Tray 1 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 1	OK	Spinning at normal speed
	Fan Tray 2 Fan 2	OK	Spinning at normal speed
	Fan Tray 2 Fan 3	OK	Spinning at normal speed

Fan Tray 2 Fan 4	OK	Spinning at normal speed
Fan Tray 2 Fan 5	OK	Spinning at normal speed
Fan Tray 2 Fan 6	OK	Spinning at normal speed
Fan Tray 2 Fan 7	OK	Spinning at normal speed
Fan Tray 2 Fan 8	OK	Spinning at normal speed
Fan Tray 2 Fan 9	OK	Spinning at normal speed
Fan Tray 3 Fan 1	OK	Spinning at normal speed
Fan Tray 3 Fan 2	OK	Spinning at normal speed
Fan Tray 3 Fan 3	OK	Spinning at normal speed
Fan Tray 3 Fan 4	OK	Spinning at normal speed
Fan Tray 3 Fan 5	OK	Spinning at normal speed
Fan Tray 3 Fan 6	OK	Spinning at normal speed
Fan Tray 3 Fan 7	OK	Spinning at normal speed
Fan Tray 3 Fan 8	OK	Spinning at normal speed
Fan Tray 3 Fan 9	OK	Spinning at normal speed
Fan Tray 4 Fan 1	OK	Spinning at normal speed
Fan Tray 4 Fan 2	OK	Spinning at normal speed
Fan Tray 4 Fan 3	OK	Spinning at normal speed
Fan Tray 4 Fan 4	OK	Spinning at normal speed
Fan Tray 4 Fan 5	OK	Spinning at normal speed
Fan Tray 4 Fan 6	OK	Spinning at normal speed
Fan Tray 4 Fan 7	OK	Spinning at normal speed
Fan Tray 4 Fan 8	OK	Spinning at normal speed
Fan Tray 4 Fan 9	OK	Spinning at normal speed
Fan Tray 5 Fan 1	OK	Spinning at normal speed
Fan Tray 5 Fan 2	OK	Spinning at normal speed
Fan Tray 5 Fan 3	OK	Spinning at normal speed
Fan Tray 5 Fan 4	OK	Spinning at normal speed
Fan Tray 5 Fan 5	OK	Spinning at normal speed
Fan Tray 5 Fan 6	OK	Spinning at normal speed
Fan Tray 5 Fan 7	OK	Spinning at normal speed
Fan Tray 5 Fan 8	OK	Spinning at normal speed
Fan Tray 5 Fan 9	Check	
Misc SPMB 0	OK	
SPMB 1	OK	

1cc0-re0:

Class	Item	Status	Measurement
Temp	PEM 0	OK	29 degrees C / 84 degrees F
	PEM 1	Check	29 degrees C / 84 degrees F
	SCG 0	OK	32 degrees C / 89 degrees F
	SCG 1	OK	33 degrees C / 91 degrees F
	Routing Engine 0	OK	32 degrees C / 89 degrees F
	Routing Engine 0 CPU	OK	51 degrees C / 123 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU	OK	49 degrees C / 120 degrees F
	CB 0	OK	34 degrees C / 93 degrees F
	CB 1	OK	34 degrees C / 93 degrees F
	SIB 0	OK	39 degrees C / 102 degrees F
	SIB 0 (B)	Absent	
	SIB 1	OK	39 degrees C / 102 degrees F
	SIB 1 (B)	Absent	
	SIB 2	OK	39 degrees C / 102 degrees F
	SIB 2 (B)	Absent	
	FPC 4 Top	OK	43 degrees C / 109 degrees F
	FPC 4 Bottom	OK	43 degrees C / 109 degrees F
	FPC 7 Fan Intake	OK	35 degrees C / 95 degrees F
	FPC 7 Fan Exhaust	OK	50 degrees C / 122 degrees F
	FPC 7 PMB	OK	50 degrees C / 122 degrees F
	FPC 7 LMB0	OK	55 degrees C / 131 degrees F

	FPC 7 LMB1	OK	49 degrees C / 120 degrees F
	FPC 7 LMB2	OK	39 degrees C / 102 degrees F
	FPC 7 PFE1 LU2	OK	55 degrees C / 131 degrees F
	FPC 7 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 7 PFE0 LU0	OK	62 degrees C / 143 degrees F
	FPC 7 XF1	OK	52 degrees C / 125 degrees F
	FPC 7 XF0	OK	61 degrees C / 141 degrees F
	FPC 7 XM1	OK	39 degrees C / 102 degrees F
	FPC 7 XM0	OK	56 degrees C / 132 degrees F
	FPC 7 PFE0 LU1	OK	60 degrees C / 140 degrees F
	FPC 7 PFE0 LU2	OK	55 degrees C / 131 degrees F
	FPC 7 PFE1 LU1	OK	41 degrees C / 105 degrees F
	FPM GBUS	OK	24 degrees C / 75 degrees F
	FPM Display	OK	28 degrees C / 82 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray fan 1 (Top)	OK	Spinning at normal speed
	Rear Tray fan 2	OK	Spinning at normal speed
	Rear Tray fan 3	OK	Spinning at normal speed
	Rear Tray fan 4	OK	Spinning at normal speed
	Rear Tray fan 5	OK	Spinning at normal speed
	Rear Tray fan 6	OK	Spinning at normal speed
	Rear Tray fan 7	OK	Spinning at normal speed
	Rear Tray fan 8	OK	Spinning at normal speed
	Rear Tray fan 9	OK	Spinning at normal speed
	Rear Tray fan 10	OK	Spinning at normal speed
	Rear Tray fan 11	OK	Spinning at normal speed
	Rear Tray fan 12	OK	Spinning at normal speed
	Rear Tray fan 13	OK	Spinning at normal speed
	Rear Tray fan 14	OK	Spinning at normal speed
	Rear Tray fan 15	OK	Spinning at normal speed
	Rear Tray fan 16 (Bottom)	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

### show chassis environment (EX4200 Standalone Switch)

user@switch> show chassis environment			
Class	Item	Status	Measurement
Power	FPC 0 Power Supply 0	OK	
	FPC 0 Power Supply 1	Absent	
Temp	FPC 0 CPU	OK	41 degrees C / 105 degrees F
	FPC 0 EX-PFE1	OK	42 degrees C / 107 degrees F
	FPC 0 EX-PFE2	OK	46 degrees C / 114 degrees F
	FPC 0 GEPHY Front Left	OK	25 degrees C / 77 degrees F
	FPC 0 GEPHY Front Right	OK	27 degrees C / 80 degrees F
	FPC 0 Uplink Conn	OK	29 degrees C / 84 degrees F
Fans	FPC 0 Fan 1	OK	Spinning at normal speed
	FPC 0 Fan 2	OK	Spinning at normal speed
	FPC 0 Fan 3	OK	Spinning at normal speed



## show chassis environment (EX8216 Switch)

```

user@switch> show chassis environment
Class Item                               Status      Measurement
Power PSU 0                             OK
      PSU 1                             OK
      PSU 2                             OK
      PSU 3                             Check
      PSU 4                             Absent
      PSU 5                             Absent
Temp  CB 0 Intake                         OK          23 degrees C / 73 degrees F
      CB 0 Exhaust                       OK          26 degrees C / 78 degrees F
      CB 1 Intake                         OK          22 degrees C / 71 degrees F
      CB 1 Exhaust                       OK          25 degrees C / 77 degrees F
      FPC 4 Intake                       OK          49 degrees C / 120 degrees F
      FPC 4 Exhaust                     OK          59 degrees C / 138 degrees F
      SIB 5 Intake                       OK          25 degrees C / 77 degrees F
      SIB 5 Exhaust                     OK          35 degrees C / 95 degrees F
      SIB 6 Intake                       OK          25 degrees C / 77 degrees F
      SIB 6 Exhaust                     OK          38 degrees C / 100 degrees F
Fans  Top Fan 1                         OK          Spinning at normal speed
      Top Fan 2                         OK          Spinning at normal speed
      Top Fan 3                         OK          Spinning at normal speed
      Top Fan 4                         OK          Spinning at normal speed
      Top Fan 5                         OK          Spinning at normal speed
      Top Fan 6                         OK          Spinning at normal speed
      Top Fan 7                         OK          Spinning at normal speed
      Top Fan 8                         OK          Spinning at normal speed
      Top Fan 9                         OK          Spinning at normal speed
      Bottom Fan 1                     OK          Spinning at normal speed
      Bottom Fan 2                     OK          Spinning at normal speed
      Bottom Fan 3                     OK          Spinning at normal speed
      Bottom Fan 4                     OK          Spinning at normal speed
      Bottom Fan 5                     OK          Spinning at normal speed
      Bottom Fan 6                     OK          Spinning at normal speed
      Bottom Fan 7                     OK          Spinning at normal speed
      Bottom Fan 8                     OK          Spinning at normal speed
      Bottom Fan 9                     OK          Spinning at normal speed

```

## show chassis environment (EX9200 Switch)

```

user@switch> show chassis environment
Class Item                               Status      Measurement
Temp PEM 0                             Check
      PEM 1                             OK          40 degrees C / 104 degrees F
      PEM 2                             OK          40 degrees C / 104 degrees F
      PEM 3                             Absent
      Routing Engine 0                 OK          35 degrees C / 95 degrees F
      Routing Engine 0 CPU              OK          33 degrees C / 91 degrees F
      Routing Engine 1                 OK          38 degrees C / 100 degrees F
      Routing Engine 1 CPU              OK          33 degrees C / 91 degrees F
      CB 0 Intake                       OK          35 degrees C / 95 degrees F
      CB 0 Exhaust A                    OK          33 degrees C / 91 degrees F
      CB 0 Exhaust B                    OK          40 degrees C / 104 degrees F
      CB 0 ACBC                         OK          39 degrees C / 102 degrees F
      CB 0 XF A                         OK          49 degrees C / 120 degrees F
      CB 0 XF B                         OK          46 degrees C / 114 degrees F
      CB 1 Intake                       OK          37 degrees C / 98 degrees F
      CB 1 Exhaust A                    OK          32 degrees C / 89 degrees F
      CB 1 Exhaust B                    OK          39 degrees C / 102 degrees F
      CB 1 ACBC                         OK          41 degrees C / 105 degrees F

```

CB 1 XF A	OK	49 degrees C / 120 degrees F
CB 1 XF B	OK	49 degrees C / 120 degrees F
FPC 2 Intake	OK	37 degrees C / 98 degrees F
FPC 2 Exhaust A	OK	40 degrees C / 104 degrees F
FPC 2 Exhaust B	OK	34 degrees C / 93 degrees F
FPC 2 LU 0 TCAM TSen	OK	44 degrees C / 111 degrees F
FPC 2 LU 0 TCAM Chip	OK	48 degrees C / 118 degrees F
FPC 2 LU 0 TSen	OK	44 degrees C / 111 degrees F
FPC 2 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 2 MQ 0 TSen	OK	44 degrees C / 111 degrees F
FPC 2 MQ 0 Chip	OK	51 degrees C / 123 degrees F
FPC 3 Intake	OK	39 degrees C / 102 degrees F
FPC 3 Exhaust A	OK	51 degrees C / 123 degrees F

[...Output truncated...]

Fans	Top Rear Fan	OK	Spinning at intermediate-speed
	Bottom Rear Fan	OK	Spinning at intermediate-speed
	Top Middle Fan	OK	Spinning at intermediate-speed
	Bottom Middle Fan	OK	Spinning at intermediate-speed
	Top Front Fan	OK	Spinning at intermediate-speed
	Bottom Front Fan	OK	Spinning at intermediate-speed

**show chassis environment (QFX Series)**

```
user@switch> show chassis environment
```

Class	Item	Status	Measurement
Power	FPC 0 Power Supply 0	OK	
	FPC 0 Power Supply 1	OK	
Temp	FPC 0 Sensor TopLeft I	OK	26 degrees C / 78 degrees F
	FPC 0 Sensor TopRight I	OK	24 degrees C / 75 degrees F
	FPC 0 Sensor TopLeft E	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopRight E	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopMiddle I	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopMiddle E	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Bottom I	OK	34 degrees C / 93 degrees F
	FPC 0 Sensor Bottom E	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Die Temp	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Mgmt Brd I	OK	24 degrees C / 75 degrees F
	FPC 0 Sensor Switch I	OK	28 degrees C / 82 degrees F
Fans	FPC 0 Fan 1 (left)	Failed	
	FPC 0 Fan 2 (right)	OK	Spinning at normal speed
	FPC 0 Fan 3 (middle)	OK	Spinning at normal speed

**show chassis environment interconnect-device (QFabric System)**

```
user@switch> show chassis environment interconnect-device IC-A0004
```

Class	Item	Status	Measurement
	CB 0		
	CB 0 L Intake	OK	30 degrees C / 86 degrees F
	CB 0 R Intake	OK	31 degrees C / 87 degrees F
	CB 0 L Exhaust	OK	32 degrees C / 89 degrees F
	CB 0 R Exhaust	OK	33 degrees C / 91 degrees F
	Routing Engine 0 CPU temp	OK	51 degrees C / 123 degrees F
	CB 1		
	CB 1 L Intake	OK	27 degrees C / 80 degrees F
	CB 1 R Intake	OK	29 degrees C / 84 degrees F
	CB 1 L Exhaust	OK	31 degrees C / 87 degrees F
	CB 1 R Exhaust	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU temp	OK	40 degrees C / 104 degrees F
	FC 0 FPC 0		

FPC 0 L Intake	OK	25 degrees C / 77 degrees F
FPC 0 R Intake	OK	28 degrees C / 82 degrees F
FPC 0 L Exhaust	OK	28 degrees C / 82 degrees F
FPC 0 R Exhaust	OK	29 degrees C / 84 degrees F
FC 7 FPC 7		
FPC 7 L Intake	OK	25 degrees C / 77 degrees F
FPC 7 R Intake	OK	26 degrees C / 78 degrees F
FPC 7 L Exhaust	OK	28 degrees C / 82 degrees F
FPC 7 R Exhaust	OK	29 degrees C / 84 degrees F
RC 0 FPC 8		
FPC 8 L Intake	OK	25 degrees C / 77 degrees F
FPC 8 R Intake	OK	26 degrees C / 78 degrees F
FPC 8 L Exhaust	OK	32 degrees C / 89 degrees F
FPC 8 R Exhaust	OK	30 degrees C / 86 degrees F
RC 7 FPC 15		
FPC 15 L Intake	OK	24 degrees C / 75 degrees F
FPC 15 R Intake	OK	25 degrees C / 77 degrees F
FPC 15 L Exhaust	OK	33 degrees C / 91 degrees F
FPC 15 R Exhaust	OK	31 degrees C / 87 degrees F
Fans TFT 0 Fan 0	OK	Spinning at normal speed
Fans TFT 0 Fan 1	OK	Spinning at normal speed
Fans TFT 0 Fan 2	OK	Spinning at normal speed
Fans TFT 0 Fan 3	OK	Spinning at normal speed
Fans TFT 0 Fan 4	OK	Spinning at normal speed
Fans TFT 0 Fan 5	OK	Spinning at normal speed
Fans BFT 1 Fan 0	OK	Spinning at normal speed
Fans BFT 1 Fan 1	OK	Spinning at normal speed
Fans BFT 1 Fan 2	OK	Spinning at normal speed
Fans BFT 1 Fan 3	Check	
Fans BFT 1 Fan 4	OK	Spinning at normal speed
Fans BFT 1 Fan 5	OK	Spinning at normal speed
Fans SFT 0 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 3 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 3 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 3 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 3 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 3 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 3 Fan 2 Rotor 1	OK	Spinning at normal speed

Fans	SFT 3	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 3	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 3	Rotor 1	OK	Spinning at normal speed
Power	PEM 0			OK	30 degrees C / 86 degrees F
Power	PEM 1			OK	30 degrees C / 86 degrees F
Power	PEM 2			OK	30 degrees C / 86 degrees F
Power	PEM 3			Absent	
Power	PEM 4			Absent	
Power	PEM 5			Absent	

#### show chassis environment node-device (QFabric System)

```

user@switch> show chassis environment node-device node1
Class Item                               Status Measurement
Power node1 Power Supply 0              Absent
      node1 Power Supply 1              Absent
Fans  node1 Fan Tray 0                  Testing
      node1 Fan Tray 1                  Testing
      node1 Fan Tray 2                  Testing

```

#### show chassis environment pem node-device (QFabric System)

```

user@switch> show chassis environment pem node-device node1
FPC 0 PEM 0 status:
State          Check
Airflow        Front to Back
Temperature    OK
AC Input:      OK
DC Output      Voltage(V) Current(A) Power(W) Load(%)
                12         10        120      18

```

```

FPC 0 PEM 1 status:
State                Online
Airflow              Back to Front
Temperature           OK
AC Input:            OK
DC Output             Voltage(V) Current(A) Power(W) Load(%)
                    11          10       110     17

```

#### show chassis environment (PTX5000 Packet Transport Router)

```

user@host> show chassis environment
Class Item                Status      Measurement
Temp PDU 0                OK
      PDU 0 PSM 0         OK          36 degrees C / 96 degrees F
      PDU 0 PSM 1         OK          38 degrees C / 100 degrees F
      PDU 0 PSM 2         OK          38 degrees C / 100 degrees F
      PDU 0 PSM 3         OK          37 degrees C / 98 degrees F
      PDU 1                Absent
      CCG 0                OK          44 degrees C / 111 degrees F
      CCG 1                OK          44 degrees C / 111 degrees F
      Routing Engine 0     OK          62 degrees C / 143 degrees F
      Routing Engine 0 CPU OK          75 degrees C / 167 degrees F
      Routing Engine 1     OK          51 degrees C / 123 degrees F
      Routing Engine 1 CPU OK          64 degrees C / 147 degrees F
      CB 0 Intake          OK          38 degrees C / 100 degrees F
      CB 0 Exhaust A       OK          46 degrees C / 114 degrees F
      CB 0 Exhaust B       OK          42 degrees C / 107 degrees F
      CB 1 Intake          OK          35 degrees C / 95 degrees F
      CB 1 Exhaust A       OK          39 degrees C / 102 degrees F
      CB 1 Exhaust B       OK          36 degrees C / 96 degrees F
      SIB 0 Exhaust        OK          47 degrees C / 116 degrees F
      SIB 0 Junction       OK          45 degrees C / 113 degrees F
      SIB 1 Exhaust        OK          44 degrees C / 111 degrees F
      SIB 1 Junction       OK          43 degrees C / 109 degrees F
      SIB 2 Exhaust        OK          47 degrees C / 116 degrees F
      SIB 2 Junction       OK          42 degrees C / 107 degrees F
      SIB 3 Exhaust        OK          43 degrees C / 109 degrees F
      SIB 3 Junction       OK          43 degrees C / 109 degrees F
      SIB 4 Exhaust        OK          47 degrees C / 116 degrees F
      SIB 4 Junction       OK          42 degrees C / 107 degrees F
      SIB 5 Exhaust        OK          42 degrees C / 107 degrees F
      SIB 5 Junction       OK          40 degrees C / 104 degrees F
      SIB 6 Exhaust        OK          46 degrees C / 114 degrees F
      SIB 6 Junction       OK          42 degrees C / 107 degrees F
      SIB 7 Exhaust        OK          43 degrees C / 109 degrees F
      SIB 7 Junction       OK          39 degrees C / 102 degrees F
      SIB 8 Exhaust        OK          44 degrees C / 111 degrees F
      SIB 8 Junction       OK          41 degrees C / 105 degrees F
      FPC 0 PMB            OK          35 degrees C / 95 degrees F
      FPC 0 Intake         OK          33 degrees C / 91 degrees F
      FPC 0 Exhaust A      OK          51 degrees C / 123 degrees F
      FPC 0 Exhaust B      OK          43 degrees C / 109 degrees F
      FPC 0 TL0            OK          48 degrees C / 118 degrees F
      FPC 0 TQ0            OK          53 degrees C / 127 degrees F
      FPC 0 TL1            OK          56 degrees C / 132 degrees F
      FPC 0 TQ1            OK          58 degrees C / 136 degrees F
      FPC 0 TL2            OK          55 degrees C / 131 degrees F
      FPC 0 TQ2            OK          56 degrees C / 132 degrees F
      FPC 0 TL3            OK          59 degrees C / 138 degrees F
      FPC 0 TQ3            OK          59 degrees C / 138 degrees F
      FPC 2 PMB            OK          35 degrees C / 95 degrees F

```

FPC 2 Intake	OK	34 degrees C / 93 degrees F
FPC 2 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 2 TL0	OK	53 degrees C / 127 degrees F
FPC 2 TQ0	OK	53 degrees C / 127 degrees F
FPC 2 TL1	OK	57 degrees C / 134 degrees F
FPC 2 TQ1	OK	58 degrees C / 136 degrees F
FPC 2 TL2	OK	54 degrees C / 129 degrees F
FPC 2 TQ2	OK	59 degrees C / 138 degrees F
FPC 2 TL3	OK	60 degrees C / 140 degrees F
FPC 2 TQ3	OK	64 degrees C / 147 degrees F
PIC 2/0 Ambient	OK	49 degrees C / 120 degrees F
FPC 3 PMB	OK	34 degrees C / 93 degrees F
FPC 3 Intake	OK	35 degrees C / 95 degrees F
FPC 3 Exhaust A	OK	54 degrees C / 129 degrees F
FPC 3 Exhaust B	OK	49 degrees C / 120 degrees F
FPC 3 TL0	OK	49 degrees C / 120 degrees F
FPC 3 TQ0	OK	55 degrees C / 131 degrees F
FPC 3 TL1	OK	56 degrees C / 132 degrees F
FPC 3 TQ1	OK	58 degrees C / 136 degrees F
FPC 3 TL2	OK	56 degrees C / 132 degrees F
FPC 3 TQ2	OK	59 degrees C / 138 degrees F
FPC 3 TL3	OK	62 degrees C / 143 degrees F
FPC 3 TQ3	OK	63 degrees C / 145 degrees F
PIC 3/1	Absent	
FPC 5 PMB	OK	35 degrees C / 95 degrees F
FPC 5 Intake	OK	34 degrees C / 93 degrees F
FPC 5 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 5 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 5 TL0	OK	54 degrees C / 129 degrees F
FPC 5 TQ0	OK	52 degrees C / 125 degrees F
FPC 5 TL1	OK	61 degrees C / 141 degrees F
FPC 5 TQ1	OK	60 degrees C / 140 degrees F
FPC 5 TL2	OK	55 degrees C / 131 degrees F
FPC 5 TQ2	OK	55 degrees C / 131 degrees F
FPC 5 TL3	OK	59 degrees C / 138 degrees F
FPC 5 TQ3	OK	58 degrees C / 136 degrees F
PIC 5/0 Ambient	OK	51 degrees C / 123 degrees F
PIC 5/1 Ambient	OK	34 degrees C / 93 degrees F
PIC 5/1 cfp-5/1/0	OK	34 degrees C / 93 degrees F
PIC 5/1 cfp-5/1/1	OK	36 degrees C / 96 degrees F
FPC 6 PMB	OK	36 degrees C / 96 degrees F
FPC 6 Intake	OK	33 degrees C / 91 degrees F
FPC 6 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 6 Exhaust B	OK	39 degrees C / 102 degrees F
FPC 6 TL0	OK	44 degrees C / 111 degrees F
FPC 6 TQ0	OK	54 degrees C / 129 degrees F
FPC 6 TL1	OK	59 degrees C / 138 degrees F
FPC 6 TQ1	OK	58 degrees C / 136 degrees F
FPC 6 TL2	OK	60 degrees C / 140 degrees F
FPC 6 TQ2	OK	57 degrees C / 134 degrees F
FPC 6 TL3	OK	65 degrees C / 149 degrees F
FPC 6 TQ3	OK	60 degrees C / 140 degrees F
FPC 7 PMB	OK	35 degrees C / 95 degrees F
FPC 7 Intake	OK	33 degrees C / 91 degrees F
FPC 7 Exhaust A	OK	53 degrees C / 127 degrees F
FPC 7 Exhaust B	OK	40 degrees C / 104 degrees F
FPC 7 TL0	OK	46 degrees C / 114 degrees F
FPC 7 TQ0	OK	58 degrees C / 136 degrees F
FPC 7 TL1	OK	53 degrees C / 127 degrees F
FPC 7 TQ1	OK	59 degrees C / 138 degrees F

	FPC 7 TL2	OK	56 degrees C / 132 degrees F
	FPC 7 TQ2	OK	61 degrees C / 141 degrees F
	FPC 7 TL3	OK	63 degrees C / 145 degrees F
	FPC 7 TQ3	OK	63 degrees C / 145 degrees F
	FPM I2CS	OK	37 degrees C / 98 degrees F
Fans	Fan Tray 0 Fan 1	OK	3042 RPM
	Fan Tray 0 Fan 2	OK	3042 RPM
	Fan Tray 0 Fan 3	OK	3000 RPM
	Fan Tray 0 Fan 4	OK	3042 RPM
	Fan Tray 0 Fan 5	OK	3000 RPM
	Fan Tray 0 Fan 6	OK	3042 RPM
	Fan Tray 0 Fan 7	OK	3085 RPM
	Fan Tray 0 Fan 8	OK	3042 RPM
	Fan Tray 0 Fan 9	OK	3042 RPM
	Fan Tray 0 Fan 10	OK	3085 RPM
	Fan Tray 0 Fan 11	OK	3085 RPM
	Fan Tray 0 Fan 12	OK	3128 RPM
	Fan Tray 0 Fan 13	OK	3128 RPM
	Fan Tray 0 Fan 14	OK	3042 RPM
	Fan Tray 1 Fan 1	OK	2299 RPM
	Fan Tray 1 Fan 2	OK	2399 RPM
	Fan Tray 1 Fan 3	OK	2299 RPM
	Fan Tray 1 Fan 4	OK	2266 RPM
	Fan Tray 1 Fan 5	OK	2266 RPM
	Fan Tray 1 Fan 6	OK	2366 RPM
	Fan Tray 2 Fan 1	OK	2199 RPM
	Fan Tray 2 Fan 2	OK	2133 RPM
	Fan Tray 2 Fan 3	OK	2366 RPM
	Fan Tray 2 Fan 4	OK	2233 RPM
	Fan Tray 2 Fan 5	OK	2399 RPM
	Fan Tray 2 Fan 6	OK	2233 RPM
Misc	SPMB 0 Intake	OK	50 degrees C / 122 degrees F
	SPMB 1 Intake	OK	40 degrees C / 104 degrees F

#### show chassis environment (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```

user@host> show chassis environment
Class Item                               Status      Measurement
Temp  PDU 0                               OK
      PDU 0 PSM 0                       OK          41 degrees C / 105 degrees F
      PDU 0 PSM 1                       Absent
      PDU 0 PSM 2                       OK          43 degrees C / 109 degrees F
      PDU 0 PSM 3                       Absent
      PDU 0 PSM 4                       OK          44 degrees C / 111 degrees F
      PDU 0 PSM 5                       Absent
      PDU 0 PSM 6                       OK          45 degrees C / 113 degrees F
      PDU 0 PSM 7                       Absent
      PDU 1                               OK
      PDU 1 PSM 0                       Absent
      PDU 1 PSM 1                       OK          45 degrees C / 113 degrees F
      PDU 1 PSM 2                       Absent
      PDU 1 PSM 3                       OK          43 degrees C / 109 degrees F
      PDU 1 PSM 4                       Absent
      PDU 1 PSM 5                       OK          46 degrees C / 114 degrees F
      PDU 1 PSM 6                       Absent
      PDU 1 PSM 7                       OK          46 degrees C / 114 degrees F
      CCG 0                             OK          27 degrees C / 80 degrees F
      CCG 1                             OK          29 degrees C / 84 degrees F
...

```

**show chassis environment (ACX2000 Universal Access Router)**

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
	PCB Left	OK	44 degrees C / 111 degrees F
	SFP+ Xcvr	OK	50 degrees C / 122 degrees F
	FEB	OK	70 degrees C / 158 degrees F
	PCB Up	OK	63 degrees C / 145 degrees F
	PCB Mid	OK	66 degrees C / 150 degrees F
	Telecom Mod	OK	65 degrees C / 149 degrees F
	Routing Engine	OK	54 degrees C / 129 degrees F
	Heater off		

**show chassis environment (ACX4000 Universal Access Router)**

On the ACX4000 router, the MIC output of the **show chassis environment** command varies depending on the number of temperature channels present in the installed MIC.

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	33 degrees C / 91 degrees F
	PEM 1	Absent	
	PCB Bottom	OK	30 degrees C / 86 degrees F
	PCB Middle	OK	34 degrees C / 93 degrees F
	BCM56445	OK	33 degrees C / 91 degrees F
	SFP+ Xcvr	OK	32 degrees C / 89 degrees F
	Fan tray inlet	OK	39 degrees C / 102 degrees F
	Exhaust	OK	30 degrees C / 86 degrees F
	Routing Engine	OK	32 degrees C / 89 degrees F
	Heater off		
Pic	PIC 0/0 Channel 0	OK	28 degrees C / 82 degrees F
	PIC 0/0 Channel 1	OK	29 degrees C / 84 degrees F
	PIC 0/0 Channel 2	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 3	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 4	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 5	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 6	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 7	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 8	OK	0 degrees C / 32 degrees F
	PIC 0/0 Channel 9	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 0	OK	33 degrees C / 91 degrees F
	PIC 1/0 Channel 1	OK	31 degrees C / 87 degrees F
	PIC 1/0 Channel 2	OK	30 degrees C / 86 degrees F
	PIC 1/0 Channel 3	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 4	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 5	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 6	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 7	OK	0 degrees C / 32 degrees F
	PIC 1/0 Channel 8	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 0	OK	31 degrees C / 87 degrees F
	PIC 1/1 Channel 1	OK	29 degrees C / 84 degrees F
	PIC 1/1 Channel 2	OK	28 degrees C / 82 degrees F
	PIC 1/1 Channel 3	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 4	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 5	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 6	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 7	OK	0 degrees C / 32 degrees F
	PIC 1/1 Channel 8	OK	0 degrees C / 32 degrees F



Fans	Fan 1	OK	Spinning at normal speed
	Fan 2	OK	Spinning at normal speed

## show chassis environment cb

---

<b>List of Syntax</b>	<a href="#">Syntax on page 970</a> <a href="#">Syntax (TX Matrix Routers) on page 970</a> <a href="#">Syntax (TX Matrix Plus Routers) on page 970</a> <a href="#">Syntax (MX Series Routers) on page 970</a> <a href="#">Syntax (MX104 3D Universal Edge Routers) on page 970</a> <a href="#">Syntax (MX2010 and MX2020 3D Universal Edge Routers) on page 970</a> <a href="#">Syntax (QFabric System) on page 970</a>
<b>Syntax</b>	show chassis environment cb <slot>
<b>Syntax (TX Matrix Routers)</b>	show chassis environment cb <lcc number   scc> <slot>
<b>Syntax (TX Matrix Plus Routers)</b>	show chassis environment cb <lcc number   sfc number > <slot>
<b>Syntax (MX Series Routers)</b>	show chassis environment cb <slot> <all-members> <local> <member member-id>
<b>Syntax (MX104 3D Universal Edge Routers)</b>	show chassis environment cb
<b>Syntax (MX2010 and MX2020 3D Universal Edge Routers)</b>	show chassis environment cb <slot>
<b>Syntax (QFabric System)</b>	show chassis environment cb <slot interconnect-device interconnect-device-name> < interconnect-device interconnect-device-name slot>
<b>Release Information</b>	Command introduced before Junos Release 7.4. Command introduced in Junos OS Release 9.4 for EX Series switches. Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers. Command introduced in Junos OS Release 12.1 for T4000 Core Routers. sfc option introduced for the TX Matrix Plus router in Junos Release 9.6. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
<b>Description</b>	(M120, M320, MX Series, and T Series routers, EX8200 switches, and PTX Series Packet Transport Routers only) Display environmental information about the Control Boards

(CBs). For information about the meaning of “CBs” on the switches, see [“EX Series Switches Hardware and CLI Terminology Mapping” on page 315](#).

**Options** **none**—Display environmental information about all CBs. For a TX Matrix router, display environmental information about all CBs on the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display environmental information about all CBs on the TX Matrix Plus router and its attached T1600 or T4000 routers.

**all-members**—(MX Series routers only) (Optional) Display environmental information about the CBs on all the members of the Virtual Chassis configuration.

**interconnect-device**—(QFabric systems only) Display environmental information about CBs on the Interconnect device.

**lcc number**—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display environmental information about the CBs on the local Virtual Chassis member.

**member member-id**—(MX Series routers only) (Optional) Display environmental information about the CBs on the specified member of the Virtual Chassis configuration. Replace **member-id** with a value of 0 or 1.

**scc**—(TX Matrix router only) (Optional) Display environmental information about the CBs in the TX Matrix router (switch-card chassis).

**sfc number**—(TX Matrix Plus router only) (Optional) Display environmental information about the CBs in the TX Matrix Plus router (or switch-fabric chassis).

**slot**—(Optional) Display environmental information about the specified CB. On routers and PTX Series Packet Transport Routers, replace **slot** with 0 or 1. On EX Series switches replace **slot** with 0, 1, or 2. On QFX Series switches, replace **slot** with 0 or 1.

**Required Privilege Level** view

**Related Documentation** • [request chassis cb on page 877](#)

- *Understanding Switching Control Board Redundancy*
- *Routing Engine and Switching Control Board Redundancy Configuration Statements*

List of Sample Output	<a href="#">show chassis environment cb (M120 Router) on page 973</a>
	<a href="#">show chassis environment cb (M320 Router) on page 973</a>
	<a href="#">show chassis environment cb (MX80 Router) on page 974</a>
	<a href="#">show chassis environment cb (MX104 Router) on page 974</a>
	<a href="#">show chassis environment cb (MX240 Router) on page 975</a>
	<a href="#">show chassis environment cb (MX240 Router with Enhanced MX SCB) on page 975</a>
	<a href="#">show chassis environment cb (MX480 Router) on page 975</a>
	<a href="#">show chassis environment cb (MX480 Router with Enhanced MX SCB) on page 976</a>
	<a href="#">show chassis environment cb (MX960 Router) on page 976</a>
	<a href="#">show chassis environment cb (MX960 Router with Enhanced MX SCB) on page 977</a>
	<a href="#">show chassis environment cb (MX2020 Router) on page 977</a>
	<a href="#">show chassis environment cb (MX2010 Router) on page 978</a>
	<a href="#">show chassis environment cb (T4000 Core Router) on page 979</a>
	<a href="#">show chassis environment cb (TX Matrix Router) on page 979</a>
	<a href="#">show chassis environment cb (TX Matrix Plus Router) on page 980</a>
	<a href="#">show chassis environment cb (EX8200 Switch) on page 984</a>
	<a href="#">show chassis environment cb (EX8208 Switch) on page 985</a>
	<a href="#">show chassis environment cb (PTX5000 Packet Transport Router) on page 986</a>
	<a href="#">show chassis environment cb (QFabric System) on page 987</a>

**Output Fields** [Table 107 on page 972](#) lists the output fields for the **show chassis environment cb** command. Output fields are listed in the approximate order in which they appear.

**Table 107: show chassis environment cb Output Fields**

Field Name	Field Description
<b>State</b>	<p>Status of the CB. If two CBs are installed and online, one is functioning as the master, and the other is the standby.</p> <ul style="list-style-type: none"> <li>• <b>Online</b>—CB is online and running.</li> <li>• <b>Offline</b>—CB is powered down.</li> </ul> <p><b>NOTE:</b> On the EX8208 switch, the installation can include three CBs. See <a href="#">“EX Series Switches Hardware and CLI Terminology Mapping” on page 315</a>.</p>
<b>Temperature</b>	<p>Temperature in Celsius (C) and Fahrenheit (F) of the air flowing past the CB.</p> <ul style="list-style-type: none"> <li>• <b>Temperature Intake</b>—Measures the temperature of the air intake to cool the power supplies.</li> <li>• <b>Temperature Exhaust</b>—Measures the temperature of the hot air exhaust.</li> </ul> <p><b>NOTE:</b> On the MX2010 and MX2020 routers, the intake temperature measures the temperature of the air intake to cool the Control Board (CB). The MX2010 and MX2020 routers include intake and exhaust temperatures for multiple zones (<b>Intake A</b>, <b>Intake B</b>, <b>Intake C</b>, <b>Exhaust A</b>, <b>Exhaust B</b>, and <b>TCBC</b>).</p>
<b>Power</b>	<p>Power required and measured on the CB. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.</p>
<b>BUS Revision</b>	<p>Revision level of the generic bus device. (Not on switches.)</p>

Table 107: show chassis environment cb Output Fields (*continued*)

Field Name	Field Description
<b>FPGA Revision</b>	Revision level of the field-programmable gate array (FPGA). (Not on switches.)
<b>PMBus device</b> (on MX240, MX480, and MX960 routers with Enhanced MX SCB)	Enhanced SCB on MX 240, MX480, and MX960 routers allows the system to save power by supplying only the amount of voltage that is required. Configurable PMBus devices are used to provide the voltage for each individual device. There is one PMBus device for each XF ASIC so that the output can be customized to each device. The following PMBus device information is displayed for routers with Enhanced MX SCB: <ul style="list-style-type: none"> <li>• Expected voltage</li> <li>• Measured voltage</li> <li>• Measured current</li> <li>• Calculated power</li> </ul>

## Sample Output

### show chassis environment cb (M120 Router)

```

user@host> show chassis environment cb
CB 0 status:
  State                Online Master
  Temperature          33 degrees C / 91 degrees F
  Power
    1.2 V              1214 mV
    1.5 V              1495 mV
    2.5 V              2494 mV
    3.3 V              3319 mV
    5.0 V              5085 mV
    3.3 V bias         3296 mV
  Bus Revision         12
  FPGA Revision        17
CB 1 status:
  State                Online Standby
  Temperature          34 degrees C / 93 degrees F
  Power
    1.2 V              1195 mV
    1.5 V              1495 mV
    2.5 V              2504 mV
    3.3 V              3312 mV
    5.0 V              5111 mV
    3.3 V bias         3296 mV
  Bus Revision         12
  FPGA Revision        17

```

### show chassis environment cb (M320 Router)

```

user@host> show chassis environment cb
CB 0 status:
  State                Online Master
  Temperature          29 degrees C / 84 degrees F
  Power:
    1.8 V              1805 mV
    2.5 V              2501 mV
    3.3 V              3293 mV
    4.6 V              4725 mV

```

5.0 V	5032 mV
12.0 V	11975 mV
3.3 V bias	3286 mV
8.0 V bias	7589 mV
BUS Revision	40
FPGA Revision	7
CB 1 status:	
State	Online Standby
Temperature	32 degrees C / 89 degrees F
Power:	
1.8 V	1802 mV
2.5 V	2482 mV
3.3 V	3289 mV
4.6 V	4720 mV
5.0 V	5001 mV
12.0 V	11946 mV
3.3 V bias	3274 mV
8.0 V bias	7562 mV
BUS Revision	40
FPGA Revision	7

#### show chassis environment cb (MX80 Router)

```
user@host> show chassis environment cb
```

CB 0 status:	
State	Online Master
Temperature	36 degrees C / 96 degrees F
Power 1	
1.0 V	1034 mV
1.0 V MQ	1037 mV
1.0 V LU	1005 mV
1.2 V	1218 mV
1.5 V	1524 mV
1.8 V	1814 mV
2.5 V	2558 mV
3.3 V	3296 mV
5.0 V	5233 mV
5.0 V bias	5207 mV
12.0 V	12162 mV

#### show chassis environment cb (MX104 Router)

```
user@host > show chassis environment cb
```

CB 0 status:	
State	Online Master
Temperature	33 degrees C / 91 degrees F
Power 1	
0.75 V	751 mV
1.0 V	1005 mV
1.1 V	1113 mV
1.5 V	1494 mV
2.5 V	2518 mV
3.3 V	3338 mV
5.0 V	4960 mV
12.0 V	12006 mV
FPGA Revision	25
CB 1 status:	
State	Empty

**show chassis environment cb (MX240 Router)**

```

user@host> show chassis environment cb
CB 0 status:
State                               Online Standby
Temperature                         37 degrees C / 98 degrees F
Power 1
  1.2 V                             1208 mV
  1.5 V                             1521 mV
  1.8 V                             1811 mV
  2.5 V                             2513 mV
  3.3 V                             3332 mV
  5.0 V                             5059 mV
  12.0 V                             12162 mV
  1.25 V                             1260 mV
  3.3 V SM3                         3306 mV
  5.0 V RE                          5085 mV
  12.0 V RE                         11872 mV
Power 2
  11.3 V bias PEM                   11272 mV
  4.6 V bias MidPlane               4827 mV
  11.3 V bias FPD                   11272 mV
  11.3 V bias POE 0                 11292 mV
  11.3 V bias POE 1                 11253 mV
Bus Revision                         42
FPGA Revision                       1

```

**show chassis environment cb (MX240 Router with Enhanced MX SCB)**

```

user@host> show chassis environment cb
CB 0 status:
State                               Online Standby
Temperature                         37 degrees C / 98 degrees F
Power 1
  1.2 V                             1208 mV
  1.5 V                             1521 mV
  1.8 V                             1811 mV
  2.5 V                             2513 mV
  3.3 V                             3332 mV
  5.0 V                             5059 mV
  12.0 V                             12162 mV
  1.25 V                             1260 mV
  3.3 V SM3                         3306 mV
  5.0 V RE                          5085 mV
  12.0 V RE                         11872 mV
Power 2
  11.3 V bias PEM                   11272 mV
  4.6 V bias MidPlane               4827 mV
  11.3 V bias FPD                   11272 mV
  11.3 V bias POE 0                 11292 mV
  11.3 V bias POE 1                 11253 mV
Bus Revision                         42
FPGA Revision                       1
PMBus                               Expected Measured Measured Calculated
device                             voltage  voltage  current  power
XF ASIC A                         1000 mV   997 mV  11031 mA 10997 mW
XF ASIC B                         1000 mV   996 mV  12125 mA 12076 mW

```

**show chassis environment cb (MX480 Router)**

```

user@host> show chassis environment cb

```

```

CB 0 status:
State                               Online Master
Temperature                         41 degrees C / 105 degrees F
Power 1
  1.2 V                             1202 mV
  1.5 V                             1511 mV
  1.8 V                             1798 mV
  2.5 V                             2507 mV
  3.3 V                             3312 mV
  5.0 V                             5027 mV
  12.0 V                             12200 mV
  1.25 V                             1260 mV
  3.3 V SM3                         3293 mV
  5 V RE                             5040 mV
  12 V RE                             11910 mV
Power 2
  11.3 V bias PEM                   11156 mV
  4.6 V bias MidPlane               4801 mV
  11.3 V bias FPD                   11214 mV
  11.3 V bias POE 0                 11098 mV
  11.3 V bias POE 1                 11330 mV
Bus Revision                         42
FPGA Revision                       1

```

#### show chassis environment cb (MX480 Router with Enhanced MX SCB)

```

user@host> show chassis environment cb
CB 0 status:
State                               Online Master
Temperature                         41 degrees C / 105 degrees F
Power 1
  1.2 V                             1202 mV
  1.5 V                             1511 mV
  1.8 V                             1798 mV
  2.5 V                             2507 mV
  3.3 V                             3312 mV
  5.0 V                             5027 mV
  12.0 V                             12200 mV
  1.25 V                             1260 mV
  3.3 V SM3                         3293 mV
  5 V RE                             5040 mV
  12 V RE                             11910 mV
Power 2
  11.3 V bias PEM                   11156 mV
  4.6 V bias MidPlane               4801 mV
  11.3 V bias FPD                   11214 mV
  11.3 V bias POE 0                 11098 mV
  11.3 V bias POE 1                 11330 mV
Bus Revision                         42
FPGA Revision                       1
PMBus                               Expected Measured Measured Calculated
device                             voltage  voltage  current  power
  XF ASIC A                        1000 mV   997 mV  11031 mA 10997 mW
  XF ASIC B                        1000 mV   996 mV  12125 mA 12076 mW

```

#### show chassis environment cb (MX960 Router)

```

user@host> show chassis environment cb
CB 0 status:
State                               Online Master
Temperature                         24 degrees C / 75 degrees F

```



```

Power 1
  1.2 V          1965 mV
  1.5 V          2465 mV
  1.8 V          2990 mV
  2.5 V          3296 mV
  3.3 V          3296 mV
  5.0 V          6593 mV
 12.0 V        13187 mV
  3.3 V bias     3296 mV
  1.25 V         1994 mV
  3.3 V SM3      3296 mV
  5 V RE         6593 mV
 12 V RE        13174 mV
Power 2          Sensor failure
Bus Revision     4
FPGA Revision    3

```

### show chassis environment cb (MX960 Router with Enhanced MX SCB)

```

user@host> show chassis environment cb
CB 0 status:
  State          Online Master
  Temperature     24 degrees C / 75 degrees F
  Power 1
    1.2 V          1965 mV
    1.5 V          2465 mV
    1.8 V          2990 mV
    2.5 V          3296 mV
    3.3 V          3296 mV
    5.0 V          6593 mV
   12.0 V        13187 mV
    3.3 V bias     3296 mV
    1.25 V         1994 mV
    3.3 V SM3      3296 mV
    5 V RE         6593 mV
   12 V RE        13174 mV
  Power 2          Sensor failure
  Bus Revision     4
  FPGA Revision    3
  PMBus
  device           Expected voltage Measured voltage Measured current Calculated power
  XF ASIC A        1000 mV          997 mV          11031 mA       10997 mW
  XF ASIC B        1000 mV          996 mV          12125 mA       12076 mW

```

### show chassis environment cb (MX2020 Router)

```

user@host> show chassis environment cb
CB 0 status:
  State          Online Master
  IntakeA-Zone0 Temperature 44 degrees C / 111 degrees F
  IntakeB-Zone1 Temperature 34 degrees C / 93 degrees F
  IntakeC-Zone0 Temperature 45 degrees C / 113 degrees F
  ExhaustA-Zone0 Temperature 43 degrees C / 109 degrees F
  ExhaustB-Zone1 Temperature 36 degrees C / 96 degrees F
  TCBC-Zone0 Temperature 39 degrees C / 102 degrees F
  Power 1
    1.0 V          1011 mV
    1.2 V          1208 mV
    1.8 V          1801 mV
    2.5 V          2552 mV
    3.3 V          3312 mV

```

```

5.0 V          5040 mV
5.0 V RE       4988 mV
12.0 V         12065 mV
12.0 V RE      12046 mV
Bus Revision   99
FPGA Revision  270
CB 1 status:
State          Online Standby
IntakeA-Zone0 Temperature 45 degrees C / 113 degrees F
IntakeB-Zone1 Temperature 41 degrees C / 105 degrees F
IntakeC-Zone0 Temperature 46 degrees C / 114 degrees F
ExhaustA-Zone0 Temperature 44 degrees C / 111 degrees F
ExhaustB-Zone1 Temperature 41 degrees C / 105 degrees F
TCBC-Zone0 Temperature 45 degrees C / 113 degrees F
Power 1
1.0 V          1008 mV
1.2 V          1208 mV
1.8 V          1798 mV
2.5 V          2539 mV
3.3 V          3325 mV
5.0 V          5033 mV
5.0 V RE       4950 mV
12.0 V         12046 mV
12.0 V RE      11968 mV
Bus Revision   99
FPGA Revision  0

```

#### show chassis environment cb (MX2010 Router)

```

user@host> show chassis environment cb
CB 0 status:
State          Online Master
IntakeA-Zone0 Temperature 36 degrees C / 96 degrees F
IntakeB-Zone1 Temperature 30 degrees C / 86 degrees F
IntakeC-Zone0 Temperature 38 degrees C / 100 degrees F
ExhaustA-Zone0 Temperature 36 degrees C / 96 degrees F
ExhaustB-Zone1 Temperature 32 degrees C / 89 degrees F
TCBC-Zone0 Temperature 34 degrees C / 93 degrees F
Power 1
1.0 V          1015 mV
1.2 V          1205 mV
1.8 V          1804 mV
2.5 V          2552 mV
3.3 V          3325 mV
5.0 V          5020 mV
5.0 V RE       4988 mV
12.0 V         12104 mV
12.0 V RE      12026 mV
Bus Revision   100
FPGA Revision  270
CB 1 status:
State          Online
IntakeA-Zone0 Temperature 35 degrees C / 95 degrees F
IntakeB-Zone1 Temperature 28 degrees C / 82 degrees F
IntakeC-Zone0 Temperature 37 degrees C / 98 degrees F
ExhaustA-Zone0 Temperature 34 degrees C / 93 degrees F
ExhaustB-Zone1 Temperature 29 degrees C / 84 degrees F
TCBC-Zone0 Temperature 32 degrees C / 89 degrees F
Power 1
1.0 V          1011 mV
1.2 V          1208 mV

```

1.8 V	1788 mV
2.5 V	2526 mV
3.3 V	3319 mV
5.0 V	5046 mV
5.0 V RE	4975 mV
12.0 V	12046 mV
12.0 V RE	12007 mV
Bus Revision	100
FPGA Revision	0

#### show chassis environment cb (T4000 Core Router)

```

user@host> show chassis environment cb
CB 0 status:
  State                Online Master
  Temperature          33 degrees C / 91 degrees F
  Power 1
    1.8 V              1805 mV
    2.5 V              2523 mV
    3.3 V              3324 mV
    3.3 V bias         3296 mV
    4.6 V              4680 mV
    5.0 V              4893 mV
    8.0 V bias         7572 mV
    12.0 V             11916 mV
  Power 2
    1.0 V              993 mV
    1.2 V              1210 mV
    3.3 V RE           3330 mV
  Bus Revision         51
  FPGA Revision        5
CB 1 status:
  State                Online Standby
  Temperature          33 degrees C / 91 degrees F
  Power 1
    1.8 V              1810 mV
    2.5 V              2496 mV
    3.3 V              3308 mV
    3.3 V bias         3286 mV
    4.6 V              4692 mV
    5.0 V              4954 mV
    8.0 V bias         7282 mV
    12.0 V             11926 mV
  Power 2
    1.0 V              993 mV
    1.2 V              1185 mV
    3.3 V RE           3316 mV
  Bus Revision         51
  FPGA Revision        5

```

#### show chassis environment cb (TX Matrix Router)

```

user@host> show chassis environment cb
-----
CB 0 status:
  State                Online Master
  Temperature          32 degrees C / 89 degrees F
  Power:
    1.8 V              1797 mV
    2.5 V              2477 mV
    3.3 V              3311 mV

```

```

4.6 V          4727 mV
5.0 V          5015 mV
12.0 V         12185 mV
3.3 V bias     3304 mV
8.0 V bias     7870 mV
BUS Revision   40
FPGA Revision  1
CB 1 status:
State          Online Standby
...

lcc0-re0:
-----
CB 0 status:
State          Online Master
Temperature     32 degrees C / 89 degrees F
Power:
  1.8 V         1787 mV
  2.5 V         2473 mV
  3.3 V         3306 mV
  4.6 V         4793 mV
  5.0 V         5025 mV
  12.0 V        12156 mV
  3.3 V bias    3289 mV
  8.0 V bias    7609 mV
BUS Revision   40
FPGA Revision  5
CB 1 status:
State          Online Standby
....
BUS Revision   40
FPGA Revision  5

lcc2-re0:
-----
CB 0 status:
State          Online Master
...
CB 1 status:
State          Online Standby
...

```

### show chassis environment cb (TX Matrix Plus Router)

```

user@host> show chassis environment cb
sfc0-re0:
-----
CB 0 status:
State          Online Master
Temperature     38 degrees C / 100 degrees F
Power 1
  1.0 V         1005 mV
  1.1 V         1108 mV
  1.2 V         1205 mV
  1.25 V        1269 mV
  1.5 V         1508 mV
  1.8 V         1814 mV
  2.5 V         2507 mV
  3.3 V         3306 mV
  3.3 V bias    3300 mV
  9.0 V         9058 mV

```

9.0 V RE	9107 mV
Power 2	
3.9 V	3963 mV
5.0 V	5020 mV
9.0 V	9087 mV
Bus Revision	79
FPGA Revision	23
CB 1 status:	
State	Online Standby
Temperature	39 degrees C / 102 degrees F
Power 1	
1.0 V	1002 mV
1.1 V	1105 mV
1.2 V	1198 mV
1.25 V	1276 mV
1.5 V	1504 mV
1.8 V	1804 mV
2.5 V	2507 mV
3.3 V	3300 mV
3.3 V bias	3293 mV
9.0 V	9039 mV
9.0 V RE	9049 mV
Power 2	
3.9 V	3892 mV
5.0 V	5040 mV
9.0 V	9058 mV
Bus Revision	79
FPGA Revision	23

lcc0-re0:

-----

CB 0 status:	
State	Online Master
Temperature	39 degrees C / 102 degrees F
Power 1	
1.8 V	1799 mV
2.5 V	2499 mV
3.3 V	3327 mV
3.3 V bias	3299 mV
4.6 V	4673 mV
5.0 V	4918 mV
8.0 V bias	7308 mV
12.0 V	11887 mV
Power 2	
1.0 V	996 mV
1.2 V	1199 mV
3.3 V RE	3319 mV
Bus Revision	51
FPGA Revision	3
CB 1 status:	
State	Online Standby
Temperature	40 degrees C / 104 degrees F
Power 1	
1.8 V	1800 mV
2.5 V	2496 mV
3.3 V	3322 mV
3.3 V bias	3284 mV
4.6 V	4680 mV
5.0 V	4954 mV
8.0 V bias	7284 mV
12.0 V	11902 mV

```
Power 2
  1.0 V          998 mV
  1.2 V          1205 mV
  3.3 V RE       3327 mV
Bus Revision     51
FPGA Revision    3
```

1cc1-re0:

-----

CB 0 status:

```
State           Online Master
Temperature      41 degrees C / 105 degrees F
Power 1
  1.8 V          1804 mV
  2.5 V          2517 mV
  3.3 V          3300 mV
  3.3 V bias     3284 mV
  4.6 V          4681 mV
  5.0 V          4927 mV
  8.0 V bias     7357 mV
  12.0 V         11907 mV
Power 2
  1.0 V          991 mV
  1.2 V          1202 mV
  3.3 V RE       3301 mV
Bus Revision     51
FPGA Revision    3
```

CB 1 status:

```
State           Online Standby
Temperature      40 degrees C / 104 degrees F
Power 1
  1.8 V          1805 mV
  2.5 V          2528 mV
  3.3 V          3324 mV
  3.3 V bias     3289 mV
  4.6 V          4694 mV
  5.0 V          4959 mV
  8.0 V bias     7311 mV
  12.0 V         11926 mV
Power 2
  1.0 V          998 mV
  1.2 V          1200 mV
  3.3 V RE       3313 mV
Bus Revision     51
FPGA Revision    3
```

1cc2-re0:

-----

CB 0 status:

```
State           Online Master
Temperature      41 degrees C / 105 degrees F
Power 1
  1.8 V          1805 mV
  2.5 V          2494 mV
  3.3 V          3333 mV
  3.3 V bias     3296 mV
  4.6 V          4673 mV
  5.0 V          4901 mV
  8.0 V bias     7343 mV
  12.0 V         11916 mV
Power 2
```

1.0 V	993 mV
1.2 V	1213 mV
3.3 V RE	3328 mV
Bus Revision	51
FPGA Revision	3
CB 1 status:	
State	Online Standby
Temperature	41 degrees C / 105 degrees F
Power 1	
1.8 V	1804 mV
2.5 V	2523 mV
3.3 V	3334 mV
3.3 V bias	3291 mV
4.6 V	4697 mV
5.0 V	4969 mV
8.0 V bias	7308 mV
12.0 V	11936 mV
Power 2	
1.0 V	996 mV
1.2 V	1200 mV
3.3 V RE	3328 mV
Bus Revision	51
FPGA Revision	3

lcc3-re0:

---

CB 0 status:	
State	Online Master
Temperature	37 degrees C / 98 degrees F
Power 1	
1.8 V	1809 mV
2.5 V	2510 mV
3.3 V	3296 mV
3.3 V bias	3291 mV
4.6 V	4670 mV
5.0 V	4905 mV
8.0 V bias	7211 mV
12.0 V	11882 mV
Power 2	
1.0 V	996 mV
1.2 V	1188 mV
3.3 V RE	3326 mV
Bus Revision	51
FPGA Revision	5
CB 1 status:	
State	Online Standby
Temperature	38 degrees C / 100 degrees F
Power 1	
1.8 V	1813 mV
2.5 V	2510 mV
3.3 V	3322 mV
3.3 V bias	3289 mV
4.6 V	4692 mV
5.0 V	4967 mV
8.0 V bias	7194 mV
12.0 V	11916 mV
Power 2	
1.0 V	996 mV
1.2 V	1205 mV
3.3 V RE	3273 mV

```
Bus Revision          51
FPGA Revision         5
```

### show chassis environment cb (EX8200 Switch)

```
user@host> show chassis environment cb
```

#### CB 0 status:

```
State                Online Master
Temperature Intake    20 degrees C / 68 degrees F
Temperature Exhaust   24 degrees C / 75 degrees F
Power 1
  1.1 V              1086 mV
  1.2 V              1179 mV
  1.2 V *            1182 mV
  1.2 V *            1182 mV
  1.25 V             1211 mV
  1.5 V              1472 mV
  1.8 V              1756 mV
  2.5 V              2449 mV
  3.3 V              3254 mV
  3.3 V bias         3300 mV
  5.0 V              4911 mV
  12.0 V             11891 mV
```

#### Power 2

```
3.3 V bias *        3615 mV
3.3 V bias *        3615 mV
3.3 V bias *        3567 mV
3.3 V bias *        3664 mV
4.3 V bias *        4224 mV
4.3 V bias *        4215 mV
4.3 V bias *        4224 mV
4.3 V bias *        4205 mV
4.3 V bias *        4195 mV
4.3 V bias *        4215 mV
5.0 V bias          4920 mV
```

#### CB 1 status:

```
State                Online Standby
Temperature Intake    19 degrees C / 66 degrees F
Temperature Exhaust   23 degrees C / 73 degrees F
Power 1
  1.1 V              1082 mV
  1.2 V              1169 mV
  1.2 V *            1179 mV
  1.2 V *            1179 mV
  1.25 V             1214 mV
  1.5 V              1482 mV
  1.8 V              1759 mV
  2.5 V              2481 mV
  3.3 V              3248 mV
  3.3 V bias         3306 mV
  5.0 V              4911 mV
  12.0 V             11910 mV
```

#### Power 2

```
3.3 V bias *        3644 mV
3.3 V bias *        3664 mV
3.3 V bias *        3586 mV
3.3 V bias *        3654 mV
4.3 V bias *        4224 mV
4.3 V bias *        4215 mV
4.3 V bias *        4224 mV
```



```

4.3 V bias *      4205 mV
4.3 V bias *      4244 mV
4.3 V bias *      4215 mV
5.0 V bias        4930 mV
CB 2 status:
State             Online
Temperature Intake 19 degrees C / 66 degrees F
Temperature Exhaust 23 degrees C / 73 degrees F
Power 1
1.2 V             1195 mV
1.5 V             1511 mV
1.8 V             1804 mV
2.5 V             2526 mV
3.3 V             3300 mV
3.3 V bias        3306 mV
12.0 V            12220 mV

```

### show chassis environment cb (EX8208 Switch)

```

user@host> show chassis environment cb
CB 0 status:
State             Online Master
Temperature Intake 20 degrees C / 68 degrees F
Temperature Exhaust 24 degrees C / 75 degrees F
Power 1
1.1 V             1086 mV
1.2 V             1179 mV
1.2 V *           1182 mV
1.2 V *           1182 mV
1.25 V            1211 mV
1.5 V             1466 mV
1.8 V             1759 mV
2.5 V             2455 mV
3.3 V             3261 mV
3.3 V bias        3300 mV
5.0 V             4930 mV
12.0 V            11891 mV
Power 2
3.3 V bias *      3606 mV
3.3 V bias *      3615 mV
3.3 V bias *      3567 mV
3.3 V bias *      3673 mV
4.3 V bias *      4224 mV
4.3 V bias *      4215 mV
4.3 V bias *      4234 mV
4.3 V bias *      4205 mV
4.3 V bias *      4186 mV
4.3 V bias *      4215 mV
5.0 V bias        4940 mV
CB 1 status:
State             Online Standby
Temperature Intake 19 degrees C / 66 degrees F
Temperature Exhaust 23 degrees C / 73 degrees F
Power 1
1.1 V             1086 mV
1.2 V             1169 mV
1.2 V *           1179 mV
1.2 V *           1179 mV
1.25 V            1211 mV
1.5 V             1479 mV
1.8 V             1759 mV

```

```

2.5 V                2475 mV
3.3 V                3235 mV
3.3 V bias           3306 mV
5.0 V                4930 mV
12.0 V               11891 mV
Power 2
3.3 V bias *         3644 mV
3.3 V bias *         3664 mV
3.3 V bias *         3586 mV
3.3 V bias *         3654 mV
4.3 V bias *         4215 mV
4.3 V bias *         4224 mV
4.3 V bias *         4215 mV
4.3 V bias *         4215 mV
4.3 V bias *         4234 mV
4.3 V bias *         4224 mV
5.0 V bias           4920 mV
CB 2 status:
State                Online
Temperature Intake    20 degrees C / 68 degrees F
Temperature Exhaust   24 degrees C / 75 degrees F
Power 1
1.2 V                1202 mV
1.5 V                1508 mV
1.8 V                1804 mV
2.5 V                2520 mV
3.3 V                3300 mV
3.3 V bias           3300 mV
12.0 V               12200 mV

```

#### show chassis environment cb (PTX5000 Packet Transport Router)

```

user@host> show chassis environment cb
CB 0 status:
State                Online Master
Intake Temperature    38 degrees C / 100 degrees F
Exhaust A Temperature 45 degrees C / 113 degrees F
Exhaust B Temperature 42 degrees C / 107 degrees F
Power 1
1.2 V                1200 mV
1.25 V               1250 mV
2.5 V                2500 mV
3.3 V                3300 mV
Power 2
1.0 V                1000 mV
3.3 V bias           3293 mV
3.9 V                3921 mV
Bus Revision          132
FPGA Revision         27
CB 1 status:
State                Online Standby
Intake Temperature    34 degrees C / 93 degrees F
Exhaust A Temperature 39 degrees C / 102 degrees F
Exhaust B Temperature 36 degrees C / 96 degrees F
Power 1
1.2 V                1199 mV
1.25 V               1250 mV
2.5 V                2499 mV
3.3 V                3299 mV
Power 2
1.0 V                1000 mV

```

3.3 V bias	3312 mV
3.9 V	3961 mV
Bus Revision	132
FPGA Revision	28

### show chassis environment cb (QFabric System)

```
user@switch> show chassis environment cb interconnect-device IC-123 0
CB 0 status:
```

State	Online Master
Left Intake Temperature	33 degrees C / 91 degrees F
Right Intake Temperature	33 degrees C / 91 degrees F
Left Exhaust Temperature	36 degrees C / 96 degrees F
Right Exhaust Temperature	35 degrees C / 95 degrees F
Power	OK
VDD 3V3	3294 mV
VDD 2V5	2436 mV
VDD 1V8	1746 mV
VDD 1V5	1460 mV
VDD 1V25	1210 mV
VDD 1V2	1164 mV
CPU CORE 1V2	1120 mV
VDD 1V0	968 mV
VDD 5V0	5088 mV
CPU MP BIAS 4V3	4050 mV
BIAS 3V3	3180 mV
VTT 0V9	866 mV

## show chassis environment fpc

---

<b>List of Syntax</b>	<a href="#">Syntax on page 988</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 988</a> <a href="#">Syntax (MX Series Routers) on page 988</a> <a href="#">Syntax (MX2010 3D Universal Edge Routers) on page 988</a> <a href="#">Syntax (MX2020 3D Universal Edge Routers) on page 988</a> <a href="#">Syntax (QFX Series) on page 988</a>
<b>Syntax</b>	show chassis environment fpc <slot>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	show chassis environment fpc <lcc number> <slot>
<b>Syntax (MX Series Routers)</b>	show chassis environment fpc <slot> <all-members> <local> <member member-id>
<b>Syntax (MX2010 3D Universal Edge Routers)</b>	show chassis environment fpc <slot>
<b>Syntax (MX2020 3D Universal Edge Routers)</b>	show chassis environment fpc <slot>
<b>Syntax (QFX Series)</b>	show chassis environment fpc <fpc-slot> interconnect-device <i>name</i>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers. Command introduced in Junos OS Release 12.1 for T4000 Core Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
<b>Description</b>	(M40e, M120, M160, M320, MX Series, T Series routers, EX Series, QFX Series, and PTX Series routers only) Display environmental information about Flexible PIC Concentrators (FPCs).
<b>Options</b>	<b>none</b> —Display environmental information about all FPCs. On a TX Matrix router, display environmental information about all FPCs on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about all FPCs on the TX Matrix Plus router and its attached routers.

**all-members**—(MX Series routers only) (Optional) Display environmental information for the FPCs in all the members of the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.

**lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display environmental information for the FPCs in the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display environmental information for the FPCs in the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**slot or fpc-slot**—(Optional) Display environmental information about an individual FPC:

- (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, if you specify the number of the T640 router by using only the **lcc *number*** option (the recommended method), replace **slot** with a value from 0 through 7. Similarly, on a TX Matrix Plus router, if you specify the number of the router by using only the **lcc *number*** option (the recommended method), replace **slot** with a value from 0 through 7. Otherwise, replace **slot** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis environment fpc 1 lcc 1
user@host> show chassis environment fpc 9
```

- M120 router—Replace **slot** with a value from 0 through 5.
- MX240 router—Replace **slot** with a value from 0 through 2.
- MX480 router—Replace **slot** with a value from 0 through 5.
- MX960 router—Replace **slot** with a value from 0 through 11.
- MX2010 router—Replace **slot** with a value from 0 through 9.
- MX2020 router—Replace **slot** with a value from 0 through 19.
- Other routers—Replace **slot** with a value from 0 through 7.
- EX Series switches:

- EX3200 switches and EX4200 standalone switches—Replace **slot** with 0.
- EX4200 switches in a Virtual Chassis configuration—Replace **slot** with a value from 0 through 9 (switch's member ID).
- EX6210 switches—Replace **slot** with a value from 0 through 3 (line card only), 4 or 5 (line card or Switch Fabric and Rotuing Engine (SRE) module), or 6 through 9 (line card only).
- EX8208 switches—Replace **slot** with a value from 0 through 7 (line card).
- EX8216 switches—Replace **slot** with a value from 0 through 15 (line card).
- QFX3500 switches —Replace **fpc-slot** with 0 through 15.
- PTX5000 Packet Transport Router—Replace **fpc-slot** with 0 through 7.

**Required Privilege Level**    view

- Related Documentation**
- [request chassis fpc on page 882](#)
  - [show chassis fpc on page 1160](#)
  - *show chassis fpc-feb-connectivity*
  - *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
  - *MX960 Flexible PIC Concentrator Description*

- List of Sample Output**
- [show chassis environment fpc \(M120 Router\) on page 992](#)
  - [show chassis environment fpc \(M160 Router\) on page 993](#)
  - [show chassis environment fpc \(M320 Router\) on page 993](#)
  - [show chassis environment fpc \(MX2020 Router\) on page 994](#)
  - [show chassis environment fpc \(MX2010 Router\) on page 997](#)
  - [show chassis environment fpc \(MX240 Router\) on page 999](#)
  - [show chassis environment fpc \(MX480 Router\) on page 1000](#)
  - [show chassis environment fpc \(MX960 Router\) on page 1001](#)
  - [show chassis environment fpc \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 1002](#)
  - [show chassis environment fpc \(MX240, MX480, MX960 with Application Services Modular Line Card\) on page 1003](#)
  - [show chassis environment fpc \(T320, T640, and T1600 Routers\) on page 1004](#)
  - [show chassis environment fpc \(T4000 Router\) on page 1004](#)
  - [show chassis environment fpc lcc \(TX Matrix Router\) on page 1009](#)
  - [show chassis environment fpc lcc \(TX Matrix Plus Router\) on page 1010](#)
  - [show chassis environment fpc \(QFX Series\) on page 1011](#)
  - [show chassis environment fpc interconnect-device \(QFabric Systems\) on page 1011](#)
  - [show chassis environment fpc 0 \(PTX5000 Packet Transport Router\) on page 1011](#)
  - [show chassis environment fpc 07 \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 1012](#)

[show chassis environment FPC 1 \(MX Routers with Media Services Blade \[MSB\]\) on page 1013](#)

**Output Fields** [Table 108 on page 991](#) lists the output fields for the **show chassis environment fpc** command. Output fields are listed in the approximate order in which they appear.

**Table 108: show chassis environment fpc Output Fields**

Field Name	Field Description
<b>State</b>	<p>Status of the FPC:</p> <ul style="list-style-type: none"> <li>• <b>Unknown</b>—FPC is not detected by the router.</li> <li>• <b>Empty</b>—No FPC is present.</li> <li>• <b>Present</b>—FPC is detected by the chassis daemon but is either not supported by the current version of the Junos OS, or the FPC is coming up but not yet online.</li> <li>• <b>Ready</b>—FPC is in intermediate or transition state.</li> <li>• <b>Announce online</b>—Intermediate state during which the FPC is coming up but not yet online, and the chassis manager acknowledges the chassisd FPC online initiative.</li> <li>• <b>Online</b>—FPC is online and running.</li> <li>• <b>Offline</b>—FPC is powered down.</li> <li>• <b>Diagnostics</b>—FPC is set to operate in diagnostics mode.</li> </ul>
<b>Temperature</b>	(M40e and M160 routers and QFX Series only) Temperature of the air flowing past the FPC.
<b>PMB Temperature</b>	<p>(PTX Series only) Temperature of the air flowing past the PMB (bottom of the FPC).</p> <p>The PTX5000 Packet Transport Router with FPC2-PTX-PIA include multiple temperatures for PMB (<b>TEMPO</b> and <b>TEMPI</b>).</p>
<b>PMB CPU Temperature</b>	(PTX5000 Packet Transport Router with FPC2-PTX-PIA only) Temperature of the air flowing past the PMB CPU.
<b>Temperature Intake</b>	(M320 routers, MX2010 routers, MX2020 routers, and PTX Series only) Temperature of the air flowing into the chassis.
<b>Temperature Top</b>	(T Series routers only) Temperature of the air flowing past the top of the FPC.
<b>Temperature Exhaust</b>	<p>(M120 and M320 routers, MX2010 routers, MX2020 routers, and PTX Series only) Temperature of the air flowing out of the chassis.</p> <p>The PTX Series Packet Transport Routers, and the MX2010 and MX2020 routers include exhaust temperatures for multiple zones (<b>Exhaust A</b> and <b>Exhaust B</b>).</p>
<b>Temperature Bottom</b>	(T Series routers only) Temperature of the air flowing past the bottom of the FPC.
<b>TL <i>n</i> Temperature</b>	(PTX Series only) Temperature of the air flowing past the specified TL area of the packet forwarding engine (PFE) on the FPC.
<b>TQ <i>n</i> Temperature</b>	(PTX Series only) Temperature of the air flowing past the specified TQ area of the packet forwarding engine (PFE) on the FPC.
<b>Temperature MMBO</b>	(T640 router only) Temperature of the air flowing past the type 3 FPC.

Table 108: show chassis environment fpc Output Fields (*continued*)

Field Name	Field Description
<b>Temperature MMB1</b>	(M320 and T Series routers only) Temperature of the air flowing past the type 1, type 2, and type 3 FPC.
<b>Power</b>	Information about the voltage supplied to the FPC. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
<b>CMB Revision or BUS revision</b>	Revision level of the chassis management bus device (M Series router) or bus (T Series routers).

## Sample Output

### show chassis environment fpc (M120 Router)

```

user@host> show chassis environment fpc
FPC 2 status:
  State                               Online
  Temperature Exhaust A               32 degrees C / 89 degrees F
  Temperature Exhaust B               31 degrees C / 87 degrees F
  Power A-Board
    1.2 V                             1202 mV
    1.5 V                             1508 mV
    1.8 V                             1798 mV
    2.5 V                             2507 mV
    3.3 V                             3351 mV
    5.0 V                             4995 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  I2C Slave Revision                  12
FPC 3 status:
  State                               Online
  Temperature Exhaust A               31 degrees C / 87 degrees F
  Temperature Exhaust B               33 degrees C / 91 degrees F
  Power A-Board
    1.2 V                             1211 mV
    1.5 V                             1501 mV
    1.8 V                             1798 mV
    2.5 V                             2471 mV
    3.3 V                             3293 mV
    5.0 V                             4930 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  Power B-Board
    1.2 V                             1214 mV
    1.5 V                             1501 mV
    2.5 V                             2471 mV
    3.3 V                             3300 mV
    5.0 V                             4943 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  I2C Slave Revision                  12
FPC 4 status:
  State                               Online

```



```

Temperature Exhaust A      32 degrees C / 89 degrees F
Temperature Exhaust B      30 degrees C / 86 degrees F
Power A-Board
  1.2 V                    1195 mV
  1.5 V                    1504 mV
  1.8 V                    1801 mV
  2.5 V                    2504 mV
  3.3 V                    3293 mV
  5.0 V                    4917 mV
  3.3 V bias               3296 mV
  1.2 V Rocket IO          1202 mV
  1.5 V Rocket IO          1492 mV
I2C Slave Revision        12

```

#### show chassis environment fpc (M160 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
  State                      Online
  Temperature                 42 degrees C / 107 degrees F
  Power:
    1.5 V                    1500 mV
    2.5 V                    2509 mV
    3.3 V                    3308 mV
    5.0 V                    4991 mV
    5.0 V bias               4952 mV
    8.0 V bias               8307 mV
  CMB Revision                12
FPC 1 status:
  State                      Online
  Temperature                 45 degrees C / 113 degrees F
  Power:
    1.5 V                    1498 mV
    2.5 V                    2501 mV
    3.3 V                    3319 mV
    5.0 V                    5020 mV
    5.0 V bias               5025 mV
    8.0 V bias               8307 mV
  CMB Revision                12

```

#### show chassis environment fpc (M320 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
  State                      Online
  Temperature Intake          27 degrees C / 80 degrees F
  Temperature Exhaust         38 degrees C / 100 degrees F
  Temperature MMB1            31 degrees C / 87 degrees F
  Power:
    1.5 V                    1487 mV
    1.5 V *                  1494 mV
    1.8 V                    1821 mV
    2.5 V                    2533 mV
    3.3 V                    3323 mV
    5.0 V                    5028 mV
    3.3 V bias               3296 mV
    5.0 V bias               4984 mV
  CMB Revision                16
FPC 1 status:
  State                      Online
  Temperature Intake          27 degrees C / 80 degrees F

```

```

Temperature Exhaust      37 degrees C / 98 degrees F
Temperature MMB1         32 degrees C / 89 degrees F
Power:
  1.5 V                  1504 mV
  1.5 V *                1499 mV
  1.8 V                  1820 mV
  2.5 V                  2529 mV
  3.3 V                  3328 mV
  5.0 V                  5013 mV
  3.3 V bias             3294 mV
  5.0 V bias             4984 mV
CMB Revision             16
FPC 2 status:
State                    Online
Temperature Intake        28 degrees C / 82 degrees F
Temperature Exhaust       38 degrees C / 100 degrees F
Temperature MMB1         32 degrees C / 89 degrees F
Power:
  1.5 V                  1498 mV
  1.5 V *                1487 mV
  1.8 V                  1816 mV
  2.5 V                  2531 mV
  3.3 V                  3324 mV
  5.0 V                  5025 mV
  3.3 V bias             3277 mV
  5.0 V bias             5013 mV
CMB Revision             17
FPC 3 status:
...
```

### show chassis environment fpc (MX2020 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State                    Online
Temperature Intake        41 degrees C / 105 degrees F
Temperature Exhaust A     48 degrees C / 118 degrees F
Temperature Exhaust B     60 degrees C / 140 degrees F
Temperature LU 0 TSen     56 degrees C / 132 degrees F
Temperature LU 0 Chip     59 degrees C / 138 degrees F
Temperature LU 1 TSen     56 degrees C / 132 degrees F
Temperature LU 1 Chip     61 degrees C / 141 degrees F
Temperature LU 2 TSen     56 degrees C / 132 degrees F
Temperature LU 2 Chip     52 degrees C / 125 degrees F
Temperature LU 3 TSen     56 degrees C / 132 degrees F
Temperature LU 3 Chip     52 degrees C / 125 degrees F
Temperature MQ 0 TSen     49 degrees C / 120 degrees F
Temperature MQ 0 Chip     49 degrees C / 120 degrees F
Temperature MQ 1 TSen     49 degrees C / 120 degrees F
Temperature MQ 1 Chip     52 degrees C / 125 degrees F
Temperature MQ 2 TSen     49 degrees C / 120 degrees F
Temperature MQ 2 Chip     45 degrees C / 113 degrees F
Temperature MQ 3 TSen     49 degrees C / 120 degrees F
Temperature MQ 3 Chip     46 degrees C / 114 degrees F
Power
AS-BIAS3V3-z12105        3299 mV
AS-VDD1V8-z12006         1807 mV
AS-VDD2V5-z12006         2512 mV
AS-AVDD1V0-z12004         997 mV
AS-PCIE_1V0-z12004        996 mV
AS-VDD3V3-z12004        3294 mV
```

```

AS-VDD_1V5A-z12004      1501 mV
AS-VDD_1V5B-z12004      1498 mV
AS-LU0_1V0-z12004        998 mV
AS-LU1_1V0-z12004       1002 mV
AS-MQ0_1V0-z12004        999 mV
AS-MQ1_1V0-z12004        994 mV
AS-LU2_1V0-z12004       1000 mV
AS-LU3_1V0-z12004        998 mV
AS-MQ2_1V0-z12004       1002 mV
AS-MQ3_1V0-z12004        999 mV
AS-PMB_1V1-z12006       1096 mV
I2C Slave Revision      68
FPC 1 status:
State                    Online
Temperature Intake       39 degrees C / 102 degrees F
Temperature Exhaust A   48 degrees C / 118 degrees F
Temperature Exhaust B   55 degrees C / 131 degrees F
Temperature LU 0 TSen    52 degrees C / 125 degrees F
Temperature LU 0 Chip    54 degrees C / 129 degrees F
Temperature LU 1 TSen    52 degrees C / 125 degrees F
Temperature LU 1 Chip    56 degrees C / 132 degrees F
Temperature LU 2 TSen    52 degrees C / 125 degrees F
Temperature LU 2 Chip    49 degrees C / 120 degrees F
Temperature LU 3 TSen    52 degrees C / 125 degrees F
Temperature LU 3 Chip    50 degrees C / 122 degrees F
Temperature MQ 0 TSen    48 degrees C / 118 degrees F
Temperature MQ 0 Chip    48 degrees C / 118 degrees F
Temperature MQ 1 TSen    48 degrees C / 118 degrees F
Temperature MQ 1 Chip    51 degrees C / 123 degrees F
Temperature MQ 2 TSen    48 degrees C / 118 degrees F
Temperature MQ 2 Chip    45 degrees C / 113 degrees F
Temperature MQ 3 TSen    48 degrees C / 118 degrees F
Temperature MQ 3 Chip    45 degrees C / 113 degrees F
Power
AS-BIAS3V3-z12105       3291 mV
AS-VDD1V8-z12006       1786 mV
AS-VDD2V5-z12006       2496 mV
AS-AVDD1V0-z12004       1000 mV
AS-PCIE_1V0-z12004       1000 mV
AS-VDD3V3-z12004       3294 mV
AS-VDD_1V5A-z12004      1500 mV
AS-VDD_1V5B-z12004      1498 mV
AS-LU0_1V0-z12004       1003 mV
AS-LU1_1V0-z12004       1000 mV
AS-MQ0_1V0-z12004       1000 mV
AS-MQ1_1V0-z12004        995 mV
AS-LU2_1V0-z12004       1002 mV
AS-LU3_1V0-z12004        997 mV
AS-MQ2_1V0-z12004       1000 mV
AS-MQ3_1V0-z12004        998 mV
AS-PMB_1V1-z12006       1096 mV
I2C Slave Revision      68
FPC 2 status:
State                    Online
Temperature Intake       39 degrees C / 102 degrees F
Temperature Exhaust A   48 degrees C / 118 degrees F
Temperature Exhaust B   58 degrees C / 136 degrees F
Temperature LU 0 TSen    55 degrees C / 131 degrees F
Temperature LU 0 Chip    57 degrees C / 134 degrees F
Temperature LU 1 TSen    55 degrees C / 131 degrees F
Temperature LU 1 Chip    63 degrees C / 145 degrees F

```

```

Temperature LU 2 TSen      55 degrees C / 131 degrees F
Temperature LU 2 Chip      51 degrees C / 123 degrees F
Temperature LU 3 TSen      55 degrees C / 131 degrees F
Temperature LU 3 Chip      52 degrees C / 125 degrees F
Temperature MQ 0 TSen      48 degrees C / 118 degrees F
Temperature MQ 0 Chip      50 degrees C / 122 degrees F
Temperature MQ 1 TSen      48 degrees C / 118 degrees F
Temperature MQ 1 Chip      52 degrees C / 125 degrees F
Temperature MQ 2 TSen      48 degrees C / 118 degrees F
Temperature MQ 2 Chip      47 degrees C / 116 degrees F
Temperature MQ 3 TSen      48 degrees C / 118 degrees F
Temperature MQ 3 Chip      47 degrees C / 116 degrees F
Power
  AS-BIAS3V3-z12105        3299 mV
  AS-VDD1V8-z12006         1805 mV
  AS-VDD2V5-z12006         2510 mV
  AS-AVDD1V0-z12004         999 mV
  AS-PCIE_1V0-z12004         998 mV
  AS-VDD3V3-z12004         3296 mV
  AS-VDD_1V5A-z12004        1492 mV
  AS-VDD_1V5B-z12004        1497 mV
  AS-LU0_1V0-z12004         997 mV
  AS-LU1_1V0-z12004        1000 mV
  AS-MQ0_1V0-z12004         998 mV
  AS-MQ1_1V0-z12004        1001 mV
  AS-LU2_1V0-z12004         996 mV
  AS-LU3_1V0-z12004         995 mV
  AS-MQ2_1V0-z12004         998 mV
  AS-MQ3_1V0-z12004         997 mV
  AS-PMB_1V1-z12006        1100 mV
I2C Slave Revision        68
FPC 3 status:
State                      Online
Temperature Intake          41 degrees C / 105 degrees F
Temperature Exhaust A       48 degrees C / 118 degrees F
Temperature Exhaust B       58 degrees C / 136 degrees F
Temperature LU 0 TSen       56 degrees C / 132 degrees F
Temperature LU 0 Chip       59 degrees C / 138 degrees F
Temperature LU 1 TSen       56 degrees C / 132 degrees F
Temperature LU 1 Chip       61 degrees C / 141 degrees F
Temperature LU 2 TSen       56 degrees C / 132 degrees F
Temperature LU 2 Chip       51 degrees C / 123 degrees F
Temperature LU 3 TSen       56 degrees C / 132 degrees F
Temperature LU 3 Chip       53 degrees C / 127 degrees F
Temperature MQ 0 TSen       50 degrees C / 122 degrees F
Temperature MQ 0 Chip       51 degrees C / 123 degrees F
Temperature MQ 1 TSen       50 degrees C / 122 degrees F
Temperature MQ 1 Chip       55 degrees C / 131 degrees F
Temperature MQ 2 TSen       50 degrees C / 122 degrees F
Temperature MQ 2 Chip       47 degrees C / 116 degrees F
Temperature MQ 3 TSen       50 degrees C / 122 degrees F
Temperature MQ 3 Chip       50 degrees C / 122 degrees F
Power
  AS-BIAS3V3-z12105        3305 mV
  AS-VDD1V8-z12006         1810 mV
  AS-VDD2V5-z12006         2508 mV
  AS-AVDD1V0-z12004         999 mV
  AS-PCIE_1V0-z12004        1001 mV
  AS-VDD3V3-z12004         3294 mV
  AS-VDD_1V5A-z12004        1500 mV
  AS-VDD_1V5B-z12004        1498 mV

```

```

AS-LU0_1V0-z12004      998 mV
AS-LU1_1V0-z12004      998 mV
AS-MQ0_1V0-z12004      999 mV
AS-MQ1_1V0-z12004      998 mV
AS-LU2_1V0-z12004      1000 mV
AS-LU3_1V0-z12004      1001 mV
AS-MQ2_1V0-z12004      996 mV
AS-MQ3_1V0-z12004      998 mV
AS-PMB_1V1-z12006      1098 mV
I2C Slave Revision      68
FPC 4 status:
...

```

### show chassis environment fpc (MX2010 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State                               Online
Temperature Intake                  36 degrees C / 96 degrees F
Temperature Exhaust A                42 degrees C / 107 degrees F
Temperature Exhaust B                51 degrees C / 123 degrees F
Temperature LU 0 TSen                49 degrees C / 120 degrees F
Temperature LU 0 Chip                50 degrees C / 122 degrees F
Temperature LU 1 TSen                49 degrees C / 120 degrees F
Temperature LU 1 Chip                54 degrees C / 129 degrees F
Temperature LU 2 TSen                49 degrees C / 120 degrees F
Temperature LU 2 Chip                45 degrees C / 113 degrees F
Temperature LU 3 TSen                49 degrees C / 120 degrees F
Temperature LU 3 Chip                46 degrees C / 114 degrees F
Temperature MQ 0 TSen                40 degrees C / 104 degrees F
Temperature MQ 0 Chip                41 degrees C / 105 degrees F
Temperature MQ 1 TSen                40 degrees C / 104 degrees F
Temperature MQ 1 Chip                44 degrees C / 111 degrees F
Temperature MQ 2 TSen                40 degrees C / 104 degrees F
Temperature MQ 2 Chip                38 degrees C / 100 degrees F
Temperature MQ 3 TSen                40 degrees C / 104 degrees F
Temperature MQ 3 Chip                41 degrees C / 105 degrees F
Power
AS-BIAS3V3-z12105                  3300 mV
AS-VDD1V8-z12006                    1805 mV
AS-VDD2V5-z12006                    2505 mV
AS-AVDD1V0-z12004                    998 mV
AS-PCIE_1V0-z12004                    999 mV
AS-VDD3V3-z12004                     3303 mV
AS-VDD_1V5A-z12004                   1497 mV
AS-VDD_1V5B-z12004                   1497 mV
AS-LU0_1V0-z12004                     998 mV
AS-LU1_1V0-z12004                     1003 mV
AS-MQ0_1V0-z12004                     998 mV
AS-MQ1_1V0-z12004                     998 mV
AS-LU2_1V0-z12004                     997 mV
AS-LU3_1V0-z12004                     1001 mV
AS-MQ2_1V0-z12004                     996 mV
AS-MQ3_1V0-z12004                     994 mV
AS-PMB_1V1-z12006                     1097 mV
I2C Slave Revision                    68
FPC 1 status:
State                               Online
Temperature Intake                  34 degrees C / 93 degrees F
Temperature Exhaust A                46 degrees C / 114 degrees F
Temperature Exhaust B                54 degrees C / 129 degrees F

```

```

Temperature LU 0 TSen      45 degrees C / 113 degrees F
Temperature LU 0 Chip      55 degrees C / 131 degrees F
Temperature LU 1 TSen      45 degrees C / 113 degrees F
Temperature LU 1 Chip      44 degrees C / 111 degrees F
Temperature LU 2 TSen      45 degrees C / 113 degrees F
Temperature LU 2 Chip      50 degrees C / 122 degrees F
Temperature LU 3 TSen      45 degrees C / 113 degrees F
Temperature LU 3 Chip      58 degrees C / 136 degrees F
Temperature XM 0 TSen      45 degrees C / 113 degrees F
Temperature XM 0 Chip      51 degrees C / 123 degrees F
Temperature XF 0 TSen      45 degrees C / 113 degrees F
Temperature XF 0 Chip      63 degrees C / 145 degrees F
Temperature PLX Switch TSen 45 degrees C / 113 degrees F
Temperature PLX Switch Chip 47 degrees C / 116 degrees F
Power
MPC-BIAS3V3-z12105        3300 mV
MPC-VDD3V3-z16100         3294 mV
MPC-VDD2V5-z16100         2505 mV
MPC-VDD1V8-z12004         1796 mV
MPC-AVDD1V0-z12004         991 mV
MPC-VDD1V2-z16100         1196 mV
MPC-VDD1V5A-z12004        1491 mV
MPC-VDD1V5B-z12004        1492 mV
MPC-XF_OV9-z12004         996 mV
MPC-PCIE_1V0-z16100       1003 mV
MPC-LU0_1V0-z12004        996 mV
MPC-LU1_1V0-z12004        996 mV
MPC-LU2_1V0-z12004        998 mV
MPC-LU3_1V0-z12004        994 mV
MPC-12VA-BMR453           12031 mV
MPC-12VB-BMR453           12003 mV
MPC-PMB_1V1-z12006        1104 mV
MPC-PMB_1V2-z12106        1194 mV
MPC-XM_OV9-vt273m         911 mV
I2C Slave Revision        110
FPC 8 status:
State                      Online
Temperature Intake          32 degrees C / 89 degrees F
Temperature Exhaust A       44 degrees C / 111 degrees F
Temperature Exhaust B       37 degrees C / 98 degrees F
Temperature LU 0 TCAM TSen  41 degrees C / 105 degrees F
Temperature LU 0 TCAM Chip  49 degrees C / 120 degrees F
Temperature LU 0 TSen        41 degrees C / 105 degrees F
Temperature LU 0 Chip        52 degrees C / 125 degrees F
Temperature MQ 0 TSen        41 degrees C / 105 degrees F
Temperature MQ 0 Chip        47 degrees C / 116 degrees F
Temperature LU 1 TCAM TSen  39 degrees C / 102 degrees F
Temperature LU 1 TCAM Chip  42 degrees C / 107 degrees F
Temperature LU 1 TSen        39 degrees C / 102 degrees F
Temperature LU 1 Chip        46 degrees C / 114 degrees F
Temperature MQ 1 TSen        39 degrees C / 102 degrees F
Temperature MQ 1 Chip        45 degrees C / 113 degrees F
Power
MPC-BIAS3V3-z12105        3296 mV
MPC-VDD3V3-z12006         3298 mV
MPC-VDD2V5-z12006         2505 mV
MPC-TCAM_1V0-z12004        997 mV
MPC-AVDD1V0-z12006        1007 mV
MPC-VDD1V8-z12006         1803 mV
MPC-PCIE_1V0-z12006        1004 mV
MPC-LU0_1V0-z12004        1000 mV

```

```

MPC-MQ0_1V0-z12004      999 mV
MPC-VDD_1V5-z12004      1498 mV
MPC-PMB_1V1-z12006      1102 mV
MPC-9VA-BMR453          9009 mV
MPC-9VB-BMR453          8960 mV
MPC-PMB_1V2-z12105      1202 mV
MPC-LU1_1V0-z12004      1005 mV
MPC-MQ1_1V0-z12004      1000 mV
I2C Slave Revision      70
FPC 9 status:
State                    Online
Temperature Intake       34 degrees C / 93 degrees F
Temperature Exhaust A    41 degrees C / 105 degrees F
Temperature Exhaust B    54 degrees C / 129 degrees F
Temperature LU 0 TSen     51 degrees C / 123 degrees F
Temperature LU 0 Chip     52 degrees C / 125 degrees F
Temperature LU 1 TSen     51 degrees C / 123 degrees F
Temperature LU 1 Chip     55 degrees C / 131 degrees F
Temperature LU 2 TSen     51 degrees C / 123 degrees F
Temperature LU 2 Chip     47 degrees C / 116 degrees F
Temperature LU 3 TSen     51 degrees C / 123 degrees F
Temperature LU 3 Chip     47 degrees C / 116 degrees F
Temperature MQ 0 TSen     40 degrees C / 104 degrees F
Temperature MQ 0 Chip     42 degrees C / 107 degrees F
Temperature MQ 1 TSen     40 degrees C / 104 degrees F
Temperature MQ 1 Chip     44 degrees C / 111 degrees F
Temperature MQ 2 TSen     40 degrees C / 104 degrees F
Temperature MQ 2 Chip     38 degrees C / 100 degrees F
Temperature MQ 3 TSen     40 degrees C / 104 degrees F
Temperature MQ 3 Chip     40 degrees C / 104 degrees F
Power
AS-BIAS3V3-z12105        3302 mV
AS-VDD1V8-z12006         1808 mV
AS-VDD2V5-z12006         2513 mV
AS-AVDD1V0-z12004         997 mV
AS-PCIE_1V0-z12004         999 mV
AS-VDD3V3-z12004         3294 mV
AS-VDD_1V5A-z12004        1503 mV
AS-VDD_1V5B-z12004        1502 mV
AS-LU0_1V0-z12004         996 mV
AS-LU1_1V0-z12004         999 mV
AS-MQ0_1V0-z12004         997 mV
AS-MQ1_1V0-z12004         999 mV
AS-LU2_1V0-z12004         997 mV
AS-LU3_1V0-z12004         998 mV
AS-MQ2_1V0-z12004        1000 mV
AS-MQ3_1V0-z12004        1000 mV
AS-PMB_1V1-z12006        1102 mV
I2C Slave Revision      68

```

#### show chassis environment fpc (MX240 Router)

```

user@host> show chassis environment fpc
FPC 1 status:
State                    Online
Temperature Intake       34 degrees C / 93 degrees F
Temperature Exhaust A    39 degrees C / 102 degrees F
Temperature Exhaust B    53 degrees C / 127 degrees F
Temperature I3 0 TSensor  51 degrees C / 123 degrees F
Temperature I3 0 Chip     54 degrees C / 129 degrees F
Temperature I3 1 TSensor  50 degrees C / 122 degrees F

```

```

Temperature I3 1 Chip      53 degrees C / 127 degrees F
Temperature I3 2 TSensor   48 degrees C / 118 degrees F
Temperature I3 2 Chip      51 degrees C / 123 degrees F
Temperature I3 3 TSensor   45 degrees C / 113 degrees F
Temperature I3 3 Chip      48 degrees C / 118 degrees F
Temperature IA 0 TSensor   45 degrees C / 113 degrees F
Temperature IA 0 Chip      45 degrees C / 113 degrees F
Temperature IA 1 TSensor   45 degrees C / 113 degrees F
Temperature IA 1 Chip      49 degrees C / 120 degrees F
Power
  1.5 V                    1492 mV
  2.5 V                    2507 mV
  3.3 V                    3306 mV
  1.8 V PFE 0              1801 mV
  1.8 V PFE 1              1804 mV
  1.8 V PFE 2              1798 mV
  1.8 V PFE 3              1798 mV
  1.2 V PFE 0              1169 mV
  1.2 V PFE 1              1189 mV
  1.2 V PFE 2              1182 mV
  1.2 V PFE 3              1176 mV
I2C Slave Revision        42
FPC 2 status:
State                      Online
Temperature Intake          33 degrees C / 91 degrees F
Temperature Exhaust A       41 degrees C / 105 degrees F
Temperature Exhaust B       53 degrees C / 127 degrees F
Temperature I3 0 TSensor    53 degrees C / 127 degrees F
Temperature I3 0 Chip       58 degrees C / 136 degrees F
Temperature I3 1 TSensor    52 degrees C / 125 degrees F
Temperature I3 1 Chip       56 degrees C / 132 degrees F
Temperature I3 2 TSensor    50 degrees C / 122 degrees F
Temperature I3 2 Chip       52 degrees C / 125 degrees F
Temperature I3 3 TSensor    46 degrees C / 114 degrees F
Temperature I3 3 Chip       49 degrees C / 120 degrees F
Temperature IA 0 TSensor    51 degrees C / 123 degrees F
Temperature IA 0 Chip       49 degrees C / 120 degrees F
Temperature IA 1 TSensor    48 degrees C / 118 degrees F
Temperature IA 1 Chip       53 degrees C / 127 degrees F
Power
  1.5 V                    1492 mV
  2.5 V                    2445 mV
  3.3 V                    3293 mV
  1.8 V PFE 0              1827 mV
  1.8 V PFE 1              1775 mV
  1.8 V PFE 2              1788 mV
  1.8 V PFE 3              1798 mV
  1.2 V PFE 0              1250 mV
  1.2 V PFE 1              1234 mV
  1.2 V PFE 2              1231 mV
  1.2 V PFE 3              1192 mV
I2C Slave Revision        42

```

#### show chassis environment fpc (MX480 Router)

```

user@host> show chassis environment fpc
FPC 1 status:
State                      Online
Temperature Intake          36 degrees C / 96 degrees F
Temperature Exhaust A       41 degrees C / 105 degrees F
Temperature Exhaust B       55 degrees C / 131 degrees F

```



```

Temperature I3 0 TSensor 55 degrees C / 131 degrees F
Temperature I3 0 Chip    57 degrees C / 134 degrees F
Temperature I3 1 TSensor 53 degrees C / 127 degrees F
Temperature I3 1 Chip    53 degrees C / 127 degrees F
Temperature I3 2 TSensor 52 degrees C / 125 degrees F
Temperature I3 2 Chip    49 degrees C / 120 degrees F
Temperature I3 3 TSensor 47 degrees C / 116 degrees F
Temperature I3 3 Chip    47 degrees C / 116 degrees F
Temperature IA 0 TSensor 54 degrees C / 129 degrees F
Temperature IA 0 Chip    58 degrees C / 136 degrees F
Temperature IA 1 TSensor 48 degrees C / 118 degrees F
Temperature IA 1 Chip    53 degrees C / 127 degrees F
Power
  1.5 V      1479 mV
  2.5 V      2542 mV
  3.3 V      3319 mV
  1.8 V PFE 0 1811 mV
  1.8 V PFE 1 1804 mV
  1.8 V PFE 2 1804 mV
  1.8 V PFE 3 1814 mV
  1.2 V PFE 0 1192 mV
  1.2 V PFE 1 1202 mV
  1.2 V PFE 2 1205 mV
  1.2 V PFE 3 1189 mV
I2C Slave Revision 40

```

#### show chassis environment fpc (MX960 Router)

```

user@host> show chassis environment fpc
FPC 5 status:
State      Online
Temperature Intake 27 degrees C / 80 degrees F
Temperature Exhaust A 34 degrees C / 93 degrees F
Temperature Exhaust B 40 degrees C / 104 degrees F
Temperature I3 0 TSensor 39 degrees C / 102 degrees F
Temperature I3 0 Chip 41 degrees C / 105 degrees F
Temperature I3 1 TSensor 38 degrees C / 100 degrees F
Temperature I3 1 Chip 37 degrees C / 98 degrees F
Temperature I3 2 TSensor 37 degrees C / 98 degrees F
Temperature I3 2 Chip 34 degrees C / 93 degrees F
Temperature I3 3 TSensor 32 degrees C / 89 degrees F
Temperature I3 3 Chip 33 degrees C / 91 degrees F
Temperature IA 0 TSensor 39 degrees C / 102 degrees F
Temperature IA 0 Chip 44 degrees C / 111 degrees F
Temperature IA 1 TSensor 36 degrees C / 96 degrees F
Temperature IA 1 Chip 44 degrees C / 111 degrees F
Power
  1.5 V      1479 mV
  2.5 V      2523 mV
  3.3 V      3254 mV
  1.8 V PFE 0 1798 mV
  1.8 V PFE 1 1798 mV
  1.8 V PFE 2 1807 mV
  1.8 V PFE 3 1791 mV
  1.2 V PFE 0 1173 mV
  1.2 V PFE 1 1179 mV
  1.2 V PFE 2 1179 mV
  1.2 V PFE 3 1185 mV
I2C Slave Revision 6
FPC 6 status:
State      Online

```

```

Temperature Intake      25 degrees C / 77 degrees F
Temperature Exhaust A   38 degrees C / 100 degrees F
Temperature Exhaust B   38 degrees C / 100 degrees F
Temperature I3 0 TSensor 40 degrees C / 104 degrees F
Temperature I3 0 Chip    40 degrees C / 104 degrees F
Temperature I3 1 TSensor 40 degrees C / 104 degrees F
Temperature I3 1 Chip    38 degrees C / 100 degrees F
Temperature I3 2 TSensor 37 degrees C / 98 degrees F
Temperature I3 2 Chip    32 degrees C / 89 degrees F
Temperature I3 3 TSensor 34 degrees C / 93 degrees F
Temperature I3 3 Chip    33 degrees C / 91 degrees F
Temperature IA 0 TSensor 45 degrees C / 113 degrees F
Temperature IA 0 Chip    47 degrees C / 116 degrees F
Temperature IA 1 TSensor 37 degrees C / 98 degrees F
Temperature IA 1 Chip    42 degrees C / 107 degrees F
Power
  1.5 V      1485 mV
  2.5 V      2510 mV
  3.3 V      3332 mV
  1.8 V PFE 0 1801 mV
  1.8 V PFE 1 1814 mV
  1.8 V PFE 2 1804 mV
  1.8 V PFE 3 1820 mV
  1.2 V PFE 0 1192 mV
  1.2 V PFE 1 1189 mV
  1.2 V PFE 2 1202 mV
  1.2 V PFE 3 1156 mV
I2C Slave Revision      40

```

#### show chassis environment fpc (MX480 Router with 100-Gigabit Ethernet CFP)

```

user@host> show chassis environment fpc
FPC 0 status:
State      Online
Temperature Intake      32 degrees C / 89 degrees F
Temperature Exhaust A   39 degrees C / 102 degrees F
Temperature Exhaust B   37 degrees C / 98 degrees F
Temperature QX 0 TSen    44 degrees C / 111 degrees F
Temperature QX 0 Chip    48 degrees C / 118 degrees F
Temperature LU 0 TCAM TSen 44 degrees C / 111 degrees F
Temperature LU 0 TCAM Chip 47 degrees C / 116 degrees F
Temperature LU 0 TSen    44 degrees C / 111 degrees F
Temperature LU 0 Chip    48 degrees C / 118 degrees F
Temperature MQ 0 TSen    44 degrees C / 111 degrees F
Temperature MQ 0 Chip    47 degrees C / 116 degrees F
Power
MPC-BIAS3V3-z12105      3297 mV
MPC-VDD3V3-z12105      3306 mV
MPC-VDD2V5-z12105      2498 mV
MPC-TCAM_1V0-z12004      999 mV
MPC-AVDD1V0-z12006      999 mV
MPC-VDD1V8-z12006      1796 mV
MPC-PCIE_1V0-z12006      1002 mV
MPC-LU0_1V0-z12004      997 mV
MPC-MQ0_1V0-z12004      995 mV
MPC-VDD_1V5-z12004      1496 mV
MPC-PMB_1V1-z12006      1094 mV
MPC-9VA-BMR453          9054 mV
MPC-9VB-BMR453          9037 mV
MPC-PMB_1V2-z12106      1191 mV
MPC-QXM0_1V0-z12006      1000 mV

```

```

I2C Slave Revision          66
FPC 1 status:
State                       Online
Temperature Intake          35 degrees C / 95 degrees F
Temperature Exhaust A      50 degrees C / 122 degrees F
Temperature Exhaust B      56 degrees C / 132 degrees F
Temperature LU 0 TSen       46 degrees C / 114 degrees F
Temperature LU 0 Chip       59 degrees C / 138 degrees F
Temperature LU 1 TSen       46 degrees C / 114 degrees F
Temperature LU 1 Chip       45 degrees C / 113 degrees F
Temperature LU 2 TSen       46 degrees C / 114 degrees F
Temperature LU 2 Chip       60 degrees C / 140 degrees F
Temperature LU 3 TSen       46 degrees C / 114 degrees F
Temperature LU 3 Chip       71 degrees C / 159 degrees F
Temperature XM 0 TSen       46 degrees C / 114 degrees F
Temperature XM 0 Chip       -18 degrees C / 0 degrees F
Temperature XF 0 TSen       46 degrees C / 114 degrees F
Temperature XF 0 Chip       76 degrees C / 168 degrees F
Power
MPC-BIAS3V3-z12105         3292 mV
MPC-VDD3V3-z16100          3303 mV
MPC-VDD2V5-z16100          2501 mV
MPC-VDD1V8-z12004          1801 mV
MPC-AVDD1V0-z12006          996 mV
MPC-VDD1V2-z16100          1199 mV
MPC-VDD1V5A-z12004         1493 mV
MPC-VDD1V5B-z12004         1498 mV
MPC-XF_0V9-z12006          996 mV
MPC-PCIE_1V0-z16100        1000 mV
MPC-LU0_1V0-z12004          994 mV
MPC-LU1_1V0-z12004          994 mV
MPC-LU2_1V0-z12004          992 mV
MPC-LU3_1V0-z12004          993 mV
MPC-12VA-BMR453            12003 mV
MPC-12VB-BMR453            12043 mV
MPC-PMB_1V1-z12006          1091 mV
MPC-PMB_1V2-z12106          1196 mV
MPC-XM_0V9-vt273m          899 mV
I2C Slave Revision          106

```

#### show chassis environment fpc (MX240, MX480, MX960 with Application Services Modular Line Card)

```

user@host>show chassis environment fpc 1
FPC 1 status:
State                       Online
Temperature Intake          36 degrees C / 96 degrees F
Temperature Exhaust A      39 degrees C / 102 degrees F
Temperature LU TSen        52 degrees C / 125 degrees F
Temperature LU Chip        54 degrees C / 129 degrees F
Temperature XM TSen        52 degrees C / 125 degrees F
Temperature XM Chip        60 degrees C / 140 degrees F
Temperature PCIE TSen      52 degrees C / 125 degrees F
Temperature PCIE Chip      69 degrees C / 156 degrees F
Power
MPC-BIAS3V3-z12106         3302 mV
MPC-VDD3V3-z16100          3325 mV
MPC-AVDD1V0-z16100         1007 mV
MPC-PCIE_1V0-z16100         904 mV
MPC-LU0_1V0-z12004          996 mV
MPC-VDD_1V5-z12004         1498 mV
MPC-12VA-BMR453            11733 mV

```

MPC-12VB-BMR453	11728 mV
MPC-XM_0V9-vt273m	900 mV
I2C Slave Revision	81

### show chassis environment fpc (T320, T640, and T1600 Routers)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                               Online
  Temperature Top                     42 degrees C / 107 degrees F
  Temperature Bottom                  36 degrees C / 96 degrees F
  Temperature MMB1                    39 degrees C / 102 degrees F
  Power:
    1.8 V                             1959 mV
    2.5 V                             2495 mV
    3.3 V                             3344 mV
    5.0 V                             5047 mV
    1.8 V bias                        1787 mV
    3.3 V bias                        3291 mV
    5.0 V bias                        4998 mV
    8.0 V bias                        7343 mV
  BUS Revision                        40
FPC 1 status:
  State                               Online
  Temperature Top                     42 degrees C / 107 degrees F
  Temperature Bottom                  39 degrees C / 102 degrees F
  Temperature MMB1                    40 degrees C / 104 degrees F
  Power:
    1.8 V                             1956 mV
    2.5 V                             2498 mV
    3.3 V                             3340 mV
    5.0 V                             5023 mV
    1.8 V bias                        1782 mV
    3.3 V bias                        3277 mV
    5.0 V bias                        4989 mV
    8.0 V bias                        7289 mV
  BUS Revision                        40
FPC 2 status:
  State                               Online
  Temperature Top                     43 degrees C / 109 degrees F
  Temperature Bottom                  39 degrees C / 102 degrees F
  Temperature MMB1                    41 degrees C / 105 degrees F
  Power:
    1.8 V                             1963 mV
    2.5 V                             2503 mV
    3.3 V                             3340 mV
    5.0 V                             5042 mV
    1.8 V bias                        1797 mV
    3.3 V bias                        3311 mV
    5.0 V bias                        5013 mV
    8.0 V bias                        7221 mV
  BUS Revision                        40
```

### show chassis environment fpc (T4000 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                               Online
  Fan Intake                          34 degrees C / 93 degrees F
  Fan Exhaust                         48 degrees C / 118 degrees F
  PMB                                 47 degrees C / 116 degrees F
```

LMB0	50 degrees C / 122 degrees F
LMB1	41 degrees C / 105 degrees F
LMB2	35 degrees C / 95 degrees F
PFE1 LU2	46 degrees C / 114 degrees F
PFE1 LU0	41 degrees C / 105 degrees F
PFE0 LU0	57 degrees C / 134 degrees F
XF1	47 degrees C / 116 degrees F
XF0	52 degrees C / 125 degrees F
XM1	41 degrees C / 105 degrees F
XM0	50 degrees C / 122 degrees F
PFE0 LU1	56 degrees C / 132 degrees F
PFE0 LU2	45 degrees C / 113 degrees F
PFE1 LU1	37 degrees C / 98 degrees F
Power 1	
1.0 V	991 mV
1.2 V bias	1195 mV
1.8 V	1788 mV
2.5 V	2483 mV
3.3 V	3289 mV
3.3 V bias	3299 mV
12.0 V A	10608 mV
12.0 V B	10637 mV
Power 2	
0.9 V	881 mV
0.9 V PFE0	916 mV
0.9 V PFE1	903 mV
1.0 V PFE0	1012 mV
1.0 V PFE1	1002 mV
1.1 V	1095 mV
1.5 V_0	1494 mV
1.5 V_1	1479 mV
Power 3	
1.0 V PFE0	1000 mV
1.0 V PFE1	1002 mV
1.0 V PFE0 *	995 mV
1.0 V PFE1 *	995 mV
1.8 V PFE 0	1788 mV
1.8 V PFE 1	1789 mV
2.5 V	2482 mV
12.0 V	11614 mV
Power 4	
1.0 V PFE0 LU0	1003 mV
1.0 V PFE1 LU0	1003 mV
1.0 V PFE1 LU2	1004 mV
1.0 V PFE0 LU0 *	995 mV
1.0 V PFE1 LU0 *	998 mV
1.0 V PFE1 LU2 *	996 mV
12.0 V	11643 mV
12.0 V C	11711 mV
Power (Base/PMB/MMB)	
LMB0 VDD2V5	2488 mV
LMB0 VDD1V8	1788 mV
LMB0 VDD1V5	1496 mV
LMB0 PFE0 LU0 AVDD1V0	1002 mV
LMB0 PFE0 LU0 VDD1V0	1000 mV
LMB0 VDD12V0	10752 mV
LMB1 VDD2V5	2472 mV
LMB1 VDD1V8	1792 mV
LMB1 VDD1V5	1480 mV
LMB1 PFE0 LU2 AVDD1V0	994 mV
LMB1 PFE0 LU2 VDD1V0	1002 mV

```

LMB1 VDD12V0          10800 mV
LMB2 VDD2V5           2472 mV
LMB2 VDD1V8           1792 mV
LMB2 VDD1V5           1486 mV
LMB2 PFE1 LU1 AVDD1V0  996 mV
LMB2 PFE1 LU1 VDD1V0   998 mV
LMB2 VDD12V0          10704 mV
PMB 1.05v             1049 mV
PMB 1.5v              1500 mV
PMB 2.5v              2500 mV
PMB 3.3v              3299 mV
Bus Revision          113
FPC 3 status:
State                 Online
Fan Intake            37 degrees C / 98 degrees F
Fan Exhaust           51 degrees C / 123 degrees F
PMB                   43 degrees C / 109 degrees F
LMB0                  57 degrees C / 134 degrees F
LMB1                  54 degrees C / 129 degrees F
LMB2                  38 degrees C / 100 degrees F
PFE1 LU2              63 degrees C / 145 degrees F
PFE1 LU0              45 degrees C / 113 degrees F
PFE0 LU0              69 degrees C / 156 degrees F
XF1                   62 degrees C / 143 degrees F
XF0                   63 degrees C / 145 degrees F
XM1                   43 degrees C / 109 degrees F
XM0                   67 degrees C / 152 degrees F
PFE0 LU1              63 degrees C / 145 degrees F
PFE0 LU2              66 degrees C / 150 degrees F
PFE1 LU1              41 degrees C / 105 degrees F
Power 1
  1.0 V                1002 mV
  1.2 V bias           1201 mV
  1.8 V                1785 mV
  2.5 V                2485 mV
  3.3 V                3288 mV
  3.3 V bias           3285 mV
  12.0 V A             10412 mV
  12.0 V B             10515 mV
Power 2
  0.9 V                882 mV
  0.9 V PFE0           920 mV
  0.9 V PFE1           905 mV
  1.0 V PFE0           1015 mV
  1.0 V PFE1           1001 mV
  1.1 V                1094 mV
  1.5 V_0              1495 mV
  1.5 V_1              1478 mV
Power 3
  0.92 V PFE1          998 mV
  1.0 V PFE0           997 mV
  1.0 V PFE0 *         992 mV
  1.0 V PFE1 *         991 mV
  1.8 V PFE 0          1780 mV
  1.8 V PFE 1          1797 mV
  2.5 V                2492 mV
  12.0 V               11604 mV
Power 4
  1.0 V PFE0 LU0       1003 mV
  1.0 V PFE1 LU0       1004 mV
  1.0 V PFE1 LU2       1003 mV

```

```

1.0 V PFE0 LU0 *      1000 mV
1.0 V PFE1 LU0 *      1001 mV
1.0 V PFE1 LU2 *      1003 mV
12.0 V                 11653 mV
12.0 V C               11672 mV
Power (Base/PMB/MMB)
LMB0 VDD2V5           2512 mV
LMB0 VDD1V8           1790 mV
LMB0 VDD1V5           1500 mV
LMB0 PFE0 LU0 AVDD1V0 1004 mV
LMB0 PFE0 LU0 VDD1V0  1002 mV
LMB0 VDD12V0          10608 mV
LMB1 VDD2V5           2472 mV
LMB1 VDD1V8           1788 mV
LMB1 VDD1V5           1480 mV
LMB1 PFE0 LU2 AVDD1V0 1000 mV
LMB1 PFE0 LU2 VDD1V0  1004 mV
LMB1 VDD12V0          10672 mV
LMB2 VDD2V5           2488 mV
LMB2 VDD1V8           1798 mV
LMB2 VDD1V5           1494 mV
LMB2 PFE1 LU1 AVDD1V0 1000 mV
LMB2 PFE1 LU1 VDD1V0  1004 mV
LMB2 VDD12V0          10528 mV
PMB 1.05v             1050 mV
PMB 1.5v              1500 mV
PMB 2.5v              2499 mV
PMB 3.3v              3299 mV
Bus Revision           113
FPC 5 status:
State                  Online
Temperature Top         39 degrees C / 102 degrees F
Temperature Bottom      38 degrees C / 100 degrees F
Power
1.8 V                  1804 mV
1.8 V bias             1802 mV
3.3 V                  3294 mV
3.3 V bias             3277 mV
5.0 V bias             5008 mV
5.0 V TOP              5067 mV
8.0 V bias             6642 mV
Power (Base/PMB/MMB)
1.2 V                  1202 mV
1.5 V                  1504 mV
5.0 V BOT              5079 mV
12.0 V TOP Base        11848 mV
12.0 V BOT Base        11780 mV
1.1 V PMB              1111 mV
1.2 V PMB              1189 mV
1.5 V PMB              1494 mV
1.8 V PMB              1819 mV
2.5 V PMB              2503 mV
3.3 V PMB              3294 mV
5.0 V PMB              5035 mV
12.0 V PMB             11788 mV
0.75 MMB TOP           766 mV
1.5 V MMB TOP          1484 mV
1.8 V MMB TOP          1772 mV
2.5 V MMB TOP          2485 mV
1.2 V MMB TOP          1137 mV
5.0 V MMB TOP          4946 mV

```

12.0 V MMB TOP	11772 mV
3.3 V MMB TOP	3289 mV
0.75 MMB BOT	759 mV
1.5 V MMB BOT	1482 mV
1.8 V MMB BOT	1792 mV
2.5 V MMB BOT	2490 mV
1.2 V MMB BOT	1145 mV
5.0 V MMB BOT	4922 mV
12.0 V MMB BOT	11625 mV
3.3 V MMB BOT	3282 mV
APS 00	2495 mV
APS 01	3308 mV
APS 02	3301 mV
5.0 V PIC 0	4967 mV
APS 10	2512 mV
APS 11	3316 mV
APS 12	3304 mV
5.0 V PIC 1	5081 mV
Bus Revision	49
FPC 6 status:	
State	Online
Fan Intake	34 degrees C / 93 degrees F
Fan Exhaust	49 degrees C / 120 degrees F
PMB	40 degrees C / 104 degrees F
LMB0	60 degrees C / 140 degrees F
LMB1	58 degrees C / 136 degrees F
LMB2	40 degrees C / 104 degrees F
PFE1 LU2	69 degrees C / 156 degrees F
PFE1 LU0	45 degrees C / 113 degrees F
PFE0 LU0	71 degrees C / 159 degrees F
XF1	58 degrees C / 136 degrees F
XF0	65 degrees C / 149 degrees F
XM1	40 degrees C / 104 degrees F
XM0	66 degrees C / 150 degrees F
PFE0 LU1	69 degrees C / 156 degrees F
PFE0 LU2	68 degrees C / 154 degrees F
PFE1 LU1	42 degrees C / 107 degrees F
Power 1	
1.0 V	998 mV
1.2 V bias	1191 mV
1.8 V	1781 mV
2.5 V	2487 mV
3.3 V	3302 mV
3.3 V bias	3300 mV
12.0 V A	10388 mV
12.0 V B	10388 mV
Power 2	
0.9 V	902 mV
0.9 V PFE0	921 mV
0.9 V PFE1	907 mV
1.0 V PFE0	996 mV
1.0 V PFE1	974 mV
1.1 V	1095 mV
1.5 V_0	1495 mV
1.5 V_1	1478 mV
Power 3	
1.0 V PFE0	997 mV
1.0 V PFE1	998 mV
1.0 V PFE0 *	993 mV
1.0 V PFE1 *	991 mV
1.8 V PFE 0	1796 mV



```

1.8 V PFE 1          1789 mV
2.5 V                2465 mV
12.0 V              11609 mV
Power 4
1.0 V PFE0 LU0       1003 mV
1.0 V PFE1 LU0       1006 mV
1.0 V PFE1 LU2       1002 mV
1.0 V PFE0 LU0 *     1000 mV
1.0 V PFE1 LU0 *      998 mV
1.0 V PFE1 LU2 *      998 mV
12.0 V              11638 mV
12.0 V C             11702 mV
Power (Base/PMB/MMB)
LMB0 VDD2V5          2484 mV
LMB0 VDD1V8          1780 mV
LMB0 VDD1V5          1496 mV
LMB0 PFE0 LU0 AVDD1V0 998 mV
LMB0 PFE0 LU0 VDD1V0  1004 mV
LMB0 VDD12V0         10528 mV
LMB1 VDD2V5          2472 mV
LMB1 VDD1V8          1776 mV
LMB1 VDD1V5          1474 mV
LMB1 PFE0 LU2 AVDD1V0 994 mV
LMB1 PFE0 LU2 VDD1V0  1004 mV
LMB1 VDD12V0         10544 mV
LMB2 VDD2V5          2476 mV
LMB2 VDD1V8          1790 mV
LMB2 VDD1V5          1492 mV
LMB2 PFE1 LU1 AVDD1V0 996 mV
LMB2 PFE1 LU1 VDD1V0  1010 mV
LMB2 VDD12V0         10528 mV
PMB 1.05v            1050 mV
PMB 1.5v              1499 mV
PMB 2.5v              2500 mV
PMB 3.3v              3300 mV
Bus Revision          80

```

### show chassis environment fpc lcc (TX Matrix Router)

```

user@host> show chassis environment fpc lcc 0
lcc0-re0:
-----
FPC 1 status:
State                Online
Temperature Top       30 degrees C / 86 degrees F
Temperature Bottom    25 degrees C / 77 degrees F
Temperature MMBO       Absent
Temperature MMB1       27 degrees C / 80 degrees F
Power:
1.8 V                1813 mV
2.5 V                2504 mV
3.3 V                3338 mV
5.0 V                5037 mV
1.8 V bias           1797 mV
3.3 V bias           3301 mV
5.0 V bias           5013 mV
8.0 V bias           7345 mV
BUS Revision          40
FPC 2 status:
State                Online
Temperature Top       37 degrees C / 98 degrees F

```

Temperature Bottom	26 degrees C / 78 degrees F
Temperature MMB0	32 degrees C / 89 degrees F
Temperature MMB1	27 degrees C / 80 degrees F
Power:	
1.8 V	1791 mV
2.5 V	2517 mV
3.3 V	3308 mV
5.0 V	5052 mV
1.8 V bias	1797 mV
3.3 V bias	3289 mV
5.0 V bias	4991 mV
8.0 V bias	7477 mV
BUS Revision	40

#### show chassis environment fpc lcc (TX Matrix Plus Router)

```
user@host> show chassis environment fpc lcc 0
lcc0-re0:
```

```
-----
FPC 1 status:
State                               Online
Temperature Top                     46 degrees C / 114 degrees F
Temperature Bottom                   47 degrees C / 116 degrees F
Power
  1.8 V                             1788 mV
  1.8 V bias                         1787 mV
  3.3 V                             3321 mV
  3.3 V bias                         3306 mV
  5.0 V bias                         5018 mV
  5.0 V TOP                          5037 mV
  8.0 V bias                         7223 mV
Power (Base/PMB/MMB)
  1.2 V                             1205 mV
  1.5 V                             1503 mV
  5.0 V BOT                          5084 mV
  12.0 V TOP Base                    11775 mV
  12.0 V BOT Base                    11794 mV
  1.1 V PMB                          1108 mV
  1.2 V PMB                          1196 mV
  1.5 V PMB                          1499 mV
  1.8 V PMB                          1811 mV
  2.5 V PMB                          2515 mV
  3.3 V PMB                          3318 mV
  5.0 V PMB                          5030 mV
  12.0 V PMB                         11832 mV
  0.75 MMB TOP                       752 mV
  1.5 V MMB TOP                      1489 mV
  1.8 V MMB TOP                      1782 mV
  2.5 V MMB TOP                      2498 mV
  1.2 V MMB TOP                      1155 mV
  5.0 V MMB TOP                      4902 mV
  12.0 V MMB TOP                     11721 mV
  3.3 V MMB TOP                      3316 mV
  0.75 MMB BOT                       754 mV
  1.5 V MMB BOT                      1482 mV
  1.8 V MMB BOT                      1758 mV
  2.5 V MMB BOT                      2488 mV
  1.2 V MMB BOT                      1157 mV
  5.0 V MMB BOT                      4962 mV
  12.0 V MMB BOT                     11691 mV
  3.3 V MMB BOT                      3308 mV
```

APS 00	1484 mV
APS 01	2503 mV
APS 02	3313 mV
5.0 V PIC 0	5025 mV
APS 10	1501 mV
APS 11	2466 mV
APS 12	3311 mV
5.0 V PIC 1	5081 mV
Bus Revision	49

#### show chassis environment fpc (QFX Series)

```
user@switch> show chassis environment fpc 0
FPC 0 status:
  State                Online
  Temperature          42 degrees C / 107 degrees F
```

#### show chassis environment fpc interconnect-device (QFabric Systems)

```
user@switch> show chassis environment fpc interconnect-device interconnect1 0
FC 0 FPC 0 status:
  State                Online
  Left Intake Temperature 24 degrees C / 75 degrees F
  Right Intake Temperature 24 degrees C / 75 degrees F
  Left Exhaust Temperature 27 degrees C / 80 degrees F
  Right Exhaust Temperature 27 degrees C / 80 degrees F
  Power
    BIAS 3V3            3330 mV
    VDD 3V3             3300 mV
    VDD 2V5             2502 mV
    VDD 1V5             1496 mV
    VDD 1V2             1194 mV
    VDD 1V0             1000 mV
    SW0 VDD 1V0         1020 mV
    SW0 CVDD 1V025      1032 mV
    SW1 VDD 1V0         1022 mV
    SW1 CVDD 1V025      1030 mV
    VDD 12V0 DIV3_33    3414 mV
```

#### show chassis environment fpc 0 (PTX5000 Packet Transport Router)

```
user@host> show chassis environment fpc 0
FPC 0 status:
  State                Online
  PMB Temperature      35 degrees C / 95 degrees F
  Intake Temperature   33 degrees C / 91 degrees F
  Exhaust A Temperature 51 degrees C / 123 degrees F
  Exhaust B Temperature 43 degrees C / 109 degrees F
  TL0 Temperature      48 degrees C / 118 degrees F
  TQ0 Temperature      53 degrees C / 127 degrees F
  TL1 Temperature      56 degrees C / 132 degrees F
  TQ1 Temperature      58 degrees C / 136 degrees F
  TL2 Temperature      55 degrees C / 131 degrees F
  TQ2 Temperature      57 degrees C / 134 degrees F
  TL3 Temperature      59 degrees C / 138 degrees F
  TQ3 Temperature      59 degrees C / 138 degrees F
  Power
    PMB 1.05v          1049 mV
    PMB 1.5v           1500 mV
    PMB 2.5v           2500 mV
    PMB 3.3v           3299 mV
    PFE0 1.5v          1500 mV
```

PFE0	1.0v	999 mV
TQ0	0.9v	900 mV
TL0	0.9v	900 mV
PFE1	1.5v	1499 mV
PFE1	1.0v	999 mV
TQ1	0.9v	899 mV
TL1	0.9v	900 mV
PFE2	1.5v	1500 mV
PFE2	1.0v	1000 mV
TQ2	0.9v	900 mV
TL2	0.9v	900 mV
PFE3	1.5v	1499 mV
PFE3	1.0v	1000 mV
TQ3	0.9v	900 mV
TL3	0.9v	900 mV
Bias	3.3v	3327 mV
FPC	3.3v	3300 mV
FPC	2.5v	2500 mV
SAM	0.9v	900 mV
A	12.0v	2014 mV
B	12.0v	2030 mV

**show chassis environment fpc 07 (PTX5000 Packet Transport Router with FPC2-PTX-PIA)**

```
user@host> show chassis environment fpc 07
```

```
FPC 7 status:
```

State	Online
PMB TEMP0 Temperature	32 degrees C / 89 degrees F
PMB TEMP1 Temperature	28 degrees C / 82 degrees F
PMB CPU Temperature	46 degrees C / 114 degrees F
Intake Temperature	35 degrees C / 95 degrees F
Exhaust A Temperature	55 degrees C / 131 degrees F
Exhaust B Temperature	54 degrees C / 129 degrees F
TL5 Temperature	59 degrees C / 138 degrees F
TQ5 Temperature	57 degrees C / 134 degrees F
TL6 Temperature	57 degrees C / 134 degrees F
TQ6 Temperature	51 degrees C / 123 degrees F
TL1 Temperature	76 degrees C / 168 degrees F
TQ1 Temperature	58 degrees C / 136 degrees F
TL2 Temperature	75 degrees C / 167 degrees F
TQ2 Temperature	57 degrees C / 134 degrees F
TL4 Temperature	52 degrees C / 125 degrees F
TQ4 Temperature	66 degrees C / 150 degrees F
TL7 Temperature	52 degrees C / 125 degrees F
TQ7 Temperature	60 degrees C / 140 degrees F
TL0 Temperature	72 degrees C / 161 degrees F
TQ0 Temperature	73 degrees C / 163 degrees F
TL3 Temperature	64 degrees C / 147 degrees F
TQ3 Temperature	70 degrees C / 158 degrees F
Power	
PMB 1.05v	1049 mV
PMB 3.3v	3299 mV
PMB 1.1v-a	1100 mV
PMB 1.5v	1499 mV
PMB 1.1v-b	1100 mV
Base 3.3v	3300 mV
FPC Base 2.5v	2499 mV
TL1 0.9v	897 mV
TQ1 0.9v	897 mV
PFE1 1.0v	999 mV
PFE1 1.5v	1499 mV

TL2	0.9v	897 mV
TQ2	0.9v	897 mV
PFE2	1.0v	999 mV
PFE2	1.5v	1499 mV
FPC Base	1.0v	1000 mV
FPC Base	1.2v	1199 mV
TL5	0.9v	898 mV
TQ5	0.9v	898 mV
PFE5	1.0v	1000 mV
PFE5	1.5v	1500 mV
TL6	0.9v	897 mV
TQ6	0.9v	897 mV
PFE6	1.0v	1000 mV
PFE6	1.5v	1499 mV
Mezz Base	2.5v	2500 mV
TL0	0.9v	896 mV
TQ0	0.9v	896 mV
PFE0	1.0v	999 mV
PFE0	1.5v	1499 mV

### show chassis environment FPC 1 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis environment fpc 1
```

```
FPC 1 status:
```

State	Online
Temperature Intake	36 degrees C / 96 degrees F
Temperature Exhaust A	39 degrees C / 102 degrees F
Temperature LU TSen	52 degrees C / 125 degrees F
Temperature LU Chip	54 degrees C / 129 degrees F
Temperature XM TSen	52 degrees C / 125 degrees F
Temperature XM Chip	60 degrees C / 140 degrees F
Temperature PCIE TSen	52 degrees C / 125 degrees F
Temperature PCIE Chip	69 degrees C / 156 degrees F
Power	
MPC-BIAS3V3-z12106	3302 mV
MPC-VDD3V3-z16100	3325 mV
MPC-AVDD1V0-z16100	1007 mV
MPC-PCIE_1V0-z16100	904 mV
MPC-LU0_1V0-z12004	996 mV
MPC-VDD_1V5-z12004	1498 mV
MPC-12VA-BMR453	11733 mV
MPC-12VB-BMR453	11728 mV
MPC-XM_0V9-vt273m	900 mV
I2C Slave Revision	81

## show chassis environment routing-engine

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1014</a> <a href="#">Syntax (TX Matrix Routers) on page 1014</a> <a href="#">Syntax (TX Matrix Plus Routers) on page 1014</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 1014</a> <a href="#">Syntax (MX Series Routers) on page 1014</a> <a href="#">Syntax (QFX Series) on page 1014</a>
<b>Syntax</b>	show chassis environment routing-engine <slot>
<b>Syntax (TX Matrix Routers)</b>	show chassis environment routing-engine <lcc number   scc> <slot>
<b>Syntax (TX Matrix Plus Routers)</b>	show chassis environment routing-engine <lcc number   sfc number> <slot>
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)</b>	show chassis environment routing-engine <slot>
<b>Syntax (MX Series Routers)</b>	show chassis environment routing-engine <slot> <all-members> <local> <member member-id>
<b>Syntax (QFX Series)</b>	show chassis environment routing-engine interconnect-device <i>name</i>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Routers. Command introduced in Junos OS Release 12.1 for the T4000 Core Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
<b>Description</b>	Display Routing Engine environmental status information.
<b>Options</b>	<b>none</b> —Display environmental information about all Routing Engines. For a TX Matrix router, display environmental information about all Routing Engines on the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display

environmental information about all Routing Engines on the TX Matrix Plus router and its attached routers.

**all-members**—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in all member routers in the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display environmental information about the Routing Engines for the Interconnect device.

**lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in the specified member in the Virtual Chassis configuration. Replace *member-id* with the value of 0 or 1.

**scc**—(TX Matrix router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix router (switch-card chassis).

**sfc**—(TX Matrix Plus router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix Plus router (or switch-fabric chassis).

**slot**—(Optional) Display environmental information about an individual Routing Engine. On M10i, M20, M40e, M120, M160, M320, MX Series, MX104 routers, MX2010 routers, MX2020 routers, and T Series routers, replace *slot* with **0** or **1**. On M5, M7i, M10, and M40 routers and on the J Series router, replace *slot* with **0**. On EX3200 and EX4200 standalone switches, replace *slot* with **0**. On EX4200 switches in a Virtual Chassis configuration and on EX8208 and EX8216 switches, replace *slot* with **0** or **1**. On the QFX3500 switch, there is only one Routing Engine, so you do not need to specify the slot number. On PTX Series Packet Transport Routers, replace *slot* with **0** or **1**.

**Required Privilege Level** view

**Related Documentation**

- [request chassis routing-engine master on page 96](#)
- [show chassis routing-engine on page 1227](#)

**List of Sample Output**

- [show chassis environment routing-engine \(Nonredundant\) on page 1016](#)
- [show chassis environment routing-engine \(Redundant\) on page 1016](#)
- [show chassis environment routing-engine \(MX104 Router\) on page 1017](#)
- [show chassis environment routing-engine \(MX2010 Router\) on page 1017](#)
- [show chassis environment routing-engine \(MX2020 Router\) on page 1017](#)
- [show chassis environment routing-engine \(TX Matrix Plus Router\) on page 1017](#)
- [show chassis environment routing-engine \(T4000 Core Router\) on page 1017](#)
- [show chassis environment routing-engine \(QFX Series\) on page 1018](#)
- [show chassis environment routing-engine interconnect-device \(QFabric System\) on page 1018](#)
- [show chassis environment routing-engine \(PTX5000 Packet Transport Router\) on page 1018](#)

**Output Fields** Table 109 on page 1016 lists the output fields for the **show chassis environment routing-engine** command. Output fields are listed in the approximate order in which they appear.

Table 109: show chassis environment routing-engine Output Fields

Field Name	Field Description
Routing engine <i>slot</i> status	Number of the Routing Engine slot: 0 or 1.
State	Status of the Routing Engine: <ul style="list-style-type: none"> <li>• <b>Online Master</b>—Routing Engine is online, operating as Master.</li> <li>• <b>Online Standby</b>—Routing Engine is online, operating as Standby.</li> <li>• <b>Offline</b>—Routing Engine is offline.</li> </ul>
Temperature	Temperature of the air flowing past the Routing Engine.
CPU Temperature	(PTX Series and T4000 Core Routers only) Temperature of the air flowing past the Routing Engine CPU.

## Sample Output

### show chassis environment routing-engine (Nonredundant)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                Online Master
  Temperature          27 degrees C / 80 degrees
```

### show chassis environment routing-engine (Redundant)

```
user@host> show chassis environment routing-engine
Route Engine 0 status:
  State:                Online Master
  Temperature:          26 degrees C / 78 degrees F
Route Engine 1 status:
  State:                Online Standby
  Temperature:          26 degrees C / 78 degrees F
```



**show chassis environment routing-engine (MX104 Router)**

```

user@ host >show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      34 degrees C / 93 degrees F
  CPU Temperature  43 degrees C / 109 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      33 degrees C / 91 degrees F
  CPU Temperature  39 degrees C / 102 degrees F

```

**show chassis environment routing-engine (MX2010 Router)**

```

user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      37 degrees C / 98 degrees F
  CPU Temperature  37 degrees C / 98 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      35 degrees C / 95 degrees F
  CPU Temperature  34 degrees C / 93 degrees F

```

**show chassis environment routing-engine (MX2020 Router)**

```

user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      35 degrees C / 95 degrees F
  CPU Temperature  34 degrees C / 93 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      44 degrees C / 111 degrees F
  CPU Temperature  43 degrees C / 109 degrees F

```

**show chassis environment routing-engine (TX Matrix Plus Router)**

```

user@host> show chassis environment routing-engine
sfc0-re0:
-----
Routing Engine 0 status:
  State           Online Master
  Temperature      26 degrees C / 78 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      28 degrees C / 82 degrees F

lcc0-re0:
-----
Routing Engine 0 status:
  State           Online Master
  Temperature      30 degrees C / 86 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      29 degrees C / 84 degrees F

```

**show chassis environment routing-engine (T4000 Core Router)**

```

user@host> show chassis environment routing-engine

```

```
Routing Engine 0 status:
  State           Online Master
  Temperature      33 degrees C / 91 degrees F
  CPU Temperature  50 degrees C / 122 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      33 degrees C / 91 degrees F
  CPU Temperature  46 degrees C / 114 degrees F
```

#### show chassis environment routing-engine (QFX Series)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      42 degrees C / 107 degrees F
```

#### show chassis environment routing-engine interconnect-device (QFabric System)

```
user@switch> show chassis environment routing-engine interconnect-device interconnect1
routing-engine interconnect-device interconnect1
Routing Engine 0 status:
  State           Online Standby
  Temperature      52 degrees C / 125 degrees F
Routing Engine 1 status:
  State           Online Master
  Temperature      57 degrees C / 134 degrees F
```

#### show chassis environment routing-engine (PTX5000 Packet Transport Router)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      55 degrees C / 131 degrees F
  CPU Temperature  66 degrees C / 150 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      52 degrees C / 125 degrees F
  CPU Temperature  64 degrees C / 147 degrees F
```

## show chassis ethernet-switch

<b>List of Syntax</b>	<a href="#">Syntax on page 1019</a> <a href="#">Syntax (EX8200 Switch) on page 1019</a> <a href="#">Syntax (T4000 Router) on page 1019</a> <a href="#">Syntax (TX Matrix Router) on page 1019</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1019</a> <a href="#">Syntax (MX Series Router) on page 1019</a> <a href="#">Syntax (MX2010 and MX2020 3D Universal Edge Routers) on page 1019</a> <a href="#">Syntax (PTX Series Packet Transport Routers) on page 1019</a>
<b>Syntax</b>	show chassis ethernet-switch <errors <port>>
<b>Syntax (EX8200 Switch)</b>	show chassis ethernet-switch <statistics <port>   switch <number>
<b>Syntax (T4000 Router)</b>	show chassis ethernet-switch <errors <port>   statistics <port>>
<b>Syntax (TX Matrix Router)</b>	show chassis ethernet-switch <errors <port>   statistics <port>> <lcc <number>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show chassis ethernet-switch <errors <port>   switch <number> <lcc number   sfc number> <statistics <port>   switch <number>
<b>Syntax (MX Series Router)</b>	show chassis ethernet-switch <all-members> <errors <port>> <local> <member member-id>
<b>Syntax (MX2010 and MX2020 3D Universal Edge Routers)</b>	show chassis ethernet-switch <errors <port>   statistics <port>> <old-rom-packet-count>
<b>Syntax (PTX Series Packet Transport Routers)</b>	show chassis ethernet-switch <errors <port>> <statistics <port>> <port-state <port>>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.4 for EX Series switches. sfc option introduced in Junos OS Release 9.6 for the TX Matrix Plus router. Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.

- Description** (M10i, M40e, M120, M160, M320, MX Series, and T Series routers and EX8200 and PTX Series routers only) Display information about the ports on the Control Board (CB) Ethernet switch.
- Options**
- none**—Display information about each connected port on the Ethernet switch. On a TX Matrix router, display information about each connected port on the Ethernet switch on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display information about each connected port on the Ethernet switch on the TX Matrix Plus router and its attached routers.
  - all-members**—(MX Series routers only) (Optional) Display information about the ports on the CB Ethernet switch on all the members of the Virtual Chassis configuration.
  - errors**—(Optional) Display the numbers and types of errors accumulated on all ports of the Ethernet switch.
  - errors *port***—(Optional) Display the numbers and types of errors accumulated on the specified port (0 through 15) of the Ethernet switch. On the TX Matrix router, replace ***port*** with a value from 0 through 15. On the TX Matrix Plus router and EX8200 switch, replace ***port*** with a value from 0 through 27. On the PTX Series Packet Transport Routers, replace ***port*** with a value from 0 through 25. On the T4000 routers, MX2020 routers, and MX2010 routers, replace ***port*** with a value from 0 through 27.
  - errors switch *number***—(TX Matrix Plus router only) (Optional) Display the numbers and types of errors accumulated on the specified switch. Replace ***number*** with a value from 0 through 2.
  - lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.  
Replace ***number*** with the following values depending on the LCC configuration:
    - 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
    - 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
    - 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
    - 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
  - local**—(MX Series routers only) (Optional) Display information about the ports on the CB Ethernet switch on the local Virtual Chassis member.
  - member *member-id***—(MX Series routers only) (Optional) Display information about the ports on the CB Ethernet switch on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

**old-rom-packet-count**—(MX 2020 Routers only) (Optional) Display information about installed linecards. A non-zero number indicates that the bootrom on that linecard needs to be updated.

**port-state**—(PTX Series only) (Optional) Display information about current port operation (**Blocking**, **Listening**, or **Disabled**).

**scc**—(TX Matrix router only) (Optional) Display information about the ports on the CB's Ethernet switch on the TX Matrix router (switch-card chassis).

**sfc number**—(TX Matrix Plus router only) (Optional) Display information about the ports on the CB's Ethernet switch on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with **0**.

**statistics**—(Optional) Display traffic statistics for each connected port on the Ethernet switch.

**statistics port**—(Optional) Display traffic statistics for the specified port on the Ethernet switch. On the TX Matrix router, replace *port* with a value from **0** through **25**. On the TX Matrix Plus router or EX8200 switch, replace *port* with a value from **0** through **27**. On the PTX Series Packet Transport Routers, replace *port* with a value from **0** through **25**. On the T4000 routers, MX2020 routers, and MX2010 routers, replace *port* with a value from **0** through **27**.

**statistics switch number**—(TX Matrix Plus routers and EX8200 switch only) (Optional) Display traffic statistics for the specified Ethernet switch number. On the TX Matrix Plus router and EX8216 switch, replace *number* with a value from **0** through **2**. On the EX8208 switch, replace *number* with a value from **0** through **1**.

**Required Privilege Level** view

#### List of Sample Output

[show chassis ethernet-switch on page 1026](#)  
[show chassis ethernet-switch \(MX480 Router with MPC4E\) on page 1026](#)  
[show chassis ethernet-switch \(MX2010 Router\) on page 1027](#)  
[show chassis ethernet-switch statistics \(MX2010 Router\) on page 1029](#)  
[show chassis ethernet-switch \(MX2020 Router\) on page 1036](#)  
[show chassis ethernet-switch statistics \(MX2020 Router\) on page 1038](#)  
[show chassis ethernet-switch \(MX2020 Router with MPC4E\) on page 1046](#)  
[show chassis ethernet-switch \(TX Matrix Router\) on page 1047](#)  
[show chassis ethernet-switch errors on page 1049](#)  
[show chassis ethernet-switch statistics on page 1049](#)  
[show chassis ethernet-switch errors \(TX Matrix Plus Router\) on page 1050](#)  
[show chassis ethernet-switch sfc errors \(TX Matrix Plus Router\) on page 1051](#)  
[show chassis ethernet-switch statistics \(TX Matrix Plus Router\) on page 1052](#)  
[show chassis ethernet-switch \(T4000 Router\) on page 1056](#)  
[show chassis ethernet-switch errors \(T4000 Router\) on page 1057](#)  
[show chassis ethernet-switch \(PTX5000 Packet Transport Router\) on page 1058](#)  
[show chassis ethernet-switch statistics \(PTX5000 Packet Transport Router\) on page 1059](#)

[show chassis ethernet-switch port-state \(PTX5000 Packet Transport Router\) on page 1062](#)

**Output Fields** [Table 110 on page 1022](#) lists the output fields for the **show chassis ethernet-switch** command. Output fields are listed in the approximate order in which they appear.

**Table 110: show chassis ethernet-switch Output Fields**

Field Name	Field Description
Link is good on port <i>n</i> connected to device	Information about the link between each port on the CB's Ethernet switch and one of the following devices:
or	<ul style="list-style-type: none"> <li>FPC0 (Flexible PIC Concentrator 0) through FPC7</li> <li>Local controller</li> </ul>
Link is good on Fast Ethernet port <i>n</i> connected to device	<ul style="list-style-type: none"> <li>Routing Engine</li> <li>Other Routing Engine (on a system with two Routing Engines)</li> <li>SPMB (Switch Processor Mezzanine Board)</li> </ul>
or	<ul style="list-style-type: none"> <li>(TX Matrix router only) LCC0 (line-card chassis 0) through LCC3</li> </ul>
Link is good on Gigabit Ethernet port <i>n</i> connected to device	
or	
Link is down on Gigabit Ethernet port connected to device	
Speed is	Speed at which the Ethernet link is running: <b>10 Mb</b> or <b>100 Mb</b> . When the device is <b>RE</b> or <b>Other RE</b> on the TX Matrix router, the speed is <b>1000 Mb</b> .  <b>NOTE:</b> Irrespective of the device, the speed is <b>1000 Mb</b> on the MX2010 and MX2020 routers.
Duplex is	Duplex type of the Ethernet link: <b>full</b> or <b>half</b> .
Autonegotiate is Enabled (or Disabled)	By default, built-in Fast Ethernet ports on a PIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the <b>no-concatenate</b> statement at the <b>[edit chassis]</b> hierarchy level, as described in the <i>Junos OS System Basics Configuration Guide</i> ).
Flow Control TX is Enabled (or Disabled)	(MX2010 routers, MX2020 routers, and PTX Series) Flow control in the transmit direction is enabled (or disabled). Flow control regulates the flow of packets from the switch to the remote side of the connection.
Flow Control RX is Enabled (or Disabled)	(MX2010 routers, MX2020 routers, and PTX Series) Flow control in the receive direction is enabled (or disabled). Flow control regulates the flow of packets from the remote side of the connection to the switch.
MLT3	Number of multilevel threshold-3 (MLT-3) Fast Ethernet errors detected.
<b>Accumulated error counts for port <i>n</i> connected to device FPC<i>n</i>: (error output only)</b>	
Lock	Number of lock errors detected.
Xmit	Number of transmission errors detected.

Table 110: show chassis ethernet-switch Output Fields (*continued*)

Field Name	Field Description
<b>ESD</b>	Number of electrostatic discharge (ESD) errors detected.
<b>False Carrier</b>	Number of false carrier errors detected. This number is increased by one if a FRU is removed.
<b>Disconnects</b>	Number of disconnect errors detected.
<b>FX mode</b>	Number of errors detected on an Ethernet link over optical fiber.
<b>Statistics for port <i>n</i> connected to device FPC<i>n</i> (statistics output only)</b>	
<b>TX Packets 64 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 64 octets transmitted.
<b>TX Packets 65 - 127 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 65 through 127 octets transmitted.
<b>TX Packets 128 - 255 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 128 through 255 octets transmitted.
<b>TX Packets 256 - 511 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 256 through 511 octets transmitted.
<b>TX Packets 512 - 1023 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 512 through 1023 octets transmitted.
<b>TX Packets 1024 - 1518 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 1024 through 1518 octets transmitted.
<b>TX Packets 1519 - 2047 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 1519 through 2047 octets transmitted.
<b>TX Packets 2048 - 4095 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 2048 through 4095 octets transmitted.
<b>TX Packets 4096 - 9216 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 4096 through 9216 octets transmitted.
<b>TX 1519 - 1522 Good Vlan frms</b>	(MX2010 and MX2020 routers) Number of transmitted frames of size 1519 through 1522 octets that are good VLAN frames.
<b>TX Octets</b>	Number of octets sent.
<b>TX Unicast packets</b>	Number of unicast packets sent.
<b>TX Multicast packets</b>	Number of multicast packets sent.
<b>TX Broadcast packets</b>	Number of broadcast packets sent.
<b>TX Single Collision frames</b>	(MX2010 and MX2020 routers) Number of packets sent after one collision.

Table 110: show chassis ethernet-switch Output Fields (*continued*)

Field Name	Field Description
<b>TX Mult. Collision frames</b>	(MX2010 and MX2020 routers) Number of packets sent after multiple collisions.
<b>TX Late collisions</b>	Number of packets aborted during sending because of collisions after 64 bytes.
<b>TX Excessive collisions</b>	Number of packets not sent because of too many collisions.
<b>TX Dropped packets</b>	Number of transmitted packets that were dropped.
<b>TX PAUSEMAC Ctrl Frames</b>	Number of Media Access Control (MAC) frames containing PAUSE commands that were sent.
<b>TX Oversize Packets</b>	Number of oversize packets that were sent.
<b>TX FCS Error Counter</b>	Number of packets discarded because of frame check sequence errors.
<b>TX Fragment Counter</b>	Number of fragmented packets sent.
<b>TX Byte Counter</b>	Number of bytes sent.
<b>TX Packet OK Counter</b>	Number of viable packets sent.
<b>TX Pause Packet Counter</b>	Number of PAUSE packets sent.
<b>RX Packets 64 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 64 octets received.
<b>RX Packets 65 - 127 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 65 through 127 octets received.
<b>RX Packets 128 - 255 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 128 through 255 octets received.
<b>RX Packets 256 - 511 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 256 through 511 octets received.
<b>RX Packets 512 - 1023 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 512 through 1023 octets received.
<b>RX Packets 1024 - 1518 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 1024 through 1518 octets received.
<b>RX Packets 1519 - 2047 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 1519 through 2047 octets received.
<b>RX Packets 2048 - 4095 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 2048 through 4095 octets received.



Table 110: show chassis ethernet-switch Output Fields (*continued*)

Field Name	Field Description
<b>RX Packets 4096 - 9216 Octets</b>	(MX2010 and MX2020 routers) Number of packets of size 4096 through 9216 octets received.
<b>RX Octets</b>	Number of octets received.
<b>RX Unicast packets</b>	Number of unicast packets received.
<b>RX Multicast packets</b>	Number of multicast packets received.
<b>RX Broadcast packets</b>	Number of broadcast packets received.
<b>RX FCS Errors</b>	Number of packets discarded because of frame check sequence errors.
<b>RX Alignment Errors</b>	Number of incomplete octets received.
<b>RX Dropped Packets</b>	Number of incoming packets that were dropped.
<b>RX Fragments</b>	Number of fragmented packets received.
<b>RX Symbol Errors</b>	Number of symbols received that the router did not correctly decode.
<b>RX MAC Control</b>	Number of Media Access Control (MAC) packets received.
<b>RX Oversize Packets</b>	Number of oversize packets received.
<b>RX Undersize Packets</b>	Number of undersize packets received.
<b>RX Jabbers</b>	Total number of frames received that exceed the maximum byte count and contain CRC errors .
<b>RX Control Frame Counter</b>	Number of control frames received.
<b>RX Pause Frame Counter</b>	Number of pause frames received.
<b>RX FCS Errors</b>	Number of packets discarded because of frame check sequence errors.
<b>RX Fragments</b>	Number of fragmented packets received.
<b>RX Byte Counter</b>	Number of bytes received.
<b>RX Packet OK Counter</b>	Number of viable packets received.

## Sample Output

### show chassis ethernet-switch

```
user@host> show chassis ethernet-switch
Link is good on port 0 connected to device: FPC0
  Speed is 100 MB
  Duplex is full

Link is good on port 1 connected to device: FPC1
  Speed is 100 MB
  Duplex is full

Link is good on port 2 connected to device: FPC2
  Speed is 100 MB
  Duplex is full

Link is good on port 3 connected to device: FPC3
  Speed is 100 MBb
  Duplex is full

Link is good on port 7 connected to device: Local controller
  Speed is 100 MB
  Duplex is full

Link is good on port 9 connected to device: SPMB
  Speed is 100 MB
  Duplex is full

Link is good on port 13 connected to device: FPC5
  Speed is 100 MB
  Duplex is full
```

### show chassis ethernet-switch (MX480 Router with MPC4E)

```
user@host > show chassis ethernet-switch
Displaying summary for switch 0
Link is down on GE port 0 connected to device: FPC0

Link is down on GE port 1 connected to device: FPC1

Link is good on GE port 2 connected to device: FPC2
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 3 connected to device: FPC3
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 4 connected to device: FPC4
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
```

```

Flow Control RX is Disabled

Link is down on GE port 5 connected to device: FPC5

Link is down on GE port 6 connected to device: FPC6

Link is down on GE port 7 connected to device: FPC7

Link is down on GE port 8 connected to device: FPC8

Link is down on GE port 9 connected to device: FPC9

Link is down on GE port 10 connected to device: FPC10

Link is down on GE port 11 connected to device: FPC11

Link is good on GE port 12 connected to device: Other RE
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on GE port 13 connected to device: RE-GigE
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is down on GE port 14 connected to device: Debug-GigE

```

#### show chassis ethernet-switch (MX2010 Router)

```

user@host > show chassis ethernet-switch
Displaying summary for switch 0
Link is good on GE port 0 connected to device: FPC0
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on GE port 1 connected to device: FPC1
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on GE port 2 connected to device: FPC3
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on GE port 3 connected to device: FPC2
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled

```

Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 4 connected to device: FPC5  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 5 connected to device: FPC4  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 6 connected to device: FPC6  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 7 connected to device: FPC7  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 8 connected to device: FPC8  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 9 connected to device: FPC9  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 20 connected to device: Other RE-GigE  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 21 connected to device: RE-GigE  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is down on GE port 22 connected to device: Debug-GigE

```

Link is good on GE port 23 connected to device: SPMB
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is down on XE port 24 connected to device: SFP+ 0

Link is down on XE port 25 connected to device: SFP+ 1

Link is down on XE port 26 connected to device: RE-10GigE

Link is down on XE port 27 connected to device: Other RE-10GigE

```

#### show chassis ethernet-switch statistics (MX2010 Router)

```

user@host > show chassis ethernet-switch statistics
Displaying port statistics for switch 0
Statistics for port 0 connected to device FPC0:
TX Packets 64 Octets      5088623
TX Packets 65-127 Octets  2637257
TX Packets 128-255 Octets 84829
TX Packets 256-511 Octets 120193
TX Packets 512-1023 Octets 252371
TX Packets 1024-1518 Octets 7189736
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets 15373009
TX Multicast Packets 14
TX Broadcast Packets 1679654
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions 0
TX Excessive Collisions 0
TX Collision frames 0
TX PAUSEMAC Ctrl Frames 0
TX MAC ctrl frames 0
TX Frame deferred Xmsns 0
TX Frame excessive deferl 0
TX Oversize Packets 0
TX Jabbers 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 3041239292
RX Packets 64 Octets 874260
RX Packets 65-127 Octets 26066124
RX Packets 128-255 Octets 1386532
RX Packets 256-511 Octets 150539
RX Packets 512-1023 Octets 4636799
RX Packets 1024-1518 Octets 92601
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 33206855
RX Multicast Packets 0
RX Broadcast Packets 279416
RX FCS Errors 0
RX Align Errors 0

```

```
RX Fragments                0
RX Symbol errors            0
RX Unsupported opcodes      0
RX Out of Range Length      0
RX False Carrier Errors     0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter       0
RX Control Frame Counter    0
RX Pause Frame Counter      0
RX Byte Counter             958929187
Statistics for port 1 connected to device FPC1:
TX Packets 64 Octets        5109146
TX Packets 65-127 Octets    2779473
TX Packets 128-255 Octets   2441286
TX Packets 256-511 Octets   173102
TX Packets 512-1023 Octets  1547504
TX Packets 1024-1518 Octets 7190581
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets                   19241092
TX Multicast Packets        14
TX Broadcast Packets        1673369
TX Single Collision frames  0
TX Mult. Collision frames   0
TX Late Collisions          0
TX Excessive Collisions     0
TX Collision frames         0
TX PAUSEMAC Ctrl Frames     0
TX MAC ctrl frames          0
TX Frame deferred Xtns      0
TX Frame excessive deferl    0
TX Oversize Packets         0
TX Jabbers                  0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             4213380187
RX Packets 64 Octets        865914
RX Packets 65-127 Octets    26612151
RX Packets 128-255 Octets   1090153
RX Packets 256-511 Octets   25126
RX Packets 512-1023 Octets  101158
RX Packets 1024-1518 Octets 78092
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                   28772594
RX Multicast Packets        0
RX Broadcast Packets        285669
RX FCS Errors               0
RX Align Errors             0
RX Fragments                0
RX Symbol errors            0
RX Unsupported opcodes      0
RX Out of Range Length      0
RX False Carrier Errors     0
RX Undersize Packets        0
```

```

RX Oversize Packets      0
RX Jabbers               0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter    0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter          2327283837

```

Link is down on GE port 2 connected to device: FPC3

Link is down on GE port 3 connected to device: FPC2

Link is down on GE port 4 connected to device: FPC5

Link is down on GE port 5 connected to device: FPC4

Link is down on GE port 6 connected to device: FPC6

Link is down on GE port 7 connected to device: FPC7

Statistics for port 8 connected to device FPC8:

```

TX Packets 64 Octets      5341094
TX Packets 65-127 Octets  2625310
TX Packets 128-255 Octets 3315158
TX Packets 256-511 Octets 174805
TX Packets 512-1023 Octets 976908
TX Packets 1024-1518 Octets 7181498
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets                  19614773
TX Multicast Packets       14
TX Broadcast Packets       1673831
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions         0
TX Excessive Collisions    0
TX Collision frames        0
TX PAUSEMAC Ctrl Frames    0
TX MAC ctrl frames         0
TX Frame deferred Xtns     0
TX Frame excessive deferl 0
TX Oversize Packets        0
TX Jabbers                 0
TX FCS Error Counter       0
TX Fragment Counter        0
TX Byte Counter            3946762991
RX Packets 64 Octets       955509
RX Packets 65-127 Octets   27568588
RX Packets 128-255 Octets  1460936
RX Packets 256-511 Octets  153248
RX Packets 512-1023 Octets 2856206
RX Packets 1024-1518 Octets 76419
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                  33070906
RX Multicast Packets       0
RX Broadcast Packets       285183
RX FCS Errors              0

```

RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	4256093824

## Statistics for port 9 connected to device FPC9:

TX Packets 64 Octets	5237213
TX Packets 65-127 Octets	3268775
TX Packets 128-255 Octets	2320476
TX Packets 256-511 Octets	1789844
TX Packets 512-1023 Octets	501022
TX Packets 1024-1518 Octets	7800455
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	20917785
TX Multicast Packets	14
TX Broadcast Packets	1673368
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xms	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	747012161
RX Packets 64 Octets	1036527
RX Packets 65-127 Octets	27590367
RX Packets 128-255 Octets	1590059
RX Packets 256-511 Octets	328257
RX Packets 512-1023 Octets	75975
RX Packets 1024-1518 Octets	73556
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	30694741
RX Multicast Packets	0
RX Broadcast Packets	285586
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0



RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	2727836941

Statistics for port 20 connected to device Other RE-GigE:

TX Packets 64 Octets	1682540
TX Packets 65-127 Octets	3454
TX Packets 128-255 Octets	659
TX Packets 256-511 Octets	0
TX Packets 512-1023 Octets	1
TX Packets 1024-1518 Octets	0
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	1686654
TX Multicast Packets	6
TX Broadcast Packets	1673798
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xms	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	108042476
RX Packets 64 Octets	710214
RX Packets 65-127 Octets	35785510
RX Packets 128-255 Octets	4616
RX Packets 256-511 Octets	232
RX Packets 512-1023 Octets	565
RX Packets 1024-1518 Octets	28798
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	36529935
RX Multicast Packets	8
RX Broadcast Packets	285546
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0

RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	2676440958

## Statistics for port 21 connected to device RE-GigE:

TX Packets 64 Octets	4805310
TX Packets 65-127 Octets	143798628
TX Packets 128-255 Octets	5532385
TX Packets 256-511 Octets	671059
TX Packets 512-1023 Octets	7684123
TX Packets 1024-1518 Octets	344021
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	162835526
TX Multicast Packets	8
TX Broadcast Packets	1673409
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xtns	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	105857355
RX Packets 64 Octets	14537137
RX Packets 65-127 Octets	11445505
RX Packets 128-255 Octets	8161767
RX Packets 256-511 Octets	2257944
RX Packets 512-1023 Octets	3277807
RX Packets 1024-1518 Octets	29373209
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	69053369
RX Multicast Packets	6
RX Broadcast Packets	285935
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	2980410755

Link is down on GE port 22 connected to device: Debug-GigE

Statistics for port 23 connected to device SPMB:

TX Packets 64 Octets	1885878
TX Packets 65-127 Octets	138845
TX Packets 128-255 Octets	18
TX Packets 256-511 Octets	1
TX Packets 512-1023 Octets	2
TX Packets 1024-1518 Octets	16391
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	2041135
TX Multicast Packets	14
TX Broadcast Packets	1707267
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xms	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	148066476
RX Packets 64 Octets	374994
RX Packets 65-127 Octets	183398
RX Packets 128-255 Octets	749
RX Packets 256-511 Octets	13658
RX Packets 512-1023 Octets	13421
RX Packets 1024-1518 Octets	9
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	586229
RX Multicast Packets	0
RX Broadcast Packets	252034
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	51431942

Link is down on XE port 24 connected to device: SFP+ 0

Link is down on XE port 25 connected to device: SFP+ 1

Link is down on XE port 26 connected to device: RE-10GigE

Link is down on XE port 27 connected to device: Other RE-10GigE

#### show chassis ethernet-switch (MX2020 Router)

```
user@host > show chassis ethernet-switch
```

Displaying summary for switch 0

Link is good on GE port 0 connected to device: FPC0

Speed is 1000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

Link is good on GE port 1 connected to device: FPC1

Speed is 1000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

Link is good on GE port 2 connected to device: FPC3

Speed is 1000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

Link is good on GE port 3 connected to device: FPC2

Speed is 1000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

Link is good on GE port 4 connected to device: FPC5

Speed is 1000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

Link is good on GE port 5 connected to device: FPC4

Speed is 1000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

Link is good on GE port 6 connected to device: FPC6

Speed is 1000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled

Flow Control RX is Disabled

Link is good on GE port 7 connected to device: FPC7

Speed is 1000Mb

Duplex is full

Autonegotiate is Enabled

Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 8 connected to device: FPC8  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 9 connected to device: FPC9  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 10 connected to device: FPC10  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 11 connected to device: FPC11  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 12 connected to device: FPC13  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 13 connected to device: FPC12  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 14 connected to device: FPC14  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 15 connected to device: FPC15  
Speed is 1000Mb  
Duplex is full  
Autonegotiate is Enabled  
Flow Control TX is Disabled  
Flow Control RX is Disabled

Link is good on GE port 16 connected to device: FPC17  
Speed is 1000Mb

```
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on GE port 17 connected to device: FPC16
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on GE port 18 connected to device: FPC18
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on GE port 19 connected to device: FPC19
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on GE port 20 connected to device: Other RE-GigE
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on GE port 21 connected to device: RE-GigE
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is down on GE port 22 connected to device: Debug-GigE

Link is good on GE port 23 connected to device: SPMB
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is down on XE port 24 connected to device: SFP+ 0

Link is down on XE port 25 connected to device: SFP+ 1

Link is down on XE port 26 connected to device: RE-10GigE

Link is down on XE port 27 connected to device: Other RE-10GigE
```

#### show chassis ethernet-switch statistics (MX2020 Router)

```
user@host > show chassis ethernet-switch statistics
```

```

Displaying port statistics for switch 0
Statistics for port 0 connected to device FPC0:
TX Packets 64 Octets      1468564
TX Packets 65-127 Octets  153896
TX Packets 128-255 Octets 237
TX Packets 256-511 Octets 286
TX Packets 512-1023 Octets 599
TX Packets 1024-1518 Octets 22803
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets      1646385
TX Multicast Packets      6
TX Broadcast Packets      970939
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions      0
TX Excessive Collisions  0
TX Collision frames      0
TX PAUSEMAC Ctrl Frames  0
TX MAC ctrl frames      0
TX Frame deferred Xmsns  0
TX Frame excessive deferl 0
TX Oversize Packets      0
TX Jabbers      0
TX FCS Error Counter      0
TX Fragment Counter      0
TX Byte Counter      130470290
RX Packets 64 Octets      180266
RX Packets 65-127 Octets  519030
RX Packets 128-255 Octets 1390
RX Packets 256-511 Octets 42857
RX Packets 512-1023 Octets 3482
RX Packets 1024-1518 Octets 8147
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets      755172
RX Multicast Packets      0
RX Broadcast Packets      42822
RX FCS Errors      0
RX Align Errors      0
RX Fragments      0
RX Symbol errors      0
RX Unsupported opcodes  0
RX Out of Range Length  0
RX False Carrier Errors  0
RX Undersize Packets      0
RX Oversize Packets      0
RX Jabbers      0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter      0
RX Control Frame Counter  0
RX Pause Frame Counter      0
RX Byte Counter      75374021
Statistics for port 1 connected to device FPC1:
TX Packets 64 Octets      1493739
TX Packets 65-127 Octets  126996
TX Packets 128-255 Octets  241
TX Packets 256-511 Octets  283

```

```
TX Packets 512-1023 Octets 604
TX Packets 1024-1518 Octets 33687
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets 1655550
TX Multicast Packets 6
TX Broadcast Packets 969032
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions 0
TX Excessive Collisions 0
TX Collision frames 0
TX PAUSEMAC Ctrl Frames 0
TX MAC ctrl frames 0
TX Frame deferred Xmsns 0
TX Frame excessive deferl 0
TX Oversize Packets 0
TX Jabbers 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 141832690
RX Packets 64 Octets 155655
RX Packets 65-127 Octets 545561
RX Packets 128-255 Octets 1394
RX Packets 256-511 Octets 42811
RX Packets 512-1023 Octets 3514
RX Packets 1024-1518 Octets 8171
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 757106
RX Multicast Packets 0
RX Broadcast Packets 44509
RX FCS Errors 0
RX Align Errors 0
RX Fragments 0
RX Symbol errors 0
RX Unsupported opcodes 0
RX Out of Range Length 0
RX False Carrier Errors 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 75691392
Statistics for port 2 connected to device FPC3:
TX Packets 64 Octets 1465749
TX Packets 65-127 Octets 152849
TX Packets 128-255 Octets 238
TX Packets 256-511 Octets 289
TX Packets 512-1023 Octets 602
TX Packets 1024-1518 Octets 38903
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
```



```

TX Octets                                1658630
TX Multicast Packets                     6
TX Broadcast Packets                     968873
TX Single Collision frames               0
TX Mult. Collision frames                0
TX Late Collisions                      0
TX Excessive Collisions                 0
TX Collision frames                     0
TX PAUSEMAC Ctrl Frames                 0
TX MAC ctrl frames                      0
TX Frame deferred Xms                   0
TX Frame excessive deferl               0
TX Oversize Packets                     0
TX Jabbers                              0
TX FCS Error Counter                    0
TX Fragment Counter                     0
TX Byte Counter                         147427010
RX Packets 64 Octets                     181636
RX Packets 65-127 Octets                 517526
RX Packets 128-255 Octets                1405
RX Packets 256-511 Octets                42806
RX Packets 512-1023 Octets               3515
RX Packets 1024-1518 Octets              8168
RX Packets 1519-2047 Octets              0
RX Packets 2048-4095 Octets              0
RX Packets 4096-9216 Octets              0
RX Octets                               755056
RX Multicast Packets                     0
RX Broadcast Packets                     44490
RX FCS Errors                           0
RX Align Errors                         0
RX Fragments                            0
RX Symbol errors                        0
RX Unsupported opcodes                   0
RX Out of Range Length                   0
RX False Carrier Errors                  0
RX Undersize Packets                     0
RX Oversize Packets                     0
RX Jabbers                              0
RX 1519-1522 Good Vlan frms             0
RX MTU Exceed Counter                    0
RX Control Frame Counter                 0
RX Pause Frame Counter                  0
RX Byte Counter                         75381869
Statistics for port 3 connected to device FPC2:
TX Packets 64 Octets                     1473828
TX Packets 65-127 Octets                 145643
TX Packets 128-255 Octets                 253
TX Packets 256-511 Octets                 285
TX Packets 512-1023 Octets               612
TX Packets 1024-1518 Octets              26603
TX Packets 1519-2047 Octets              0
TX Packets 2048-4095 Octets              0
TX Packets 4096-9216 Octets              0
TX 1519-1522 Good Vlan frms              0
TX Octets                               1647224
TX Multicast Packets                     6
TX Broadcast Packets                     968925
TX Single Collision frames               0
TX Mult. Collision frames                0
TX Late Collisions                      0

```

TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xms	0
TX Frame excessive deferl	0
TX Oversize Packets	0
TX Jabbers	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	134293832
RX Packets 64 Octets	174230
RX Packets 65-127 Octets	525756
RX Packets 128-255 Octets	1404
RX Packets 256-511 Octets	42815
RX Packets 512-1023 Octets	3530
RX Packets 1024-1518 Octets	8176
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Octets	755911
RX Multicast Packets	0
RX Broadcast Packets	44499
RX FCS Errors	0
RX Align Errors	0
RX Fragments	0
RX Symbol errors	0
RX Unsupported opcodes	0
RX Out of Range Length	0
RX False Carrier Errors	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX 1519-1522 Good Vlan frms	0
RX MTU Exceed Counter	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	75517355

Statistics for port 4 connected to device FPC5:

TX Packets 64 Octets	1466664
TX Packets 65-127 Octets	151155
TX Packets 128-255 Octets	238
TX Packets 256-511 Octets	277
TX Packets 512-1023 Octets	615
TX Packets 1024-1518 Octets	54674
TX Packets 1519-2047 Octets	0
TX Packets 2048-4095 Octets	0
TX Packets 4096-9216 Octets	0
TX 1519-1522 Good Vlan frms	0
TX Octets	1673623
TX Multicast Packets	6
TX Broadcast Packets	968610
TX Single Collision frames	0
TX Mult. Collision frames	0
TX Late Collisions	0
TX Excessive Collisions	0
TX Collision frames	0
TX PAUSEMAC Ctrl Frames	0
TX MAC ctrl frames	0
TX Frame deferred Xms	0
TX Frame excessive deferl	0

```

TX Oversize Packets      0
TX Jabbers               0
TX FCS Error Counter    0
TX Fragment Counter      0
TX Byte Counter          164247790
RX Packets 64 Octets     180006
RX Packets 65-127 Octets 518217
RX Packets 128-255 Octets 1406
RX Packets 256-511 Octets 42787
RX Packets 512-1023 Octets 3515
RX Packets 1024-1518 Octets 8164
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                754095
RX Multicast Packets     0
RX Broadcast Packets     44457
RX FCS Errors            0
RX Align Errors          0
RX Fragments             0
RX Symbol errors         0
RX Unsupported opcodes   0
RX Out of Range Length   0
RX False Carrier Errors  0
RX Undersize Packets     0
RX Oversize Packets      0
RX Jabbers               0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter    0
RX Control Frame Counter 0
RX Pause Frame Counter   0
RX Byte Counter          75311970
Statistics for port 5 connected to device FPC4:
TX Packets 64 Octets     1464770
TX Packets 65-127 Octets 154498
TX Packets 128-255 Octets 225
TX Packets 256-511 Octets 280
TX Packets 512-1023 Octets 637
TX Packets 1024-1518 Octets 26355
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets                1646765
TX Multicast Packets     6
TX Broadcast Packets     968730
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions       0
TX Excessive Collisions  0
TX Collision frames      0
TX PAUSEMAC Ctrl Frames  0
TX MAC ctrl frames       0
TX Frame deferred Xtns    0
TX Frame excessive deferl 0
TX Oversize Packets      0
TX Jabbers               0
TX FCS Error Counter     0
TX Fragment Counter      0
TX Byte Counter          134058606
RX Packets 64 Octets     169269

```

```
RX Packets 65-127 Octets      515285
RX Packets 128-255 Octets    1527
RX Packets 256-511 Octets    42804
RX Packets 512-1023 Octets   3521
RX Packets 1024-1518 Octets  9142
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                    741548
RX Multicast Packets         0
RX Broadcast Packets         44470
RX FCS Errors                0
RX Align Errors              0
RX Fragments                 0
RX Symbol errors             0
RX Unsupported opcodes       0
RX Out of Range Length       0
RX False Carrier Errors      0
RX Undersize Packets         0
RX Oversize Packets          0
RX Jabbers                   0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter        0
RX Control Frame Counter     0
RX Pause Frame Counter       0
RX Byte Counter              75498393
Statistics for port 6 connected to device FPC6:
TX Packets 64 Octets         1475260
TX Packets 65-127 Octets     143324
TX Packets 128-255 Octets    260
TX Packets 256-511 Octets    274
TX Packets 512-1023 Octets   603
TX Packets 1024-1518 Octets  40631
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan frms 0
TX Octets                    1660352
TX Multicast Packets         6
TX Broadcast Packets         968466
TX Single Collision frames   0
TX Mult. Collision frames    0
TX Late Collisions           0
TX Excessive Collisions      0
TX Collision frames          0
TX PAUSEMAC Ctrl Frames     0
TX MAC ctrl frames           0
TX Frame deferred Xtns       0
TX Frame excessive deferl    0
TX Oversize Packets          0
TX Jabbers                   0
TX FCS Error Counter         0
TX Fragment Counter          0
TX Byte Counter              149212764
RX Packets 64 Octets         172275
RX Packets 65-127 Octets     526519
RX Packets 128-255 Octets    1394
RX Packets 256-511 Octets    42777
RX Packets 512-1023 Octets   3514
RX Packets 1024-1518 Octets  8161
RX Packets 1519-2047 Octets  0
```

```

RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                    754640
RX Multicast Packets         0
RX Broadcast Packets         44443
RX FCS Errors                0
RX Align Errors              0
RX Fragments                 0
RX Symbol errors             0
RX Unsupported opcodes       0
RX Out of Range Length       0
RX False Carrier Errors      0
RX Undersize Packets         0
RX Oversize Packets          0
RX Jabbers                   0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter        0
RX Control Frame Counter     0
RX Pause Frame Counter       0
RX Byte Counter              75386517
Statistics for port 7 connected to device FPC7:
TX Packets 64 Octets         1472361
TX Packets 65-127 Octets     145646
TX Packets 128-255 Octets    251
TX Packets 256-511 Octets    250
TX Packets 512-1023 Octets   580
TX Packets 1024-1518 Octets  49530
TX Packets 1519-2047 Octets  0
TX Packets 2048-4095 Octets  0
TX Packets 4096-9216 Octets  0
TX 1519-1522 Good Vlan frms 0
TX Octets                    1668618
TX Multicast Packets         6
TX Broadcast Packets         968317
TX Single Collision frames   0
TX Mult. Collision frames    0
TX Late Collisions           0
TX Excessive Collisions      0
TX Collision frames          0
TX PAUSEMAC Ctrl Frames     0
TX MAC ctrl frames           0
TX Frame deferred Xtns       0
TX Frame excessive deferl    0
TX Oversize Packets          0
TX Jabbers                   0
TX FCS Error Counter         0
TX Fragment Counter          0
TX Byte Counter              158689814
RX Packets 64 Octets         174618
RX Packets 65-127 Octets     523421
RX Packets 128-255 Octets    1393
RX Packets 256-511 Octets    42764
RX Packets 512-1023 Octets   3514
RX Packets 1024-1518 Octets  8158
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Octets                    753868
RX Multicast Packets         0
RX Broadcast Packets         44429
RX FCS Errors                0

```

```
RX Align Errors          0
RX Fragments             0
RX Symbol errors         0
RX Unsupported opcodes   0
RX Out of Range Length   0
RX False Carrier Errors  0
RX Undersize Packets     0
RX Oversize Packets      0
RX Jabbers               0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter    0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter          75309863
Statistics for port 8 connected to device FPC8:
...
```

#### show chassis ethernet-switch (MX2020 Router with MPC4E)

```
user@ host > show chassis ethernet-switch
Displaying summary for switch 0
Link is good on GE port 0 connected to device: FPC0
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is down on GE port 1 connected to device: FPC1

Link is down on GE port 2 connected to device: FPC3

Link is down on GE port 3 connected to device: FPC2

Link is down on GE port 4 connected to device: FPC5

Link is down on GE port 5 connected to device: FPC4

Link is down on GE port 6 connected to device: FPC6

Link is down on GE port 7 connected to device: FPC7

Link is down on GE port 8 connected to device: FPC8

Link is good on GE port 9 connected to device: FPC9
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on GE port 10 connected to device: FPC10
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is down on GE port 11 connected to device: FPC11

Link is down on GE port 12 connected to device: FPC13
```

```
Link is down on GE port 13 connected to device: FPC12

Link is good on GE port 14 connected to device: FPC14
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 15 connected to device: FPC15

Link is down on GE port 16 connected to device: FPC17

Link is down on GE port 17 connected to device: FPC16

Link is down on GE port 18 connected to device: FPC18

Link is good on GE port 19 connected to device: FPC19
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 20 connected to device: Other RE-GigE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is good on GE port 21 connected to device: RE-GigE
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on GE port 22 connected to device: Debug-GigE

Link is good on GE port 23 connected to device: SPMB
  Speed is 1000Mb
  Duplex is full
  Autonegotiate is Enabled
  Flow Control TX is Disabled
  Flow Control RX is Disabled

Link is down on XE port 24 connected to device: SFP+ 0

Link is down on XE port 25 connected to device: SFP+ 1

Link is down on XE port 26 connected to device: RE-10GigE

Link is down on XE port 27 connected to device: Other RE-10GigE
```

#### show chassis ethernet-switch (TX Matrix Router)

```
user@host> show chassis ethernet-switch
scc-re0:
```

-----

Link is good on FE port 4 connected to device: LCC0  
Speed is 100 MB  
Duplex is full  
Autonegotiate is Enabled

Link is good on FE port 6 connected to device: LCC2  
Speed is 100 MB  
Duplex is full  
Autonegotiate is Enabled

Link is good on FE port 8 connected to device: SPMB  
Speed is 100 MB  
Duplex is full  
Autonegotiate is Enabled

lcc0-re0:

-----  
Link is good on FE port 1 connected to device: FPC1  
Speed is 100 MB  
Duplex is full  
Autonegotiate is Enabled

Link is good on FE port 2 connected to device: FPC2  
Speed is 100 MB  
Duplex is full  
Autonegotiate is Enabled

Link is good on FE port 8 connected to device: SPMB  
Speed is 100 MB  
Duplex is full  
Autonegotiate is Enabled

Link is good on FE port 10 connected to device: SCC  
Speed is 100 MB  
Duplex is full  
Autonegotiate is Enabled

lcc2-re0:

-----  
Link is good on FE port 0 connected to device: FPC0  
Speed is 100 MB  
Duplex is full  
Autonegotiate is Enabled

Link is good on FE port 1 connected to device: FPC1  
Speed is 100 MB  
Duplex is full  
Autonegotiate is Enabled

Link is good on FE port 2 connected to device: FPC2  
Speed is 100 MB  
Duplex is full  
Autonegotiate is Enabled

Link is good on FE port 8 connected to device: SPMB  
Speed is 100 MB  
Duplex is full  
Autonegotiate is Enabled

Link is good on FE port 10 connected to device: SCC  
Speed is 100 MB



```
Duplex is full
Autonegotiate is Enabled
```

### show chassis ethernet-switch errors

```
user@host> show chassis ethernet-switch errors
Accumulated error counts for port 0 connected to device FPC0:
  MLT3      2
  Lock      0
  Xmit      0
  ESD       0
  False carrier 2
  Disconnects 0
  FX mode   0
Accumulated error counts for port 1 connected to device FPC1:
  MLT3      2
  Lock      0
  Xmit      0
  ESD       0
  False carrier 2
  Disconnects 0
  FX mode   0
Accumulated error counts for port 2 connected to device FPC2:
  MLT3      2
  Lock      0
  Xmit      0
  ESD       0
  False carrier 3
  Disconnects 0
  FX mode   0
Accumulated error counts for port 3 connected to device FPC3:
  MLT3      0
  Lock      0
  Xmit      0
  ESD       0
  False carrier 0
  Disconnects 0
Accumulated error counts for port 4 connected to device Nothing:
  MLT3      0
  Lock      0
  Xmit      0
  ESD       0
  False carrier 0
  Disconnects 0
  FX mode   0
...
```

### show chassis ethernet-switch statistics

```
user@host> show chassis ethernet-switch statistics
Statistics for port 0 connected to device FPC0:
  TX Unicast packets      68113
  TX Multicast packets    0
  TX Broadcast packets    20851
  TX Late collisions      0
  TX Excessive collisions 0
  TX Dropped packets      0

  RX Unicast packets      67410
  RX Multicast packets    0
  RX Broadcast packets    20852
```

RX FCS Errors	0
RX Alignment Errors	0
RX Dropped Packets	0
RX Fragments	0
RX Symbol Errors	0

## Statistics for port 1 connected to device FPC1:

TX Unicast packets	66496
TX Multicast packets	0
TX Broadcast packets	20080
TX Late collisions	0
TX Excessive collisions	0
TX Dropped packets	0

RX Unicast packets	66037
RX Multicast packets	0
RX Broadcast packets	20080
RX FCS Errors	0
RX Alignment Errors	0
RX Dropped Packets	0
RX Fragments	0
RX Symbol Errors	0

## Statistics for port 2 connected to device FPC2:

TX Unicast packets	64206
TX Multicast packets	0
TX Broadcast packets	21183
TX Late collisions	0
TX Excessive collisions	0
TX Dropped packets	0

RX Unicast packets	63671
RX Multicast packets	0
RX Broadcast packets	21183
RX FCS Errors	0
RX Alignment Errors	0
RX Dropped Packets	0
RX Fragments	0
RX Symbol Errors	0

## Statistics for port 3 connected to device FPC3:

...

**show chassis ethernet-switch errors (TX Matrix Plus Router)**

user@host&gt; show chassis ethernet-switch errors

sfc0-re0:

-----  
Displaying error for switch 0

Displaying error for switch 1

Accumulated error counts for port 0 connected to device LCC0:

MLT3	0
Lock	0
Xmit	0
ESD	0
False carrier	0
Disconnects	0
FX mode	0

lcc0-re0:  
-----

```

Displaying error for switch 0
Accumulated error counts for port 6 connected to device FPC0:
  MLT3      0
  Lock      0
  Xmit      0
  ESD       0
  False carrier 5
  Disconnects 0
  FX mode   0
Accumulated error counts for port 7 connected to device FPC1:
  MLT3      0
  Lock      0
  Xmit      0
  ESD       0
  False carrier 7
  Disconnects 0
  FX mode   0
Accumulated error counts for port 19 connected to device Other RE:
  MLT3      0
  Lock      0
  Xmit      0
  ESD       0
  False carrier 0
  Disconnects 0
  FX mode   0
Accumulated error counts for port 20 connected to device SFC0:
  MLT3      0
  Lock      0
  Xmit      0
  ESD       0
  False carrier 0
  Disconnects 0
  FX mode   0

```

#### show chassis ethernet-switch sfc errors (TX Matrix Plus Router)

```

user@host> show chassis ethernet-switch errors switch sfc
sfc0-re0:
-----
Displaying error for switch 1
Accumulated error counts for port 0 connected to device LCC0:
  MLT3      0
  Lock      0
  Xmit      0
  ESD       0
  False carrier 0
  Disconnects 0
  FX mode   0
Accumulated error counts for port 2 connected to device LCC1:
  MLT3      0
  Lock      0
  Xmit      0
  ESD       0
  False carrier 0
  Disconnects 0
  FX mode   0
Accumulated error counts for port 4 connected to device LCC2:
  MLT3      0
  Lock      0
  Xmit      0
  ESD       0

```

```
False carrier 0
Disconnects   0
FX mode       0
Accumulated error counts for port 6 connected to device LCC3:
MLT3         0
Lock         0
Xmit         0
ESD          0
False carrier 0
Disconnects   0
FX mode       0
```

```
lcc0-re0:
```

```
-----
error: command is not valid on the t1600
```

```
lcc1-re0:
```

```
-----
error: command is not valid on the t1600
```

```
lcc2-re0:
```

```
-----
error: command is not valid on the t1600
```

```
lcc3-re0:
```

```
-----
error: command is not valid on the t1600
```

#### show chassis ethernet-switch statistics (TX Matrix Plus Router)

```
user@host> show chassis ethernet-switch statistics
```

```
sfc0-re0:
```

```
-----
Displaying port statistics for switch 0
Statistics for port 1 connected to device 1GSW:
```

```
TX Packets 64 Octets      5183577
TX Packets 65-127 Octets  67820
TX Packets 128-255 Octets 772
TX Packets 256-511 Octets 136
TX Packets 512-1023 Octets 68
TX Packets 1024-1518 Octets 10881
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                5263254
TX Multicast Packets      16
TX Broadcast Packets      723403
TX PAUSEMAC Ctrl Frames   0
TX Oversize Packets       0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           349922253
TX Packet OK Counter      5263254
TX Pause Packet Counter   0
TX Unicast Counter        4539835
RX Packets 64 Octets      6513629
RX Packets 65-127 Octets  88761
RX Packets 128-255 Octets 6382
RX Packets 256-511 Octets 22027
RX Packets 512-1023 Octets 4319
```

```

RX Packets 1024-1518 Octets  49922
RX Packets 1519-2047 Octets  0
RX Packets 2048-4095 Octets  0
RX Packets 4096-9216 Octets  0
RX Packets 9217-16383 Octets  0
RX Octets 6685040
RX Multicast Packets 4
RX Broadcast Packets 2137376
RX FCS Errors 0
RX Fragments 0
RX MAC Control Packets 0
RX Out of Range Length 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 509224602
RX Unicast Frame Count 4547660
RX Packet OK Count 6685040
Statistics for port 9 connected to device RE1:
TX Packets 64 Octets 2500318
TX Packets 65-127 Octets 443
TX Packets 128-255 Octets 0
TX Packets 256-511 Octets 0
TX Packets 512-1023 Octets 0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets 2500761
TX Multicast Packets 4
TX Broadcast Packets 2500757
TX PAUSEMAC Ctrl Frames 0
TX Oversize Packets 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 160049670
TX Packet OK Counter 0
TX Pause Packet Counter 0
TX Unicast Counter 0
RX Packets 64 Octets 701191
RX Packets 65-127 Octets 5882
RX Packets 128-255 Octets 2
RX Packets 256-511 Octets 0
RX Packets 512-1023 Octets 17965
RX Packets 1024-1518 Octets 7
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets 725047
RX Multicast Packets 8
RX Broadcast Packets 2500757
RX FCS Errors 0
RX Fragments 0
RX MAC Control Packets 0
RX Out of Range Length 0
RX Undersize Packets 0
RX Oversize Packets 0

```

```
RX Jabbers                                0
RX Control Frame Counter                  0
RX Pause Frame Counter                    0
RX Byte Counter                          62402656
RX Unicast Frame Count                    0
RX Packet OK Count                        0
Statistics for port 17 connected to device RE0:
TX Packets 64 Octets                      7214818
TX Packets 65-127 Octets                  94640
TX Packets 128-255 Octets                 6384
TX Packets 256-511 Octets                 22027
TX Packets 512-1023 Octets                22284
TX Packets 1024-1518 Octets              49929
TX Packets 1519-2047 Octets               0
TX Packets 2048-4095 Octets               0
TX Packets 4096-9216 Octets               0
TX Packets 9217-16383 Octets              0
TX Octets                                7410082
TX Multicast Packets                      12
TX Broadcast Packets                     2497247
TX PAUSEMAC Ctrl Frames                  0
TX Oversize Packets                      0
TX FCS Error Counter                     0
TX Fragment Counter                      0
TX Byte Counter                          571626932
TX Packet OK Counter                     0
TX Pause Packet Counter                   0
TX Unicast Counter                       0
RX Packets 64 Octets                      4823701
RX Packets 65-127 Octets                  67812
RX Packets 128-255 Octets                 772
RX Packets 256-511 Octets                 136
RX Packets 512-1023 Octets                68
RX Packets 1024-1518 Octets              10881
RX Packets 1519-2047 Octets               0
RX Packets 2048-4095 Octets               0
RX Packets 4096-9216 Octets               0
RX Packets 9217-16383 Octets              0
RX Octets                                4903370
RX Multicast Packets                      8
RX Broadcast Packets                     2497247
RX FCS Errors                            0
RX Fragments                             0
RX MAC Control Packets                   0
RX Out of Range Length                   0
RX Undersize Packets                     0
RX Oversize Packets                      0
RX Jabbers                               0
RX Control Frame Counter                  0
RX Pause Frame Counter                    0
RX Byte Counter                          326889517
RX Unicast Frame Count                    0
RX Packet OK Count                        0
```

```
Displaying port statistics for switch 1
Statistics for port 0 connected to device LCC0:
TX Packets 64 Octets                      5053443
TX Packets 65-127 Octets                  59737
TX Packets 128-255 Octets                 768
TX Packets 256-511 Octets                 87
TX Packets 512-1023 Octets                68
```

```

TX Packets 1024-1518 Octets 85
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets 5114188
TX Multicast Packets 16
TX Broadcast Packets 1125742
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions 0
TX Excessive Collisions 0
TX Collision frames 0
TX PAUSEMAC Ctrl Frames 0
TX MAC ctrl frames 0
TX Frame deferred Xms 0
TX Frame excessive deferl 0
TX Oversize Packets 0
TX Jabbers 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 329291449
RX Packets 64 Octets 5640175
RX Packets 65-127 Octets 79875
RX Packets 128-255 Octets 6338
RX Packets 256-511 Octets 165
RX Packets 512-1023 Octets 4317
RX Packets 1024-1518 Octets 10
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets 5730880
RX Multicast Packets 4
RX Broadcast Packets 1735007
RX FCS Errors 0
RX Align Errors 0
RX Fragments 0
RX Symbol errors 0
RX Unsupported opcodes 0
RX Out of Range Length 0
RX False Carrier Errors 0
RX Undersize Packets 0
RX Oversize Packets 0
RX Jabbers 0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter 0
RX Control Frame Counter 0
RX Pause Frame Counter 0
RX Byte Counter 371282850
Statistics for port 18 connected to device SPMB:
TX Packets 64 Octets 2990326
TX Packets 65-127 Octets 8572
TX Packets 128-255 Octets 4
TX Packets 256-511 Octets 49
TX Packets 512-1023 Octets 0
TX Packets 1024-1518 Octets 10793
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX 1519-1522 Good Vlan frms 0
TX Octets 3009744

```

```
TX Multicast Packets      20
TX Broadcast Packets     2458322
TX Single Collision frames 0
TX Mult. Collision frames 0
TX Late Collisions        0
TX Excessive Collisions   0
TX Collision frames       0
TX PAUSEMAC Ctrl Frames   0
TX MAC ctrl frames        0
TX Frame deferred Xmsns   0
TX Frame excessive deferl 0
TX Oversize Packets       0
TX Jabbers                0
TX FCS Error Counter      0
TX Fragment Counter       0
TX Byte Counter           203712524
RX Packets 64 Octets      873454
RX Packets 65-127 Octets  8886
RX Packets 128-255 Octets 44
RX Packets 256-511 Octets 21862
RX Packets 512-1023 Octets 2
RX Packets 1024-1518 Octets 49912
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Octets                 954160
RX Multicast Packets      0
RX Broadcast Packets      402369
RX FCS Errors             0
RX Align Errors           0
RX Fragments              0
RX Symbol errors          0
RX Unsupported opcodes    0
RX Out of Range Length    0
RX False Carrier Errors   0
RX Undersize Packets      0
RX Oversize Packets       0
RX Jabbers                0
RX 1519-1522 Good Vlan frms 0
RX MTU Exceed Counter     0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter           137941752
...
```

#### show chassis ethernet-switch (T4000 Router)

```
user@host> show chassis ethernet-switch
Displaying summary for switch 0
Link is good on GE port 6 connected to device: FPC0
  Speed is 100Mb
  Duplex is full
  Autonegotiate is Enabled
  False carrier sense count = 04

Link is good on GE port 9 connected to device: FPC3
  Speed is 100Mb
  Duplex is full
  Autonegotiate is Enabled
  False carrier sense count = 03
```



```

Link is good on GE port 11 connected to device: FPC5
Speed is 100Mb
Duplex is full
Autonegotiate is Enabled
False carrier sense count = 03

Link is good on GE port 12 connected to device: FPC6
Speed is 100Mb
Duplex is full
Autonegotiate is Enabled
False carrier sense count = 03

Link is good on GE port 14 connected to device: SPMB
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled

Link is good on GE port 18 connected to device: RE
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled

Link is good on GE port 19 connected to device: Other RE
Speed is 1000Mb
Duplex is full
Autonegotiate is Enabled

```

#### show chassis ethernet-switch errors (T4000 Router)

```

user@host> show chassis ethernet-switch errors

Displaying error for switch 0
Accumulated error counts for port 6 connected to device FPC0:
MLT3          0
Lock          0
Xmit          0
ESD           0
False carrier 4
Disconnects   0
FX mode       0
Accumulated error counts for port 9 connected to device FPC3:
MLT3          0
Lock          0
Xmit          0
ESD           0
False carrier 3
Disconnects   0
FX mode       0
Accumulated error counts for port 11 connected to device FPC5:
MLT3          0
Lock          0
Xmit          0
ESD           0
False carrier 3
Disconnects   0
FX mode       0
Accumulated error counts for port 12 connected to device FPC6:
MLT3          0
Lock          0
Xmit          0
ESD           0

```

```
False carrier  3
Disconnects    0
FX mode        0
Accumulated error counts for port 19 connected to device Other RE:
MLT3           0
Lock           0
Xmit           0
ESD            0
False carrier  0
Disconnects    0
FX mode        0
```

#### show chassis ethernet-switch (PTX5000 Packet Transport Router)

```
user@host> show chassis ethernet-switch
Displaying summary for switch 0
Link is good on XE port 2 connected to device: SPMB
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on XE port 11 connected to device: FPC7
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on XE port 12 connected to device: FPC6
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on XE port 13 connected to device: FPC5
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on XE port 15 connected to device: FPC3
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on XE port 16 connected to device: FPC2
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled

Link is good on XE port 18 connected to device: FPC0
Speed is 1000Mb
Duplex is full
```

```

Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled

```

```

Link is good on XE port 19 connected to device: OTHER RE
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled

```

```

Link is good on XE port 20 connected to device: RE
Speed is 1000Mb
Duplex is full
Autonegotiate is Disabled
Flow Control TX is Disabled
Flow Control RX is Disabled

```

#### show chassis ethernet-switch statistics (PTX5000 Packet Transport Router)

```

user@host> show chassis ethernet-switch statistics
Displaying port statistics for switch 0
Statistics for port 2 connected to device SPMB:
TX Packets 64 Octets      10942
TX Packets 65-127 Octets  843
TX Packets 128-255 Octets 2
TX Packets 256-511 Octets 2
TX Packets 512-1023 Octets 0
TX Packets 1024-1518 Octets 6862
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets      18651
TX Multicast Packets 6
TX Broadcast Packets 10331
TX PAUSEMAC Ctrl Frames 0
TX Oversize Packets 0
TX FCS Error Counter 0
TX Fragment Counter 0
TX Byte Counter 8105166
TX Packet OK Counter 0
TX Pause Packet Counter 0
TX Unicast Counter 0
RX Packets 64 Octets      8679
RX Packets 65-127 Octets  2364
RX Packets 128-255 Octets 531
RX Packets 256-511 Octets 112
RX Packets 512-1023 Octets 26
RX Packets 1024-1518 Octets 8
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets      11720
RX Multicast Packets 0
RX Broadcast Packets 10331
RX FCS Errors 0
RX Fragments 0
RX MAC Control Packets 0
RX Out of Range Length 0

```

```
RX Undersize Packets      0
RX Oversize Packets      0
RX Jabbers                0
RX Control Frame Counter  0
RX Pause Frame Counter    0
RX Byte Counter          938105
RX Unicast Frame Count    0
RX Packet OK Count        0
Statistics for port 11 connected to device FPC7:
TX Packets 64 Octets      14492
TX Packets 65-127 Octets  3542
TX Packets 128-255 Octets 6
TX Packets 256-511 Octets 45
TX Packets 512-1023 Octets 60
```

Continued...

```
Statistics for port 18 connected to device FPC0:
TX Packets 64 Octets      15212
TX Packets 65-127 Octets  3810
TX Packets 128-255 Octets 6
TX Packets 256-511 Octets 43
TX Packets 512-1023 Octets 66
TX Packets 1024-1518 Octets 169
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                  19306
TX Multicast Packets        0
TX Broadcast Packets        10886
TX PAUSEMAC Ctrl Frames    0
TX Oversize Packets         0
TX FCS Error Counter        0
TX Fragment Counter         0
TX Byte Counter             1569412
TX Packet OK Counter        0
TX Pause Packet Counter     0
TX Unicast Counter          0
RX Packets 64 Octets        17994
RX Packets 65-127 Octets    8006
RX Packets 128-255 Octets   230
RX Packets 256-511 Octets   19
RX Packets 512-1023 Octets  53
RX Packets 1024-1518 Octets 11
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                   26313
RX Multicast Packets        0
RX Broadcast Packets        10886
RX FCS Errors               0
RX Fragments                0
RX MAC Control Packets      0
RX Out of Range Length      0
RX Undersize Packets        0
RX Oversize Packets         0
RX Jabbers                  0
RX Control Frame Counter    2
RX Pause Frame Counter      2
```

```

RX Byte Counter          1836287
RX Unicast Frame Count   0
RX Packet OK Count       0
Statistics for port 19 connected to device OTHER RE:
TX Packets 64 Octets     10234
TX Packets 65-127 Octets 162
TX Packets 128-255 Octets 0
TX Packets 256-511 Octets 0
TX Packets 512-1023 Octets 0
TX Packets 1024-1518 Octets 0
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0
TX Octets                10396
TX Multicast Packets     8
TX Broadcast Packets     10317
TX PAUSEMAC Ctrl Frames  0
TX Oversize Packets      0
TX FCS Error Counter     0
TX Fragment Counter      0
TX Byte Counter          666260
TX Packet OK Counter     0
TX Pause Packet Counter  0
TX Unicast Counter       0
RX Packets 64 Octets     4073
RX Packets 65-127 Octets 325
RX Packets 128-255 Octets 1
RX Packets 256-511 Octets 0
RX Packets 512-1023 Octets 0
RX Packets 1024-1518 Octets 72
RX Packets 1519-2047 Octets 0
RX Packets 2048-4095 Octets 0
RX Packets 4096-9216 Octets 0
RX Packets 9217-16383 Octets 0
RX Octets                4471
RX Multicast Packets     0
RX Broadcast Packets     10317
RX FCS Errors            0
RX Fragments            0
RX MAC Control Packets   0
RX Out of Range Length   0
RX Undersize Packets     0
RX Oversize Packets      0
RX Jabbers               0
RX Control Frame Counter 0
RX Pause Frame Counter   0
RX Byte Counter          387333
RX Unicast Frame Count   0
RX Packet OK Count       0
Statistics for port 20 connected to device RE:
TX Packets 64 Octets     658856
TX Packets 65-127 Octets 45535
TX Packets 128-255 Octets 1900
TX Packets 256-511 Octets 532
TX Packets 512-1023 Octets 372
TX Packets 1024-1518 Octets 191
TX Packets 1519-2047 Octets 0
TX Packets 2048-4095 Octets 0
TX Packets 4096-9216 Octets 0
TX Packets 9217-16383 Octets 0

```

TX Octets	707386
TX Multicast Packets	0
TX Broadcast Packets	10421
TX PAUSEMAC Ctrl Frames	0
TX Oversize Packets	0
TX FCS Error Counter	0
TX Fragment Counter	0
TX Byte Counter	46608676
TX Packet OK Counter	0
TX Pause Packet Counter	0
TX Unicast Counter	0
RX Packets 64 Octets	27394
RX Packets 65-127 Octets	20271
RX Packets 128-255 Octets	78
RX Packets 256-511 Octets	215
RX Packets 512-1023 Octets	269
RX Packets 1024-1518 Octets	253370
RX Packets 1519-2047 Octets	0
RX Packets 2048-4095 Octets	0
RX Packets 4096-9216 Octets	0
RX Packets 9217-16383 Octets	0
RX Octets	301597
RX Multicast Packets	8
RX Broadcast Packets	10421
RX FCS Errors	0
RX Fragments	0
RX MAC Control Packets	0
RX Out of Range Length	0
RX Undersize Packets	0
RX Oversize Packets	0
RX Jabbers	0
RX Control Frame Counter	0
RX Pause Frame Counter	0
RX Byte Counter	275043436
RX Unicast Frame Count	0
RX Packet OK Count	0

Continued ...

#### show chassis ethernet-switch port-state (PTX5000 Packet Transport Router)

```
user@host> show chassis ethernet-switch port-state
Displaying port state for switch 0
Port      : 02
Target    : SPMB

Error reading port 2 connected to device: SPMB
```

## show chassis fabric fpcs

<b>List of Syntax</b>	<a href="#">Syntax on page 1063</a> <a href="#">Syntax (MX Series Routers) on page 1063</a> <a href="#">Syntax (MX2010 and MX2020 3D Universal Edge Routers) on page 1063</a> <a href="#">Syntax (T4000 Core Router) on page 1063</a> <a href="#">Syntax (PTX Series Packet Transport Routers) on page 1063</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1063</a>
<b>Syntax</b>	show chassis fabric fpcs <fcc number>
<b>Syntax (MX Series Routers)</b>	show chassis fabric fpcs <all-members> <local> <member member-id>
<b>Syntax (MX2010 and MX2020 3D Universal Edge Routers)</b>	show chassis fabric fpcs
<b>Syntax (T4000 Core Router)</b>	show chassis fabric fpcs
<b>Syntax (PTX Series Packet Transport Routers)</b>	show chassis fabric fpcs <slot fpc-slot>
<b>Syntax (TX Matrix Plus Router)</b>	show chassis fabric fpcs <fcc number>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
<b>Description</b>	(M320, MX Series, and T Series routers, EX8200 switches, and PTX Series Packet Transport Routers only) Display the state of the electrical switch fabric links between the Flexible PIC Concentrators (FPCs) and the Switch Interface Boards (SIBs).
<b>Options</b>	<p><b>none</b>—Display the switch fabric link state. On a TX Matrix router, display the switching fabric link states for the FPCs in all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display the switching fabric link states for the FPCs in all routers connected to the TX Matrix Plus router.</p> <p><b>all-members</b>—(MX Series routers only) (Optional) Display the switching fabric link states for the FPCs in all members of the Virtual Chassis configuration.</p> <p><b>fcc number</b>—(TX Matrix router and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the switch fabric link state for the FPCs in the specified T640 router</p>

(line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the switch fabric link state for the FPCs in the specified router (line-card chassis) that is connected to the TX Matrix Plus router. Replace *number* with a following value depending on the LCC configurations:

- From **0** through **3** on a T640 router on the routing matrix with TX Matrix routers.
- From **0** through **3** on a T1600 router on the routing matrix with TX Matrix Plus routers.
- From **0** through **7** on a T1600 router in a routing matrix with TX Matrix Plus router with 3D SIBs.
- **0, 2, 4, 6** on a T4000 router in a routing matrix with TX Matrix Plus router with 3D SIBs.

**local**—(MX Series routers only) (Optional) Display the switching fabric link states for the FPCs in the local Virtual Chassis member.

**member member-id**—(MX Series routers only) (Optional) Display the switching fabric link states for the FPCs in the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**slot fpc-slot**—(PTX Series Packet Transport Routers only) (Optional) Display the fabric state of the specified FPC slot. If no value is provided, display the status of all FPCs.

**Required Privilege Level** view

**Related Documentation**

- *request chassis fabric fpc*
- [show chassis fpc on page 1160](#)
- *Displaying Information About DPCs or FPCs in an MX Series Router*

**List of Sample Output**

- [show chassis fabric fpcs \(M320 Router\) on page 1066](#)
- [show chassis fabric fpcs \(MX240 Router\) on page 1067](#)
- [show chassis fabric fpcs \(MX480 Router\) on page 1067](#)
- [show chassis fabric fpcs \(MX960 Router\) on page 1068](#)
- [show chassis fabric fpcs \(MX240 with AS MLC Modular Carrier Card\) on page 1070](#)
- [show chassis fabric fpcs \(MX480 with AS MLC Modular Carrier Card\) on page 1070](#)
- [show chassis fabric fpcs \(MX480 Router with MPC4E\) on page 1071](#)
- [show chassis fabric fpcs \(MX960 with AS MLC Modular Carrier Card on page 1072](#)
- [show chassis fabric fpcs \(MX2010 Router\) on page 1074](#)
- [show chassis fabric fpcs \(MX2020 Router\) on page 1077](#)
- [show chassis fabric fpcs \(MX2020 Router with MPC4E\) on page 1080](#)
- [show chassis fabric fpcs \(T320 Router\) on page 1081](#)
- [show chassis fabric fpcs \(T640 Router\) on page 1082](#)
- [show chassis fabric fpcs \(TX Matrix Router\) on page 1082](#)
- [show chassis fabric fpcs \(TX Matrix Router with 3D SIBs\) on page 1084](#)
- [show chassis fabric fpcs lcc \(TX Matrix Router with 3D SIBs\) on page 1087](#)
- [show chassis fabric fpcs \(T1600 Router\) on page 1087](#)



[show chassis fabric fpcs \(T4000 Core Router\) on page 1089](#)  
[show chassis fabric fpcs \(TX Matrix Plus Router\) on page 1090](#)  
[show chassis fabric fpcs lcc \(TX Matrix Plus Router\) on page 1098](#)  
[show chassis fabric fpcs \(EX8200 Switch\) on page 1098](#)  
[show chassis fabric fpcs \(PTX3000 Router\) on page 1099](#)

**Output Fields** [Table 111 on page 1066](#) lists the output fields for the **show chassis fabric fpcs** command. Output fields are listed in the approximate order in which they appear.

Table 111: show chassis fabric fpcs Output Fields

Field Name	Field Description
<b>Fabric management FPC state</b>	<p>Switching fabric link (link from SIB to FPC) state for each FPC:</p> <ul style="list-style-type: none"> <li>• <b>Unused</b>—FPC is not present. (On MX240 and MX480 routers with AS- MLC modular carrier card or MPC4E only) the fabric plane from the pair that share physical links (1 and 5, and 3 and 7) is inactive.</li> <li>• <b>Destination error on PFEs <i>list of PFE numbers</i></b>—Destination errors to the listed Packet Forwarding Engines. Indicates that the link is not carrying traffic to the listed Packet Forwarding Engines.  <b>NOTE:</b> In Junos OS Release 9.6 and later, the list of Packet Forwarding Engines with destination errors is displayed in the output.  In Junos OS Releases before 9.6, the output only indicates that there are destination errors. However, the list of Packet Forwarding Engines with destination errors is not displayed.</li> <li>• <b>Links ok</b>—Link between the spare SIB and FPC is eligible to carry traffic.</li> <li>• <b>Link error</b>—Link between the SIB and FPC has CRC errors. However, the link is still eligible to carry traffic.</li> <li>• <b>Plane disabled</b>—Fabric plane has been disabled for the following reasons: <ul style="list-style-type: none"> <li>• Destination errors have exceeded the thresholds.</li> <li>• Run-time link errors have exceeded the thresholds.</li> <li>• Initialization time link errors detected, and link training was unsuccessful.</li> <li>• <b>Plane Disabled, Links Error</b> (PTX Series Packet Transport Routers only)—The plane is disabled because of link errors detected at the FPC RX.</li> </ul> </li> <li>• <b>Plane Disabled, Links Down</b> (PTX Series Packet Transport Routers only)—The plane is disabled because of link errors detected at the SIB RX.</li> <li>• <b>Plane enabled</b>—Link between the active SIB and FPC is eligible to carry traffic.  <b>NOTE:</b> On the Enhanced MX SCB with MPC, a maximum of 4 planes are operational and running. On all the other SCBs with MPC, all the planes are operational and running.</li> <li>• <b>Plane Enabled, Links OK</b> (PTX Series Packet Transport Routers only)—The FPC CCL RX link is eligible to carry traffic.</li> <li>• <b>Plane Enabled, Links OK</b> (TX Matrix and TX Matrix Plus routers only)—The FPC HSL RX link is eligible to carry traffic.</li> </ul>

## Sample Output

### show chassis fabric fpcs (M320 Router)

```

user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC #2

```

```

PFE #1
  SIB #0      Plane enabled
  SIB #1      Plane enabled
  SIB #2      Plane enabled
  SIB #3      Plane enabled

```

#### show chassis fabric fpcs (MX240 Router)

```

user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC 2
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok
  PFE #2
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok
  PFE #3
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok

```

#### show chassis fabric fpcs (MX480 Router)

```

user@host> show chassis fabric fpcs

FPC 0
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled

```

```
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
FPC 1
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
```

#### show chassis fabric fpcs (MX960 Router)

```
user@host> show chassis fabric fpcs
FPC 0
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
```

```
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
FPC 1
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
FPC 2
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
...
```

**show chassis fabric fpcs (MX240 with AS MLC Modular Carrier Card)**

In the following output, FPC 1 is the AS MLC modular carrier card (AS MCC).

```
user@host>show chassis fabric fpcs
FPC 1
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Unused
    Plane 6: Plane enabled
    Plane 7: Unused
FPC 2
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
```

**show chassis fabric fpcs (MX480 with AS MLC Modular Carrier Card)**

In the following output, FPC 5 is the AS MLC modular carrier card (AS MCC).

```
user@host>show chassis fabric fpcs
FPC 2
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 4
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok
  PFE #2
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
    Plane 6: Links ok
    Plane 7: Links ok
FPC 5
```

```

PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Unused
Plane 6: Plane enabled
Plane 7: Unused

```

### show chassis fabric fpcs (MX480 Router with MPC4E)

In the following output, **FPC4** is the MPC4E (MPC4E-3D-32XGE-SFPP) card.

```
user@host > show chassis fabric fpcs
```

```
Fabric management FPC state:
```

```

FPC 0
PFE #0
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
PFE #1
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
FPC 1
PFE #0
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
PFE #1
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
PFE #2
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok

```

```
Plane 6: Plane enabled
Plane 7: Links ok
PFE #3
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled

FPC 3
PFE #0
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok

FPC 4
PFE #0
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Unused
Plane 6: Plane enabled
Plane 7: Unused

PFE #1
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Unused
Plane 6: Plane enabled
Plane 7: Unused
```

### show chassis fabric fpcs (MX960 with AS MLC Modular Carrier Card)

In the following output, FPC 5 is the AS MLC modular carrier card (AS MCC).

```
user@host>show chassis fabric fpcs
Fabric management FPC state:
FPC 0
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
```



```
Plane 5: Links ok
FPC 1
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
FPC 4
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
  PFE #2
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
  PFE #3
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
FPC 5
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
FPC 8
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Links ok
    Plane 5: Links ok
```

```
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok

PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
```

### show chassis fabric fpcs (MX2010 Router)

```
user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC 0
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 1
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 2
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
```

```
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 3
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #2
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #3
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 4
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 5
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
```

```

        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #1
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane disabled
        Plane 4: Plane enabled
    Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    FPC 6
    PFE #0
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane disabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #1
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane disabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #2
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane disabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #3
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane disabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    FPC 7
    PFE #0
        Plane 0: Plane enabled
        Plane 1: Plane enabled
        Plane 2: Plane enabled
        Plane 3: Plane disabled
        Plane 4: Plane enabled
        Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
    PFE #1
        Plane 0: Plane enabled
        Plane 1: Plane enabled

```

```

        Plane 2: Plane enabled
        Plane 3: Plane disabled
        Plane 4: Plane enabled
    Plane 5: Plane enabled
        Plane 6: Plane enabled
        Plane 7: Plane enabled
FPC 8
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 9
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane disabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled

```

#### show chassis fabric fpcs (MX2020 Router)

```

user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC 0
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #2
    Plane 0: Plane enabled

```

```

Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 1
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 2
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled

```

```
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 3
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
```

```
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 4
...
```

#### show chassis fabric fpcs (MX2020 Router with MPC4E)

```
user@host > show chassis fabric fpcs
Fabric management FPC state:
FPC 0
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 9
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 10
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
FPC 14
```



```

PFE #0
  Plane 0: Plane enabled
  Plane 1: Plane enabled
  Plane 2: Plane enabled
  Plane 3: Plane enabled
  Plane 4: Plane enabled
  Plane 5: Plane enabled
  Plane 6: Plane enabled
  Plane 7: Plane enabled
PFE #1
  Plane 0: Plane enabled
  Plane 1: Plane enabled
  Plane 2: Plane enabled
  Plane 3: Plane enabled
  Plane 4: Plane enabled
  Plane 5: Plane enabled
  Plane 6: Plane enabled
  Plane 7: Plane enabled
FPC 19
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #2
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled
  PFE #3
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    Plane 3: Plane enabled
    Plane 4: Plane enabled
    Plane 5: Plane enabled
    Plane 6: Plane enabled
    Plane 7: Plane enabled

```

### show chassis fabric fpcs (T320 Router)

```

user@host> show chassis fabric fpcs
FPC #3
  PFE #1

```

```

SIB #0
    Links ok
SIB #1
    Plane enabled
SIB #2
    Plane enabled
FPC #5
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Plane enabled
        SIB #2
            Plane enabled
FPC #7
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Plane enabled
        SIB #2
            Plane enabled

```

#### show chassis fabric fpcs (T640 Router)

```

user@host> show chassis fabric fpcs
Fabric management FPC state:

```

```

FPC #2
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Plane enabled
        SIB #2
            Plane enabled
        SIB #3
            Plane enabled
        SIB #4
            Plane enabled
FPC #3
    PFE #1
        SIB #2
            Plane enabled
        SIB #3
            Link error
            Destination error on PFES
            8   9   10   11   12   13   14   15   16   17   18   19   20   21
        SIB #4
            Destination error on PFES
            8   9   10   11   12   13   14   15   16   17   18   19   20   21
...

```

#### show chassis fabric fpcs (TX Matrix Router)

```

user@host> show chassis fabric fpcs
1cc0-re0:
-----
Fabric management FPC state:
FPC #0
    PFE #1

```

```

SIB #0
    Links ok
SIB #2
    Links ok
SIB #3
    Links ok
SIB #4
    Links ok
FPC #2
    PFE #1
        SIB #0
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
FPC #3
    PFE #1
        SIB #2
            Plane enabled
        SIB #3
            Link error
            Destination error on PFes
            0  1  2  3  4  5  6  7
            8  9 10 11 12 13 14 15 16 17 18 19 20 21
        SIB #4
            Destination error on PFes
            0  1  2  3  4  5  6  7
            8  9 10 11 12 13 14 15 16 17 18 19 20 21
...
FPC #4
    PFE #0
        SIB #4 Links ok
    PFE #1
        SIB #4 Links ok
FPC #5
    PFE #1
        SIB #4 Links ok
FPC #6
    PFE #1
        SIB #4 Links ok

lcc2-re0:
-----
Fabric management FPC state:
FPC #0
    PFE #1
        SIB #4 Links ok
FPC #1
    PFE #1
        SIB #4 Links ok
FPC #2
    PFE #0
        SIB #4 Links ok
    PFE #1
        SIB #4 Links ok
FPC #4
    PFE #0
        SIB #4 Links ok
    PFE #1
        SIB #4 Links ok
FPC #5

```

```
PFE #1
  SIB #4 Links ok
```

### show chassis fabric fpcs (TX Matrix Router with 3D SIBs)

```
user@host> show chassis fabric fpcs
1cc0-re0:
```

```
-----
Fabric management FPC state:
```

```
FPC #0
```

```
  PFE #0
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

```
  PFE #1
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

```
FPC #3
```

```
  PFE #0
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

```
  PFE #1
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

```
FPC #4
```

```
  PFE #0
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
PFE #1
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #5
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
PFE #1
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #6
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
PFE #1
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
```

```

        SIB #4
            Links ok

1cc2-re0:
-----

1cc4-re0:
-----
Fabric management FPC state:
FPC #2
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
FPC #3
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok

1cc6-re0:
-----

```

**show chassis fabric fpcs lcc (TX Matrix Router with 3D SIBs)**

```
user@host> show chassis fabric fpcs lcc 4
lcc4-re0:
```

```
-----
Fabric management FPC state:
```

```
FPC #2
```

```
  PFE #0
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

```
  PFE #1
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

```
FPC #3
```

```
  PFE #0
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

```
  PFE #1
```

```
    SIB #0
```

```
      Links ok
```

```
    SIB #1
```

```
      Links ok
```

```
    SIB #2
```

```
      Links ok
```

```
    SIB #3
```

```
      Links ok
```

```
    SIB #4
```

```
      Links ok
```

**show chassis fabric fpcs (T1600 Router)**

```
user@host> show chassis fabric fpcs
```

```
Fabric management FPC state:
```

```
FPC #0
```

```
  PFE #0
```

```
    SIB #0
```

```
      Links ok
```

```
SIB #1
    Plane enabled
SIB #2
    Plane enabled
SIB #3
    Plane enabled
SIB #4
    Plane enabled
PFE #1
    SIB #0
        Links ok
    SIB #1
        Plane enabled
    SIB #2
        Plane enabled
    SIB #3
        Plane enabled
    SIB #4
        Plane enabled
FPC #1
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Plane enabled
        SIB #2
            Plane enabled
        SIB #3
            Plane enabled
        SIB #4
            Plane enabled
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Plane enabled
        SIB #2
            Plane enabled
        SIB #3
            Plane enabled
        SIB #4
            Plane enabled
FPC #2
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Plane enabled
        SIB #2
            Plane enabled
        SIB #3
            Plane enabled
        SIB #4
            Plane enabled
FPC #4
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Plane enabled
        SIB #2
```



```

        Plane enabled
    SIB #3
        Plane enabled
    SIB #4
        Plane enabled
PFE #1
    SIB #0
        Links ok
    SIB #1
        Plane enabled
    SIB #2
        Plane enabled
    SIB #3
        Plane enabled
    SIB #4
        Plane enabled
FPC #3
PFE #1
    SIB #2
        Plane enabled
    SIB #3
        Link error
        Destination error on PFEs
            0  1  2  3  4  5  6  7
            8  9 10 11 12 13 14 15 16 17 18 19 20 21
    SIB #4
        Destination error on PFEs
            0  1  2  3  4  5  6  7
            8  9 10 11 12 13 14 15 16 17 18 19 20 21

```

#### show chassis fabric fpcs (T4000 Core Router)

```

Fabric management FPC state:
FPC #2
PFE #0
    SIB #0
        Links ok
    SIB #1
        Plane enabled
    SIB #2
        Plane enabled
    SIB #3
        Plane enabled
    SIB #4
        Plane enabled
FPC #3
PFE #0
    SIB #0
        Links ok
    SIB #1
        Plane enabled
    SIB #2
        Plane enabled
    SIB #3
        Plane enabled
    SIB #4
        Plane enabled
FPC #5
PFE #0
    SIB #0
        Links ok
    SIB #1
        Plane enabled

```

```
SIB #2
    Plane enabled
SIB #3
    Plane enabled
SIB #4
    Plane enabled
PFE #1
    SIB #0
        Links ok
    SIB #1
        Plane enabled
    SIB #2
        Plane enabled
    SIB #3
        Plane enabled
    SIB #4
        Plane enabled
FPC #6
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Plane enabled
        SIB #2
            Plane enabled
        SIB #3
            Plane enabled
        SIB #4
            Plane enabled
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Plane enabled
        SIB #2
            Plane enabled
        SIB #3
            Plane enabled
        SIB #4
            Plane enabled
```

#### show chassis fabric fpcs (TX Matrix Plus Router)

```
user@host> show chassis fabric fpcs
lcc0-re0:
```

```
-----
Fabric management FPC state:
```

```
FPC #0
    PFE #1
        SIB #0
            Unused
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
FPC #2
    PFE #0
```

```

SIB #0
    Unused
SIB #1
    Links ok
SIB #2
    Links ok
SIB #3
    Links ok
SIB #4
    Links ok
PFE #1
    SIB #0
        Unused
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
FPC #3
    PFE #1
        SIB #2
            Plane enabled
        SIB #3
            Link error
            Destination error on PFes
            0  1  2  3  4  5  6  7
            8  9 10 11 12 13 14 15 16 17 18 19 20 21
        SIB #4
            Destination error on PFes
            0  1  2  3  4  5  6  7
            8  9 10 11 12 13 14 15 16 17 18 19 20 21
FPC #4
    PFE #0
        SIB #0
            Unused
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
    PFE #1
        SIB #0
            Unused
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
FPC #6
    PFE #0
        SIB #0
            Unused
        SIB #1
            Links ok

```

```

SIB #2
    Links ok
SIB #3
    Links ok
SIB #4
    Links ok
PFE #1
    SIB #0
        Unused
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
FPC #7
    PFE #0
        SIB #0
            Unused
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok

```

lcc1-re0:

-----  
Fabric management FPC state:

```

FPC #2
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
FPC #4
    PFE #0
        SIB #0
            Links ok
        SIB #1

```

```

        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
PFE #1
    SIB #0
        Links ok
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Destination error on PFES      1      8      9     29     40     65     72     73
                                         93    104
    SIB #4
        Links ok
FPC #6
PFE #0
    SIB #0
        Links ok
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
PFE #1
    SIB #0
        Links ok
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
FPC #7
PFE #0
    SIB #0
        Links ok
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok

lcc2-re0:
-----
Fabric management FPC state:
FPC #0
    PFE #0

```

```

SIB #0
SIB #1
SIB #2
SIB #3
SIB #4
PFE #1
SIB #0
SIB #1
SIB #2
SIB #3
SIB #4
FPC #2
PFE #0
SIB #0
SIB #1
SIB #2
SIB #3
SIB #4
PFE #1
SIB #0
SIB #1
SIB #2
SIB #3
SIB #4
FPC #4
PFE #0
SIB #0
SIB #1
SIB #2
SIB #3
SIB #4
FPC #5
PFE #0
SIB #0
SIB #1

```

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

Links ok

```

        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
PFE #1
    SIB #0
        Links ok
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
FPC #6
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
FPC #7
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
```

lcc3-re0:

-----  
Fabric management FPC state:

```
FPC #0
    PFE #0
        SIB #0
            Links ok
```

```
SIB #1
    Links ok
SIB #2
    Links ok
SIB #3
    Links ok
SIB #4
    Links ok
PFE #1
    SIB #0
        Links ok
    SIB #1
        Links ok
    SIB #2
        Links ok
    SIB #3
        Links ok
    SIB #4
        Links ok
FPC #2
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
FPC #4
    PFE #0
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
        SIB #3
            Links ok
        SIB #4
            Links ok
    PFE #1
        SIB #0
            Links ok
        SIB #1
            Links ok
        SIB #2
            Links ok
```



```
SIB #3
Links ok
SIB #4
Links ok
FPC #5
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
PFE #1
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #6
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
PFE #1
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #7
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
```

```
SIB #4
    Links ok
```

### show chassis fabric fpcs lcc (TX Matrix Plus Router)

```
user@host> show chassis fabric fpcs lcc 0
lcc0-re1:
-----
Fabric management FPC state:
FPC #3
  PFE #1
    SIB #2
      Plane enabled
    SIB #3
      Link error
      Destination error on PFes
      8   9   10  11  12  13  14  15  16  17  18  19  20  21
    SIB #4
      Destination error on PFes
      8   9   10  11  12  13  14  15  16  17  18  19  20  21
FPC #4
  PFE #0
    SIB #0 Links ok
    SIB #1 Links ok
    SIB #2 Links ok
    SIB #3 Links ok
    SIB #4 Links ok
  PFE #1
    SIB #0 Links ok
    SIB #1 Links ok
    SIB #2 Links ok
    SIB #3 Links ok
    SIB #4 Links ok
FPC #6
  PFE #0
    SIB #0 Links ok
    SIB #1 Links ok
    SIB #2 Links ok
    SIB #3 Links ok
    SIB #4 Links ok
  PFE #1
    SIB #0 Links ok
    SIB #1 Links ok
    SIB #2 Links ok
    SIB #3 Links ok
    SIB #4 Links ok
FPC #7
  PFE #0
    SIB #0 Links ok
    SIB #1 Links ok
    SIB #2 Links ok
    SIB #3 Links ok
    SIB #4 Links ok
```

### show chassis fabric fpcs (EX8200 Switch)

```
user@host> show chassis fabric fpcs
Fabric management FPC state
FPC 6
  PFE #0
    Plane 0: Plane enabled
```

```

Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
Plane 8: Plane enabled
Plane 9: Plane enabled
Plane 10: Plane enabled
Plane 11: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
Plane 8: Plane enabled
Plane 9: Plane enabled
Plane 10: Plane enabled
Plane 11: Plane enabled
FPC 7
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
Plane 8: Plane enabled
Plane 9: Plane enabled
Plane 10: Plane enabled
Plane 11: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
Plane 6: Links ok
Plane 7: Links ok
Plane 8: Plane enabled
Plane 9: Plane enabled
Plane 10: Plane enabled
Plane 11: Plane enabled

```

#### show chassis fabric fpcs (PTX3000 Router)

```

user@host> show chassis fabric fpcs slot 8
Fabric management FPC state:
FPC #8
PFE #0
SIB0_Fcore0 (plane 0)  Plane Enabled, Links OK
SIB0_Fcore1 (plane 1)  Plane Enabled, Links OK
SIB1_Fcore0 (plane 2)  Plane Enabled, Links OK

```

SIB1_Fcore1	(plane 3)	Plane Enabled, Links OK
SIB2_Fcore0	(plane 4)	Plane Enabled, Links OK
SIB2_Fcore1	(plane 5)	Plane Enabled, Links OK
SIB3_Fcore0	(plane 6)	Plane Enabled, Links OK
SIB3_Fcore1	(plane 7)	Plane Enabled, Links OK
SIB4_Fcore0	(plane 8)	Plane Enabled, Links OK
SIB4_Fcore1	(plane 9)	Plane Enabled, Links OK
SIB5_Fcore0	(plane 10)	Plane Enabled, Links OK
SIB5_Fcore1	(plane 11)	Plane Enabled, Links OK
SIB6_Fcore0	(plane 12)	Plane Enabled, Links OK
SIB6_Fcore1	(plane 13)	Plane Enabled, Links OK
SIB7_Fcore0	(plane 14)	Plane Enabled, Links OK
SIB7_Fcore1	(plane 15)	Plane Enabled, Links OK
SIB8_Fcore0	(plane 16)	Plane Enabled, Links OK
SIB8_Fcore1	(plane 17)	Plane Enabled, Links OK
PFE #1		
SIB0_Fcore0	(plane 0)	Plane Enabled, Links OK
SIB0_Fcore1	(plane 1)	Plane Enabled, Links OK
SIB1_Fcore0	(plane 2)	Plane Enabled, Links OK
SIB1_Fcore1	(plane 3)	Plane Enabled, Links OK
SIB2_Fcore0	(plane 4)	Plane Enabled, Links OK
SIB2_Fcore1	(plane 5)	Plane Enabled, Links OK
SIB3_Fcore0	(plane 6)	Plane Enabled, Links OK
SIB3_Fcore1	(plane 7)	Plane Enabled, Links OK
SIB4_Fcore0	(plane 8)	Plane Enabled, Links OK
SIB4_Fcore1	(plane 9)	Plane Enabled, Links OK
SIB5_Fcore0	(plane 10)	Plane Enabled, Links OK
SIB5_Fcore1	(plane 11)	Plane Enabled, Links OK
SIB6_Fcore0	(plane 12)	Plane Enabled, Links OK
SIB6_Fcore1	(plane 13)	Plane Enabled, Links OK
SIB7_Fcore0	(plane 14)	Plane Enabled, Links OK
SIB7_Fcore1	(plane 15)	Plane Enabled, Links OK
SIB8_Fcore0	(plane 16)	Plane Enabled, Links OK
SIB8_Fcore1	(plane 17)	Plane Enabled, Links OK

## show chassis fabric map

<b>List of Syntax</b>	<a href="#">Syntax on page 1101</a> <a href="#">Syntax (MX Series Router) on page 1101</a>
<b>Syntax</b>	<pre>show chassis fabric map plane &lt;plane-number&gt;</pre>
<b>Syntax (MX Series Router)</b>	<pre>show chassis fabric map &lt;all-members&gt; &lt;local&gt; &lt;member member-id&gt; &lt;plane plane-number&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.4 for EX Series switches.</p>
<b>Description</b>	<p>(M120 and MX Series routers and EX8200 switches only) On the M120 router, display the state of the switching fabric map for connections from the Forwarding Engine Boards (FEBs) to the ports on the fabric planes, as interpreted by the fabric plane. On the MX Series router and the EX8200 switch, display the state of the switching fabric map for connections from each Packet Forwarding Engine on the Dense Port Concentrators (DPCs) to the ports on the fabric planes, as interpreted by the fabric plane. For information about the meaning of “fabric plane”, “DPCs”, and “SIBs” on the switches, see <a href="#">“EX Series Switches Hardware and CLI Terminology Mapping” on page 315</a>.</p>
<b>Options</b>	<p><b>none</b>—Display the switching fabric map state for the M120 or MX Series router or EX8200 switch.</p> <p><b>all-members</b>—(MX Series routers only) (Optional) Display the switching fabric map state for all the members of the Virtual Chassis configuration.</p> <p><b>local</b>—(MX Series routers only) (Optional) Display the switching fabric map state for the local Virtual Chassis member.</p> <p><b>member <i>member-id</i></b>—(MX Series routers only) (Optional) Display the switching fabric map state for the specified member of the Virtual Chassis configuration. Replace the <i>member-id</i> with a value of 0 or 1.</p> <p><b>plane <i>plane-number</i></b>—(Optional) Display the state of the fabric link for the specified plane number.</p> <ul style="list-style-type: none"> <li>For the M120 router, replace <i>plane-number</i> with a value from 0 through 3.</li> <li>For the MX480 and MX240 routers, replace <i>plane-number</i> with a value from 0 through 7.</li> <li>For the MX960 router, replace <i>plane-number</i> with a value from 0 through 5.</li> <li>For the EX8208 switch, replace <i>plane-number</i> with a value from 0 through 11.</li> <li>For the EX8216 switch, replace <i>plane-number</i> with a value from 0 through 7.</li> </ul>

**Required Privilege Level** view

**List of Sample Output** [show chassis fabric map \(M120 Router\) on page 1102](#)  
[show chassis fabric map \(MX Series Routers\) on page 1102](#)  
[show chassis fabric map plane 1 \(EX8200 Switch\) on page 1106](#)

**Output Fields** [Table 112 on page 1102](#) lists the output fields for the **show chassis fabric map** command. Output fields are listed in the approximate order in which they appear.

**Table 112: show chassis fabric map Output Fields**

Field Name	Field Description
<b>in-links</b>	Fabric map for receive side links.
<b>out-links</b>	Fabric map for transmit side links.
<b>state</b>	State of the fabric link: <ul style="list-style-type: none"> <li>• <b>RESET</b>—Link between SIB and FPC/DPC is powered down on purpose. This is done in all non-dual PFE based boards.</li> <li>• <b>UP</b>—Link between SIB and FPC/DPC is up and running.</li> <li>• <b>DOWN</b>—Link between SIB and FPC/DPC is powered down.</li> <li>• <b>FAULT</b>—SIB is in alarmed state where the SIB's plane is not operational for the following reasons: <ul style="list-style-type: none"> <li>• On-board F-chip is not operational.</li> <li>• Fiber optic connector faults.</li> <li>• FPC connector faults.</li> <li>• SIB midplane connector faults.</li> </ul> </li> </ul>

## Sample Output

### show chassis fabric map (M120 Router)

```
user@host> show chassis fabric map
FEB0->CB0F0_00 up CB0F0_08->FEB7 Down

FEB1->CB0F0_01 Down CB0F0_09->FEB6 Down

FEB6->CB0F0_02 Down CB0F0_10->FEB1 Down

FEB2->CB0F0_03 Down CB0F0_11->FEB0 up

FEB3->CB0F0_04 Down CB0F0_12->FEB3 Down

FEB4->CB0F0_05 up CB0F0_13->FEB2 Down

FEB7->CB0F0_06 Down CB0F0_14->FEB5 Down

FEB5->CB0F0_07 Down CB0F0_15->FEB4 up:
```

### show chassis fabric map (MX Series Routers)

```
user@host> show chassis fabric map
```

DPC4PFE0->CB0F0_00_0	up	CB0F0_00_0->DPC4PFE0	up
DPC4PFE1->CB0F0_00_1	up	CB0F0_00_1->DPC4PFE1	up
DPC4PFE2->CB0F0_00_2	up	CB0F0_00_2->DPC4PFE2	up
DPC4PFE3->CB0F0_00_3	up	CB0F0_00_3->DPC4PFE3	up
DPC7PFE0->CB0F0_01_0	Down	CB0F0_01_0->DPC7PFE0	Down
DPC7PFE1->CB0F0_01_1	Down	CB0F0_01_1->DPC7PFE1	Down
DPC7PFE2->CB0F0_01_2	Down	CB0F0_01_2->DPC7PFE2	Down
DPC7PFE3->CB0F0_01_3	Down	CB0F0_01_3->DPC7PFE3	Down
DPC3PFE0->CB0F0_03_0	Down	CB0F0_03_0->DPC3PFE0	Down
DPC3PFE1->CB0F0_03_1	Down	CB0F0_03_1->DPC3PFE1	Down
DPC3PFE2->CB0F0_03_2	Down	CB0F0_03_2->DPC3PFE2	Down
DPC3PFE3->CB0F0_03_3	Down	CB0F0_03_3->DPC3PFE3	Down
DPC8PFE0->CB0F0_05_0	Down	CB0F0_05_0->DPC8PFE0	Down
DPC8PFE1->CB0F0_05_1	Down	CB0F0_05_1->DPC8PFE1	Down
DPC8PFE2->CB0F0_05_2	Down	CB0F0_05_2->DPC8PFE2	Down
DPC8PFE3->CB0F0_05_3	Down	CB0F0_05_3->DPC8PFE3	Down
DPC1PFE0->CB0F0_06_0	Down	CB0F0_06_0->DPC1PFE0	Down
DPC1PFE1->CB0F0_06_1	Down	CB0F0_06_1->DPC1PFE1	Down
DPC1PFE2->CB0F0_06_2	Down	CB0F0_06_2->DPC1PFE2	Down
DPC1PFE3->CB0F0_06_3	Down	CB0F0_06_3->DPC1PFE3	Down
DPC10PFE0->CB0F0_07_0	Down	CB0F0_07_0->DPC10PFE0	Down
DPC10PFE1->CB0F0_07_1	Down	CB0F0_07_1->DPC10PFE1	Down
DPC10PFE2->CB0F0_07_2	Down	CB0F0_07_2->DPC10PFE2	Down
DPC10PFE3->CB0F0_07_3	Down	CB0F0_07_3->DPC10PFE3	Down
DPC11PFE0->CB0F0_08_0	Down	CB0F0_08_0->DPC11PFE0	Down
DPC11PFE1->CB0F0_08_1	Down	CB0F0_08_1->DPC11PFE1	Down
DPC11PFE2->CB0F0_08_2	Down	CB0F0_08_2->DPC11PFE2	Down
DPC11PFE3->CB0F0_08_3	Down	CB0F0_08_3->DPC11PFE3	Down
DPC0PFE0->CB0F0_09_0	Down	CB0F0_09_0->DPC0PFE0	Down
DPC0PFE1->CB0F0_09_1	Down	CB0F0_09_1->DPC0PFE1	Down
DPC0PFE2->CB0F0_09_2	Down	CB0F0_09_2->DPC0PFE2	Down
DPC0PFE3->CB0F0_09_3	Down	CB0F0_09_3->DPC0PFE3	Down
DPC9PFE0->CB0F0_11_0	Down	CB0F0_11_0->DPC9PFE0	Down
DPC9PFE1->CB0F0_11_1	Down	CB0F0_11_1->DPC9PFE1	Down
DPC9PFE2->CB0F0_11_2	Down	CB0F0_11_2->DPC9PFE2	Down
DPC9PFE3->CB0F0_11_3	Down	CB0F0_11_3->DPC9PFE3	Down
DPC2PFE0->CB0F0_13_0	up	CB0F0_13_0->DPC2PFE0	up
DPC2PFE1->CB0F0_13_1	up	CB0F0_13_1->DPC2PFE1	up
DPC2PFE2->CB0F0_13_2	up	CB0F0_13_2->DPC2PFE2	up
DPC2PFE3->CB0F0_13_3	up	CB0F0_13_3->DPC2PFE3	up
DPC6PFE0->CB0F0_14_0	Down	CB0F0_14_0->DPC6PFE0	Down
DPC6PFE1->CB0F0_14_1	Down	CB0F0_14_1->DPC6PFE1	Down
DPC6PFE2->CB0F0_14_2	Down	CB0F0_14_2->DPC6PFE2	Down
DPC6PFE3->CB0F0_14_3	Down	CB0F0_14_3->DPC6PFE3	Down
DPC5PFE0->CB0F0_15_0	Down	CB0F0_15_0->DPC5PFE0	Down
DPC5PFE1->CB0F0_15_1	Down	CB0F0_15_1->DPC5PFE1	Down
DPC5PFE2->CB0F0_15_2	Down	CB0F0_15_2->DPC5PFE2	Down
DPC5PFE3->CB0F0_15_3	Down	CB0F0_15_3->DPC5PFE3	Down
DPC4PFE0->CB0F1_00_0	up	CB0F1_00_0->DPC4PFE0	up
DPC4PFE1->CB0F1_00_1	up	CB0F1_00_1->DPC4PFE1	up
DPC4PFE2->CB0F1_00_2	up	CB0F1_00_2->DPC4PFE2	up
DPC4PFE3->CB0F1_00_3	up	CB0F1_00_3->DPC4PFE3	up
DPC7PFE0->CB0F1_01_0	Down	CB0F1_01_0->DPC7PFE0	Down
DPC7PFE1->CB0F1_01_1	Down	CB0F1_01_1->DPC7PFE1	Down
DPC7PFE2->CB0F1_01_2	Down	CB0F1_01_2->DPC7PFE2	Down
DPC7PFE3->CB0F1_01_3	Down	CB0F1_01_3->DPC7PFE3	Down
DPC3PFE0->CB0F1_03_0	Down	CB0F1_03_0->DPC3PFE0	Down
DPC3PFE1->CB0F1_03_1	Down	CB0F1_03_1->DPC3PFE1	Down
DPC3PFE2->CB0F1_03_2	Down	CB0F1_03_2->DPC3PFE2	Down
DPC3PFE3->CB0F1_03_3	Down	CB0F1_03_3->DPC3PFE3	Down
DPC8PFE0->CB0F1_05_0	Down	CB0F1_05_0->DPC8PFE0	Down

DPC8PFE1->CB0F1_05_1	Down	CB0F1_05_1->DPC8PFE1	Down
DPC8PFE2->CB0F1_05_2	Down	CB0F1_05_2->DPC8PFE2	Down
DPC8PFE3->CB0F1_05_3	Down	CB0F1_05_3->DPC8PFE3	Down
DPC1PFE0->CB0F1_06_0	Down	CB0F1_06_0->DPC1PFE0	Down
DPC1PFE1->CB0F1_06_1	Down	CB0F1_06_1->DPC1PFE1	Down
DPC1PFE2->CB0F1_06_2	Down	CB0F1_06_2->DPC1PFE2	Down
DPC1PFE3->CB0F1_06_3	Down	CB0F1_06_3->DPC1PFE3	Down
DPC10PFE0->CB0F1_07_0	Down	CB0F1_07_0->DPC10PFE0	Down
DPC10PFE1->CB0F1_07_1	Down	CB0F1_07_1->DPC10PFE1	Down
DPC10PFE2->CB0F1_07_2	Down	CB0F1_07_2->DPC10PFE2	Down
DPC10PFE3->CB0F1_07_3	Down	CB0F1_07_3->DPC10PFE3	Down
DPC11PFE0->CB0F1_08_0	Down	CB0F1_08_0->DPC11PFE0	Down
DPC11PFE1->CB0F1_08_1	Down	CB0F1_08_1->DPC11PFE1	Down
DPC11PFE2->CB0F1_08_2	Down	CB0F1_08_2->DPC11PFE2	Down
DPC11PFE3->CB0F1_08_3	Down	CB0F1_08_3->DPC11PFE3	Down
DPC0PFE0->CB0F1_09_0	Down	CB0F1_09_0->DPC0PFE0	Down
DPC0PFE1->CB0F1_09_1	Down	CB0F1_09_1->DPC0PFE1	Down
DPC0PFE2->CB0F1_09_2	Down	CB0F1_09_2->DPC0PFE2	Down
DPC0PFE3->CB0F1_09_3	Down	CB0F1_09_3->DPC0PFE3	Down
DPC9PFE0->CB0F1_11_0	Down	CB0F1_11_0->DPC9PFE0	Down
DPC9PFE1->CB0F1_11_1	Down	CB0F1_11_1->DPC9PFE1	Down
DPC9PFE2->CB0F1_11_2	Down	CB0F1_11_2->DPC9PFE2	Down
DPC9PFE3->CB0F1_11_3	Down	CB0F1_11_3->DPC9PFE3	Down
DPC2PFE0->CB0F1_13_0	up	CB0F1_13_0->DPC2PFE0	up
DPC2PFE1->CB0F1_13_1	up	CB0F1_13_1->DPC2PFE1	up
DPC2PFE2->CB0F1_13_2	up	CB0F1_13_2->DPC2PFE2	up
DPC2PFE3->CB0F1_13_3	up	CB0F1_13_3->DPC2PFE3	up
DPC6PFE0->CB0F1_14_0	Down	CB0F1_14_0->DPC6PFE0	Down
DPC6PFE1->CB0F1_14_1	Down	CB0F1_14_1->DPC6PFE1	Down
DPC6PFE2->CB0F1_14_2	Down	CB0F1_14_2->DPC6PFE2	Down
DPC6PFE3->CB0F1_14_3	Down	CB0F1_14_3->DPC6PFE3	Down
DPC5PFE0->CB0F1_15_0	Down	CB0F1_15_0->DPC5PFE0	Down
DPC5PFE1->CB0F1_15_1	Down	CB0F1_15_1->DPC5PFE1	Down
DPC5PFE2->CB0F1_15_2	Down	CB0F1_15_2->DPC5PFE2	Down
DPC5PFE3->CB0F1_15_3	Down	CB0F1_15_3->DPC5PFE3	Down
DPC4PFE0->CB1F0_00_0	up	CB1F0_00_0->DPC4PFE0	up
DPC4PFE1->CB1F0_00_1	up	CB1F0_00_1->DPC4PFE1	up
DPC4PFE2->CB1F0_00_2	up	CB1F0_00_2->DPC4PFE2	up
DPC4PFE3->CB1F0_00_3	up	CB1F0_00_3->DPC4PFE3	up
DPC7PFE0->CB1F0_01_0	Down	CB1F0_01_0->DPC7PFE0	Down
DPC7PFE1->CB1F0_01_1	Down	CB1F0_01_1->DPC7PFE1	Down
DPC7PFE2->CB1F0_01_2	Down	CB1F0_01_2->DPC7PFE2	Down
DPC7PFE3->CB1F0_01_3	Down	CB1F0_01_3->DPC7PFE3	Down
DPC3PFE0->CB1F0_03_0	Down	CB1F0_03_0->DPC3PFE0	Down
DPC3PFE1->CB1F0_03_1	Down	CB1F0_03_1->DPC3PFE1	Down
DPC3PFE2->CB1F0_03_2	Down	CB1F0_03_2->DPC3PFE2	Down
DPC3PFE3->CB1F0_03_3	Down	CB1F0_03_3->DPC3PFE3	Down
DPC8PFE0->CB1F0_05_0	Down	CB1F0_05_0->DPC8PFE0	Down
DPC8PFE1->CB1F0_05_1	Down	CB1F0_05_1->DPC8PFE1	Down
DPC8PFE2->CB1F0_05_2	Down	CB1F0_05_2->DPC8PFE2	Down
DPC8PFE3->CB1F0_05_3	Down	CB1F0_05_3->DPC8PFE3	Down
DPC1PFE0->CB1F0_06_0	Down	CB1F0_06_0->DPC1PFE0	Down
DPC1PFE1->CB1F0_06_1	Down	CB1F0_06_1->DPC1PFE1	Down
DPC1PFE2->CB1F0_06_2	Down	CB1F0_06_2->DPC1PFE2	Down
DPC1PFE3->CB1F0_06_3	Down	CB1F0_06_3->DPC1PFE3	Down
DPC10PFE0->CB1F0_07_0	Down	CB1F0_07_0->DPC10PFE0	Down
DPC10PFE1->CB1F0_07_1	Down	CB1F0_07_1->DPC10PFE1	Down
DPC10PFE2->CB1F0_07_2	Down	CB1F0_07_2->DPC10PFE2	Down
DPC10PFE3->CB1F0_07_3	Down	CB1F0_07_3->DPC10PFE3	Down
DPC11PFE0->CB1F0_08_0	Down	CB1F0_08_0->DPC11PFE0	Down
DPC11PFE1->CB1F0_08_1	Down	CB1F0_08_1->DPC11PFE1	Down



DPC11PFE2->CB1F0_08_2	Down	CB1F0_08_2->DPC11PFE2	Down
DPC11PFE3->CB1F0_08_3	Down	CB1F0_08_3->DPC11PFE3	Down
DPC0PFE0->CB1F0_09_0	Down	CB1F0_09_0->DPC0PFE0	Down
DPC0PFE1->CB1F0_09_1	Down	CB1F0_09_1->DPC0PFE1	Down
DPC0PFE2->CB1F0_09_2	Down	CB1F0_09_2->DPC0PFE2	Down
DPC0PFE3->CB1F0_09_3	Down	CB1F0_09_3->DPC0PFE3	Down
DPC9PFE0->CB1F0_11_0	Down	CB1F0_11_0->DPC9PFE0	Down
DPC9PFE1->CB1F0_11_1	Down	CB1F0_11_1->DPC9PFE1	Down
DPC9PFE2->CB1F0_11_2	Down	CB1F0_11_2->DPC9PFE2	Down
DPC9PFE3->CB1F0_11_3	Down	CB1F0_11_3->DPC9PFE3	Down
DPC2PFE0->CB1F0_13_0	up	CB1F0_13_0->DPC2PFE0	up
DPC2PFE1->CB1F0_13_1	up	CB1F0_13_1->DPC2PFE1	up
DPC2PFE2->CB1F0_13_2	up	CB1F0_13_2->DPC2PFE2	up
DPC2PFE3->CB1F0_13_3	up	CB1F0_13_3->DPC2PFE3	up
DPC6PFE0->CB1F0_14_0	Down	CB1F0_14_0->DPC6PFE0	Down
DPC6PFE1->CB1F0_14_1	Down	CB1F0_14_1->DPC6PFE1	Down
DPC6PFE2->CB1F0_14_2	Down	CB1F0_14_2->DPC6PFE2	Down
DPC6PFE3->CB1F0_14_3	Down	CB1F0_14_3->DPC6PFE3	Down
DPC5PFE0->CB1F0_15_0	Down	CB1F0_15_0->DPC5PFE0	Down
DPC5PFE1->CB1F0_15_1	Down	CB1F0_15_1->DPC5PFE1	Down
DPC5PFE2->CB1F0_15_2	Down	CB1F0_15_2->DPC5PFE2	Down
DPC5PFE3->CB1F0_15_3	Down	CB1F0_15_3->DPC5PFE3	Down
DPC4PFE0->CB1F1_00_0	up	CB1F1_00_0->DPC4PFE0	up
DPC4PFE1->CB1F1_00_1	up	CB1F1_00_1->DPC4PFE1	up
DPC4PFE2->CB1F1_00_2	up	CB1F1_00_2->DPC4PFE2	up
DPC4PFE3->CB1F1_00_3	up	CB1F1_00_3->DPC4PFE3	up
DPC7PFE0->CB1F1_01_0	Down	CB1F1_01_0->DPC7PFE0	Down
DPC7PFE1->CB1F1_01_1	Down	CB1F1_01_1->DPC7PFE1	Down
DPC7PFE2->CB1F1_01_2	Down	CB1F1_01_2->DPC7PFE2	Down
DPC7PFE3->CB1F1_01_3	Down	CB1F1_01_3->DPC7PFE3	Down
DPC3PFE0->CB1F1_03_0	Down	CB1F1_03_0->DPC3PFE0	Down
DPC3PFE1->CB1F1_03_1	Down	CB1F1_03_1->DPC3PFE1	Down
DPC3PFE2->CB1F1_03_2	Down	CB1F1_03_2->DPC3PFE2	Down
DPC3PFE3->CB1F1_03_3	Down	CB1F1_03_3->DPC3PFE3	Down
DPC8PFE0->CB1F1_05_0	Down	CB1F1_05_0->DPC8PFE0	Down
DPC8PFE1->CB1F1_05_1	Down	CB1F1_05_1->DPC8PFE1	Down
DPC8PFE2->CB1F1_05_2	Down	CB1F1_05_2->DPC8PFE2	Down
DPC8PFE3->CB1F1_05_3	Down	CB1F1_05_3->DPC8PFE3	Down
DPC1PFE0->CB1F1_06_0	Down	CB1F1_06_0->DPC1PFE0	Down
DPC1PFE1->CB1F1_06_1	Down	CB1F1_06_1->DPC1PFE1	Down
DPC1PFE2->CB1F1_06_2	Down	CB1F1_06_2->DPC1PFE2	Down
DPC1PFE3->CB1F1_06_3	Down	CB1F1_06_3->DPC1PFE3	Down
DPC10PFE0->CB1F1_07_0	Down	CB1F1_07_0->DPC10PFE0	Down
DPC10PFE1->CB1F1_07_1	Down	CB1F1_07_1->DPC10PFE1	Down
DPC10PFE2->CB1F1_07_2	Down	CB1F1_07_2->DPC10PFE2	Down
DPC10PFE3->CB1F1_07_3	Down	CB1F1_07_3->DPC10PFE3	Down
DPC11PFE0->CB1F1_08_0	Down	CB1F1_08_0->DPC11PFE0	Down
DPC11PFE1->CB1F1_08_1	Down	CB1F1_08_1->DPC11PFE1	Down
DPC11PFE2->CB1F1_08_2	Down	CB1F1_08_2->DPC11PFE2	Down
DPC11PFE3->CB1F1_08_3	Down	CB1F1_08_3->DPC11PFE3	Down
DPC0PFE0->CB1F1_09_0	Down	CB1F1_09_0->DPC0PFE0	Down
DPC0PFE1->CB1F1_09_1	Down	CB1F1_09_1->DPC0PFE1	Down
DPC0PFE2->CB1F1_09_2	Down	CB1F1_09_2->DPC0PFE2	Down
DPC0PFE3->CB1F1_09_3	Down	CB1F1_09_3->DPC0PFE3	Down
DPC9PFE0->CB1F1_11_0	Down	CB1F1_11_0->DPC9PFE0	Down
DPC9PFE1->CB1F1_11_1	Down	CB1F1_11_1->DPC9PFE1	Down
DPC9PFE2->CB1F1_11_2	Down	CB1F1_11_2->DPC9PFE2	Down
DPC9PFE3->CB1F1_11_3	Down	CB1F1_11_3->DPC9PFE3	Down
DPC2PFE0->CB1F1_13_0	up	CB1F1_13_0->DPC2PFE0	up
DPC2PFE1->CB1F1_13_1	up	CB1F1_13_1->DPC2PFE1	up
DPC2PFE2->CB1F1_13_2	up	CB1F1_13_2->DPC2PFE2	up

DPC2PFE3->CB1F1_13_3	up	CB1F1_13_3->DPC2PFE3	up
DPC6PFE0->CB1F1_14_0	Down	CB1F1_14_0->DPC6PFE0	Down
DPC6PFE1->CB1F1_14_1	Down	CB1F1_14_1->DPC6PFE1	Down
DPC6PFE2->CB1F1_14_2	Down	CB1F1_14_2->DPC6PFE2	Down
DPC6PFE3->CB1F1_14_3	Down	CB1F1_14_3->DPC6PFE3	Down
DPC5PFE0->CB1F1_15_0	Down	CB1F1_15_0->DPC5PFE0	Down
DPC5PFE1->CB1F1_15_1	Down	CB1F1_15_1->DPC5PFE1	Down
DPC5PFE2->CB1F1_15_2	Down	CB1F1_15_2->DPC5PFE2	Down
DPC5PFE3->CB1F1_15_3	Down	CB1F1_15_3->DPC5PFE3	Down
plane 4 is not up			
plane 5 is not up			

### show chassis fabric map plane 1 (EX8200 Switch)

```

user@host> show chassis fabric map plane 1
regress@tp-grande01> show chassis fabric map plane 1
DPC6PFE0->CB0F0_00_0    Down    CB0F0_00_0->DPC6PFE0    Down
DPC6PFE1->CB0F0_00_1    Down    CB0F0_00_1->DPC6PFE1    Down
DPC6PFE2->CB0F0_00_2    Down    CB0F0_00_2->DPC6PFE2    Down
DPC6PFE3->CB0F0_00_3    Down    CB0F0_00_3->DPC6PFE3    Down
DPC0PFE0->CB0F0_01_0    Down    CB0F0_01_0->DPC0PFE0    Down
DPC0PFE1->CB0F0_01_1    Down    CB0F0_01_1->DPC0PFE1    Down
DPC0PFE2->CB0F0_01_2    Down    CB0F0_01_2->DPC0PFE2    Down
DPC0PFE3->CB0F0_01_3    Down    CB0F0_01_3->DPC0PFE3    Down
DPC5PFE0->CB0F0_02_0    Down    CB0F0_02_0->DPC5PFE0    Down
DPC5PFE1->CB0F0_02_1    Down    CB0F0_02_1->DPC5PFE1    Down
DPC5PFE2->CB0F0_02_2    Down    CB0F0_02_2->DPC5PFE2    Down
DPC5PFE3->CB0F0_02_3    Down    CB0F0_02_3->DPC5PFE3    Down
DPC3PFE0->CB0F0_03_0    Down    CB0F0_03_0->DPC3PFE0    Down
DPC3PFE1->CB0F0_03_1    Down    CB0F0_03_1->DPC3PFE1    Down
DPC3PFE2->CB0F0_03_2    Down    CB0F0_03_2->DPC3PFE2    Down
DPC3PFE3->CB0F0_03_3    Down    CB0F0_03_3->DPC3PFE3    Down
DPC4PFE0->CB0F0_04_0    Down    CB0F0_04_0->DPC4PFE0    Down
DPC4PFE1->CB0F0_04_1    Down    CB0F0_04_1->DPC4PFE1    Down
DPC4PFE2->CB0F0_04_2    Down    CB0F0_04_2->DPC4PFE2    Down
DPC4PFE3->CB0F0_04_3    Down    CB0F0_04_3->DPC4PFE3    Down
DPC2PFE0->CB0F0_05_0    Down    CB0F0_05_0->DPC2PFE0    Down
DPC2PFE1->CB0F0_05_1    Down    CB0F0_05_1->DPC2PFE1    Down
DPC2PFE2->CB0F0_05_2    Down    CB0F0_05_2->DPC2PFE2    Down
DPC2PFE3->CB0F0_05_3    Down    CB0F0_05_3->DPC2PFE3    Down
DPC7PFE0->CB0F0_06_0    Down    CB0F0_06_0->DPC7PFE0    Down
DPC7PFE1->CB0F0_06_1    Down    CB0F0_06_1->DPC7PFE1    Down
DPC7PFE2->CB0F0_06_2    Down    CB0F0_06_2->DPC7PFE2    Down
DPC7PFE3->CB0F0_06_3    Down    CB0F0_06_3->DPC7PFE3    Down
DPC1PFE0->CB0F0_07_0    Down    CB0F0_07_0->DPC1PFE0    Down
DPC1PFE1->CB0F0_07_1    Down    CB0F0_07_1->DPC1PFE1    Down
DPC1PFE2->CB0F0_07_2    Down    CB0F0_07_2->DPC1PFE2    Down
DPC1PFE3->CB0F0_07_3    Down    CB0F0_07_3->DPC1PFE3    Down
DPC0PFE0->CB0F0_08_0    Down    CB0F0_08_0->DPC0PFE0    Down
DPC0PFE1->CB0F0_08_1    Down    CB0F0_08_1->DPC0PFE1    Down
DPC0PFE2->CB0F0_08_2    Down    CB0F0_08_2->DPC0PFE2    Down
DPC0PFE3->CB0F0_08_3    Down    CB0F0_08_3->DPC0PFE3    Down
DPC7PFE0->CB0F0_09_0    Down    CB0F0_09_0->DPC7PFE0    Down
DPC7PFE1->CB0F0_09_1    Down    CB0F0_09_1->DPC7PFE1    Down
DPC7PFE2->CB0F0_09_2    Down    CB0F0_09_2->DPC7PFE2    Down
DPC7PFE3->CB0F0_09_3    Down    CB0F0_09_3->DPC7PFE3    Down
DPC1PFE0->CB0F0_10_0    Down    CB0F0_10_0->DPC1PFE0    Down
DPC1PFE1->CB0F0_10_1    Down    CB0F0_10_1->DPC1PFE1    Down
DPC1PFE2->CB0F0_10_2    Down    CB0F0_10_2->DPC1PFE2    Down
DPC1PFE3->CB0F0_10_3    Down    CB0F0_10_3->DPC1PFE3    Down
DPC4PFE0->CB0F0_11_0    Down    CB0F0_11_0->DPC4PFE0    Down

```

DPC4PFE1->CB0F0_11_1	Down	CB0F0_11_1->DPC4PFE1	Down
DPC4PFE2->CB0F0_11_2	Down	CB0F0_11_2->DPC4PFE2	Down
DPC4PFE3->CB0F0_11_3	Down	CB0F0_11_3->DPC4PFE3	Down
DPC2PFE0->CB0F0_12_0	Down	CB0F0_12_0->DPC2PFE0	Down
DPC2PFE1->CB0F0_12_1	Down	CB0F0_12_1->DPC2PFE1	Down
DPC2PFE2->CB0F0_12_2	Down	CB0F0_12_2->DPC2PFE2	Down
DPC2PFE3->CB0F0_12_3	Down	CB0F0_12_3->DPC2PFE3	Down
DPC5PFE0->CB0F0_13_0	Down	CB0F0_13_0->DPC5PFE0	Down
DPC5PFE1->CB0F0_13_1	Down	CB0F0_13_1->DPC5PFE1	Down
DPC5PFE2->CB0F0_13_2	Down	CB0F0_13_2->DPC5PFE2	Down
DPC5PFE3->CB0F0_13_3	Down	CB0F0_13_3->DPC5PFE3	Down
DPC3PFE0->CB0F0_14_0	Down	CB0F0_14_0->DPC3PFE0	Down
DPC3PFE1->CB0F0_14_1	Down	CB0F0_14_1->DPC3PFE1	Down
DPC3PFE2->CB0F0_14_2	Down	CB0F0_14_2->DPC3PFE2	Down
DPC3PFE3->CB0F0_14_3	Down	CB0F0_14_3->DPC3PFE3	Down
DPC6PFE0->CB0F0_15_0	Down	CB0F0_15_0->DPC6PFE0	Down
DPC6PFE1->CB0F0_15_1	Down	CB0F0_15_1->DPC6PFE1	Down
DPC6PFE2->CB0F0_15_2	Down	CB0F0_15_2->DPC6PFE2	Down
DPC6PFE3->CB0F0_15_3	Down	CB0F0_15_3->DPC6PFE3	Down

## show chassis fabric plane

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1108</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1108</a> <a href="#">Syntax (MX Series Routers) on page 1108</a> <a href="#">Syntax (MX2010 and MX2020 3D Universal Edge Routers) on page 1108</a>
<b>Syntax</b>	show chassis fabric plane
<b>Syntax (TX Matrix Plus Router)</b>	show chassis fabric plane <detail   extensive   terse> <lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Routers)</b>	show chassis fabric plane <detail   extensive   terse> <all-members> <local> <member <i>member-id</i> >
<b>Syntax (MX2010 and MX2020 3D Universal Edge Routers)</b>	show chassis fabric plane
<b>Release Information</b>	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.4 for EX Series switches. <b>detail</b> , <b>extensive</b> , <b>lcc</b> , <b>sfc</b> , and <b>terse</b> options introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
<b>Description</b>	(TX Matrix Plus router, T4000, T1600, M120, and MX Series routers and EX8200 switches only) On the M120 router, display the state of all fabric plane connections to the Forwarding Engine Boards (FEBs). On MX Series routers, display the state of all fabric plane connections to the Dense Port Concentrators (DPCs) and Packet Forwarding Engines (PFEs) on the Flexible PIC Concentrators (FPCs). On the TX Matrix Plus router, and on T1600 or T4000 routers in a routing matrix, display the state of the fabric management plane and the logical planes on the switch-fabric chassis (SFC) and line-card chassis (LCC). On EX8200 switches, display the state of all fabric planes. This command can be used on the master Routing Engine only.
<b>Options</b>	<b>none</b> —(MX2010 and MX2020 Routers only) (Optional) Display the state of the fabric management plane.  <b>detail</b> —(TX Matrix Plus routers, T1600 or T4000 routers in a routing matrix, and MX Series routers only) (Optional) Display detailed output for the fabric management plane. Show Switch Interface Board (SIB) states for the TXP-F13 SIB and the TXP-F2S SIB.  <b>extensive</b> —(TX Matrix Plus routers, T1600 or T4000 routers in a routing matrix, and MX Series routers only) (Optional) Display extensive output for the fabric management plane.

**terse**—(TX Matrix Plus routers and MX Series routers only) (Optional) Display terse output for the fabric management plane.

**all-members**—(MX Series routers only) (Optional) Display the state of all fabric plane connections on all members of the Virtual Chassis configuration.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display the state of all fabric plane connections on the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display the state of all fabric plane connections on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**sfc *number***—(TX Matrix Plus router only) (Optional) Show information about the TX Matrix Plus router (SFC). Replace *number* with 0.

**Required Privilege Level** view

**Related Documentation**

- [request chassis fabric plane on page 880](#)
- [show chassis fabric plane-location on page 1150](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

- [show chassis fabric plane \(M120 Router\) on page 1116](#)
- [show chassis fabric plane \(MX240 Router\) on page 1117](#)
- [show chassis fabric plane \(MX480 Router\) on page 1118](#)
- [show chassis fabric plane \(MX960 Router\) on page 1119](#)
- [show chassis fabric plane \(MX240 with AS MLC Modular Carrier Card\) on page 1120](#)
- [show chassis fabric plane \(MX480 with AS MLC Modular Carrier Card\) on page 1121](#)
- [show chassis fabric plane \(MX480 Router with MPC4E\) on page 1122](#)
- [show chassis fabric plane \(MX960 with AS-MLC Modular Carrier Card\) on page 1124](#)
- [show chassis fabric plane \(MX2010 Router\) on page 1126](#)
- [show chassis fabric plane \(MX2020 Router\) on page 1130](#)
- [show chassis fabric plane \(MX2020 Router with MPC4E\) on page 1135](#)

[show chassis fabric plane \(TX Matrix Plus Router\) on page 1138](#)  
[show chassis fabric plane \(TX Matrix Plus Router with 3D SIBs\) on page 1138](#)  
[show chassis fabric plane detail \(TX Matrix Plus Router\) on page 1139](#)  
[show chassis fabric plane extensive \(TX Matrix Plus Router \) on page 1140](#)  
[show chassis fabric plane extensive \(TX Matrix Plus Router with 3D SIBs\) on page 1142](#)  
[show chassis fabric plane terse \(TX Matrix Plus Router\) on page 1144](#)  
[show chassis fabric plane terse \(TX Matrix Plus Router with 3D SIBs\) on page 1144](#)  
[show chassis fabric plane lcc \(TX Matrix Plus Router\) on page 1145](#)  
[show chassis fabric plane lcc \(TX Matrix Plus Router with 3D SIBs\) on page 1145](#)  
[show chassis fabric plane sfc \(TX Matrix Plus Router\) on page 1146](#)  
[show chassis fabric plane sfc \(TX Matrix Plus Router with 3D SIBs\) on page 1146](#)  
[show chassis fabric plane \(T1600 Router\) on page 1146](#)  
[show chassis fabric plane extensive \(T1600 Router\) on page 1146](#)  
[show chassis fabric plane detail \(T1600 Router\) on page 1149](#)  
[show chassis fabric plane \(EX8200 Switch\) on page 1149](#)

**Output Fields** Table 113 on page 1110 lists the output fields for the **show chassis fabric plane** command. Output fields are listed in the approximate order in which they appear.

**Table 113: show chassis fabric plane Output Fields**

Field Name	Field Description	Level of output
<b>Plane</b>	(TX Matrix Plus, MX Series routers, M120 routers, and EX8200 switches only) Number of the plane.	none
<b>Plane state</b>	<p>(MX Series and M120 routers and EX8200 switches only) State of each plane:</p> <ul style="list-style-type: none"> <li>• <b>ACTIVE</b>—SIB is operational and running.</li> </ul> <p><b>NOTE:</b> On the Enhanced MX SCB with MPCs, a maximum of 4 planes are operational and running. On all the other SCBs with MPCs, all the planes are operational and running.</p> <ul style="list-style-type: none"> <li>• <b>FAULTY</b>— SIB is in alarmed state where the SIB's plane is not operational for the following reasons: <ul style="list-style-type: none"> <li>• On-board fabric ASIC is not operational.</li> <li>• Fiber optic connector faults.</li> <li>• FPC connector faults.</li> <li>• SIB midplane connector faults.</li> </ul> </li> </ul> <p>(MX2010 and MX2020 Routers only) State of each plane:</p> <ul style="list-style-type: none"> <li>• <b>ACTIVE</b>—SFB is operational and running.</li> <li>• <b>OFFLINE</b>— SFB is in offline.</li> </ul>	none
<b>FEB</b>	<p>(M120 routers only) FEB number and state of links to each FEB:</p> <ul style="list-style-type: none"> <li>• <b>Link error</b>—Link between SIB and FPC is not operational.</li> <li>• <b>Links ok</b>—Link between SIB and FPC is active.</li> <li>• <b>Unused</b>—No FPC is present.</li> </ul>	none

Table 113: show chassis fabric plane Output Fields (*continued*)

Field Name	Field Description	Level of output
FPC	(MX Series routers only) Slot number of each Dense Port Concentrator (DPC) or Flexible PIC Concentrator (FPC). An FPC occupies two DPC slots on an MX Series router. The interface corresponds to the lowest numbered DPC slot for which the FPC is installed.	none
PFE	<p>(MX Series and M120 routers only) Slot number of each Packet Forwarding Engine and the state of the links to the DCP: <b>Links ok</b>, <b>Link error</b>, or <b>Unused</b>. Each DPC includes four Packet Forwarding Engines.</p> <ul style="list-style-type: none"> <li>• <b>Links ok</b>: Link between SIB and FPC is active.</li> <li>• <b>Link error</b>: Link between SIB and FPC is not operational.</li> <li>• <b>Unused</b>: No FPC is present.</li> </ul> <p>(On MX240 and MX480 routers with AS MLC modular carrier card and MPC4E only) Indicates that the link between the fabric plane and the hardware link on the modular carrier card or MPC4E is not operational.</p> <p>(MX2010 and MX2020 routers only) Slot number of each Packet Forwarding Engine and the state of the links to the DPC: <b>Links ok</b>, <b>Link error</b>, or <b>Unused</b>. Each DPC includes four Packet Forwarding Engines.</p> <ul style="list-style-type: none"> <li>• <b>Links ok</b>: Link between SFB and FPC is active.</li> <li>• <b>Link error</b>: Link between SFB and FPC is not operational.</li> <li>• <b>Unused</b>: No FPC is present.</li> </ul>	none

Table 113: show chassis fabric plane Output Fields (*continued*)

Field Name	Field Description	Level of output
<b>State</b>	<p>(TX Matrix Plus, and T1600 or T4000 routers in a routing matrix only)—State of the fabric plane:</p> <ul style="list-style-type: none"> <li>• <b>Online</b>: Fabric plane is operational and running and links on the SIB are operational.</li> <li>• <b>Offline</b>: Fabric plane state is <b>Offline</b> because the plane does not have four or more F2S and one F13 online.</li> <li>• <b>Empty</b>: Fabric plane state is <b>Empty</b> if all SIBs in the plane are absent.</li> <li>• <b>Spare</b>: Fabric plane is redundant and can be operational if the operational fabric plane encounters an error.</li> <li>• <b>Check</b>: Fabric plane is in alarmed state due to the following reason and the cause of the error must be resolved: <ul style="list-style-type: none"> <li>• One or more SIBs (belonging to the fabric plane) in the <b>Online</b> or <b>Spare</b> states has transitioned to the <b>Check</b> state. <b>Check</b> state of the SIB can be caused by link errors or destination errors.</li> </ul> </li> <li>• <b>Fault</b>: Fabric plane is in alarmed state if one or more SIBs belonging to the plane are in the <b>Fault</b> state. A SIB can be in the <b>Fault</b> state because of the following reasons: <ul style="list-style-type: none"> <li>• On-board fabric ASIC is not operational.</li> <li>• Fiber optic connector faults.</li> <li>• FPC connector faults.</li> <li>• SIB midplane connector faults.</li> <li>• Link errors have exceeded the threshold.</li> </ul> </li> </ul>	none
<b>Link Errors</b>	(TX Matrix Plus routers with 3D SIBs only) indicate the number of links which are marked faulty because the errors on them have crossed threshold.	none
<b>Cable Errors</b>	(TX Matrix Plus routers with 3D SIBs only) Indicate the number of mandatory cables that are not connected, or in up state for that plane	none
<b>Destination Errors</b>	(TX Matrix Plus routers with 3D SIBs only) Indicates the number of destinations that are not reachable on this plane.	none
<b>Uptime</b>	(TX Matrix Plus, and T1600 or T4000 routers in a routing matrix only)—Time the fabric plane has been up and running.	none

Fabric Management Plane State Output Fields for the show chassis fabric plane extensive Command on a TX Matrix Plus Router



Table 113: show chassis fabric plane Output Fields (*continued*)

Field Name	Field Description	Level of output
<b>PLANE number</b>	<p>State of the fabric plane:</p> <ul style="list-style-type: none"> <li>• <b>Online</b>: Fabric plane is operational and running and links on the SIB are operational.</li> <li>• <b>Offline</b>: Fabric plane state is <b>Offline</b> because the plane does not have 4 or more F2S and 1 F13 online.</li> <li>• <b>Empty</b>: Fabric plane state is <b>Empty</b> if all SIBs in the plane are absent.</li> <li>• <b>Spare</b>: Fabric plane is redundant and can be operational if the operational fabric plane encounters an error.</li> <li>• <b>Check</b>: Fabric plane is in alarmed state due to the following reasons and the cause of the error must be resolved: <ul style="list-style-type: none"> <li>• One or more SIBs (belonging to the fabric plane) in the <b>Online</b> or <b>Spare</b> states has transitioned to the <b>Check</b> state. <b>Check</b> state of the SIB can be caused because of link errors or destination errors.</li> </ul> </li> <li>• <b>Fault</b>: Fabric plane is in alarmed state if one or more SIBs belonging to the plane are in the <b>Fault</b> state. A SIB can be in the <b>Fault</b> state because of the following reasons: <ul style="list-style-type: none"> <li>• On-board fabric ASIC is not operational.</li> <li>• Fiber optic connector faults.</li> <li>• FPC connector faults.</li> <li>• SIB midplane connector faults.</li> <li>• Link errors have exceeded the threshold.</li> </ul> </li> </ul>	<b>extensive</b>
<b>SIB F13/F2S slot-number</b>	<p>State of the TXP-F13 SIB or TXP-F2S SIB:</p> <ul style="list-style-type: none"> <li>• <b>Activating</b>—Transitional state when the SIB is transitioning to the <b>Online</b> or <b>Spare</b> state.</li> <li>• <b>Deactivating</b>—Transitional state when the SIB is going offline.</li> <li>• <b>Online</b>—SIB is operational and running.</li> <li>• <b>Offline</b>—SIB is powered down.</li> <li>• <b>Spare</b>—SIB is redundant and will move to active state if one of the working SIBs fails to pass traffic.</li> <li>• <b>Empty</b>—No SIB is present.</li> <li>• <b>Fault</b>—SIB is in alarmed state because of the following reasons and the cause of the error must be resolved: <ul style="list-style-type: none"> <li>• On-board fabric ASIC is not operational.</li> <li>• Fiber optic connector faults.</li> <li>• FPC connector faults.</li> <li>• SIB midplane connector faults.</li> <li>• Link errors have exceeded the threshold</li> </ul> </li> <li>• <b>Check</b>—SIB is in alarmed state where the SIB is partially operational because of link or destination errors. Only a SIB that is <b>Online</b> or <b>Spare</b> can transition to the <b>Check</b> state.</li> </ul> <p><b>NOTE:</b> If a SIB is not inserted properly, the SIB cannot transition to the <b>Online</b> or <b>Spare</b> state, and therefore cannot transition to the <b>Check</b> state.</p>	<b>extensive</b>

Table 113: show chassis fabric plane Output Fields (*continued*)

Field Name	Field Description	Level of output
<b>SIB F13 slot-number Odd/Even</b>	State of the TXP-F13 SIB even and odd port connection optical links from the TX Matrix Plus router (SFC) to the router (LCC) in the routing matrix. The left four ports on the SFC are labeled <b>Even</b> and provide connections to one even-numbered LCC—LCC0 or LCC2. The right four ports on the SFC are labeled <b>Odd</b> and provide connections to one odd-numbered LCC—LCC1 or LCC3.	<b>extensive</b>
<b>LCC number, SIB slot-number</b>	<p>State of the SIB on the LCC that is connected to the <b>Even</b> or <b>Odd</b> port on the TXP-F13 SIB faceplate:</p> <ul style="list-style-type: none"> <li>• <b>Links ok</b>—Links between the TXP-F13 SIB on the SFC and the LCC are active.</li> <li>• <b>Links error</b>—One or more links between the TXP-F13 SIB on the SFC and the LCC, have experienced an error, but the affected links remain operational.</li> <li>• <b>Unused</b>—No SIB is present.</li> </ul>	<b>extensive</b>
<b>SG number Port number</b>	<p>State of the SG chip ports on the LCC:</p> <ul style="list-style-type: none"> <li>• <b>Links ok</b>—Link is active.</li> <li>• <b>Link error</b>—Link is operational with errors.</li> <li>• <b>Link error crc saturated</b>—CRC has exceeded the rate threshold and reached saturation without optical issues—that is, a cable has not been cut, removed, or otherwise experienced an error.</li> <li>• <b>Link error crc saturated with optical errors</b>—CRC has exceeded the rate threshold and reached saturation with optical issues—that is, a cable has been cut, removed, or otherwise experienced an error.</li> <li>• <b>Unused</b>—Port is not in use.</li> </ul>	<b>extensive</b>
<b>SIB F2S slot-number</b>	State of the intra-chassis links between the TXP-F2S and TXP-F13 SIBs.	<b>extensive</b>

Fabric Management SIB State Output Fields for the show chassis fabric plane extensive Command on a TX Matrix Plus Router

Table 113: show chassis fabric plane Output Fields (*continued*)

Field Name	Field Description	Level of output
<i>SIB slot-number</i>	<p>State of the SIBs on the T1600/T4000 router (LCC) in the routing matrix:</p> <ul style="list-style-type: none"> <li>• <b>Activating</b>—Transitional state when the SIB is coming online.</li> <li>• <b>Deactivating</b>—Transitional state when the SIB is going offline.</li> <li>• <b>Connected</b>—SIBs on an LCC are connected and trained, but are either not online or are spare, because the plane on the the TX Matrix Plus router (SFC) is still offline. The LCC SIB transitions to the <b>Connected</b> state when the F13 SIB to which it connects is online but the SFC plane (to which the LCC SIB connects) is offline for some reason; for instance, when there are insufficient number of F2 SIBs in the plane.</li> <li>• <b>Disconnected</b>—If an F13 SIB on the TX Matrix Plus router (SFC) goes offline, then the SIBs on the LCCs connected to the F13 SIB get disconnected. On the TX Matrix Plus router with 3D SIBs, the LCC SIB is also disconnected if the F13 SIB is online, but none of the cables are connected or trained. The <b>Disconnected</b> state is valid only for SIBs on an LCC. An LCC SIB transitions to the <b>Disconnected</b> state when the F13 SIB to which it connects goes <b>Offline</b>, irrespective of the state of the SFC plane. <b>SFC Error</b>—If an F13 SIB on the TX Matrix Plus router (SFC) transitions to the <b>Fault</b> state (because of link errors, for instance), and if an LCC SIB connected to the F13 SIB comes online, the LCC SIB transitions to the <b>SFC Error</b> state. This state indicates that the F13 SIB to which the LCC SIB is connected has errors. <b>NOTE:</b> The <b>Connected</b>, <b>Disconnected</b>, and <b>SFC Error</b> states are applicable only to the SIBs on an LCC.</li> <li>• <b>Online</b>—SIB is operational and running.</li> <li>• <b>Offline</b>—SIB is powered down.</li> <li>• <b>Spare</b>—SIB is redundant and will move to active state if one of the working SIBs fails to pass traffic.</li> <li>• <b>Empty</b>—No SIB is present.</li> <li>• <b>Fault</b>—SIB is in alarmed state where the SIB's plane is not operational for the following reasons: <ul style="list-style-type: none"> <li>• On-board fabric ASIC is not operational.</li> <li>• Fiber optic connector faults.</li> <li>• FPC connector faults.</li> <li>• SIB midplane connector faults.</li> <li>• Link errors have exceeded the threshold</li> </ul> </li> <li>• <b>Check</b>—SIB is in alarmed state where the SIB is partially operational because of link or destination errors. Only a SIB that is <b>Online</b> or <b>Spare</b> can transition to the <b>Check</b> state. <b>NOTE:</b> If a SIB is not inserted properly, the SIB cannot transition to the <b>Online</b> or <b>Spare</b> state, and therefore cannot transition to the <b>Check</b> state.</li> </ul>	<b>extensive</b>

Table 113: show chassis fabric plane Output Fields (*continued*)

Field Name	Field Description	Level of output
<b>LCC SIB Link State</b>	State of the LCC SIB link: <ul style="list-style-type: none"> <li>• <b>Links ok</b>—Link is active.</li> <li>• <b>Links error</b>—A link error has occurred, but the link remains operational.</li> <li>• <b>Unused</b>—SIB is not in use.</li> </ul>	<b>extensive</b>
<b>SG number Port number</b>	State of the SG chip ports on the LCC: <ul style="list-style-type: none"> <li>• <b>Links ok</b>—Link is active.</li> <li>• <b>Link error</b>—Link is operational with errors.</li> <li>• <b>Link error crc saturated</b>—CRC has exceeded the rate threshold and reached saturation without optical issues—that is, a cable has not been cut, removed, or otherwise experienced an error.</li> <li>• <b>Link error crc saturated with optical errors</b>—CRC has exceeded the rate threshold and reached saturation with optical issues—that is, a cable has been cut, removed, or otherwise experienced an error.</li> <li>• <b>Unused</b>—Port is not in use.</li> </ul>	<b>extensive</b>

## Sample Output

### show chassis fabric plane (M120 Router)

```

user@host> show chassis fabric plane
Fabric management PLANE state
Plane 0
Plane state: ACTIVE
FEB 0: Links ok
FEB 1: Links ok
FEB 2: Links ok
FEB 3: Links ok
FEB 4: Links ok
FEB 5: Links ok
Plane 1
Plane state: ACTIVE
FEB 0: Links ok
FEB 1: Links ok
FEB 2: Links ok
FEB 3: Links ok
FEB 4: Links ok
FEB 5: Links ok
Plane 2
Plane state: ACTIVE
FEB 0: Links ok
FEB 1: Links ok
FEB 2: Links ok
FEB 3: Links ok
FEB 4: Links ok
FEB 5: Links ok
Plane 3
Plane state: ACTIVE
FEB 0: Links ok
FEB 1: Links ok

```

```
FEB 2: Links ok
FEB 3: Links ok
FEB 4: Links ok
FEB 5: Links ok
```

### show chassis fabric plane (MX240 Router)

```
user@host> show chassis fabric plane
```

```
Plane 0
  Plane state: ACTIVE
    FPC 1
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
    FPC 2
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 1
  Plane state: ACTIVE
    FPC 1
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
    FPC 2
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 2
  Plane state: ACTIVE
    FPC 1
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
    FPC 2
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 3
  Plane state: ACTIVE
    FPC 1
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
    FPC 2
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 4
  Plane state: SPARE
    FPC 1
      PFE 0 :Links ok
      PFE 1 :Links ok
```

```
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 2
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
Plane 5
  Plane state: SPARE
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 2
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
Plane 6
  Plane state: SPARE
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 2
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
Plane 7
  Plane state: SPARE
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 2
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
```

#### show chassis fabric plane (MX480 Router)

```
user@host> show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
Plane 1
  Plane state: ACTIVE
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
```

```

        PFE 3 :Links ok
Plane 2
  Plane state: ACTIVE
    FPC 1
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 3
  Plane state: ACTIVE
    FPC 1
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 4
  Plane state: SPARE
    FPC 1
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 5
  Plane state: SPARE
    FPC 1
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 6
  Plane state: SPARE
    FPC 1
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 7
  Plane state: SPARE
    FPC 1
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok

```

### show chassis fabric plane (MX960 Router)

```

user@host> show chassis fabric plane
Plane 0
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
      PFE 1 :Links ok
Plane 1
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
      PFE 1 :Links ok

```

```
Plane 2
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
      PFE 1 :Links ok
Plane 3
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
      PFE 1 :Links ok
```

### show chassis fabric plane (MX240 with AS MLC Modular Carrier Card)

In the following output, FPC 1 is the AS MLC modular carrier card (AS MCC).

```
user@host>show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
    FPC 1
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
Plane 1
  Plane state: ACTIVE
    Plane state: ACTIVE
    FPC 1
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
Plane 2
  Plane state: ACTIVE
    FPC 2
      PFE 0 :Links ok
    FPC 4
      PFE 0 :Links ok
      PFE 2 :Links ok
    FPC 5
      PFE 0 :Links ok
Plane 3
  Plane state: ACTIVE
    Plane state: ACTIVE
    FPC 1
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
Plane 4
  Plane state: ACTIVE
    Plane state: ACTIVE
    FPC 1
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
Plane 5
  Plane state: ACTIVE
    FPC 1
      PFE 0 :Unused
```



```

        FPC 2
          PFE 0 :Links ok
Plane 6
  Plane state: ACTIVE
    FPC 1
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
Plane 7
  Plane state: ACTIVE
    FPC 1
      PFE 0 :Unused
    FPC 2
      PFE 0 :Links ok

```

### show chassis fabric plane (MX480 with AS MLC Modular Carrier Card)

In the following output, FPC 5 is the AS MLC modular carrier card (AS MCC).

```

user@host>show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
    FPC 2
      PFE 0 :Links ok
    FPC 4
      PFE 0 :Links ok
      PFE 2 :Links ok
    FPC 5
      PFE 0 :Links ok
Plane 1
  Plane state: ACTIVE
    FPC 2
      PFE 0 :Links ok
    FPC 4
      PFE 0 :Links ok
      PFE 2 :Links ok
    FPC 5
      PFE 0 :Links ok
Plane 2
  Plane state: ACTIVE
    FPC 2
      PFE 0 :Links ok
    FPC 4
      PFE 0 :Links ok
      PFE 2 :Links ok
    FPC 5
      PFE 0 :Links ok
Plane 3
  Plane state: ACTIVE
    FPC 2
      PFE 0 :Links ok
    FPC 4
      PFE 0 :Links ok
      PFE 2 :Links ok
    FPC 5
      PFE 0 :Links ok
Plane 4
  Plane state: ACTIVE
    FPC 2

```

```
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 2 :Links ok
    FPC 5
        PFE 0 :Links ok
Plane 5
    Plane state: ACTIVE
    FPC 2
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 2 :Links ok
    FPC 5
        PFE 0 :Unused
Plane 6
    Plane state: ACTIVE
    FPC 2
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 2 :Links ok
    FPC 5
        PFE 0 :Links ok
Plane 7
    Plane state: ACTIVE
    FPC 2
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 2 :Links ok
    FPC 5
        PFE 0 :Unused
```

#### show chassis fabric plane (MX480 Router with MPC4E)

```
user@host > show chassis fabric plane
Fabric management PLANE state
Plane 0
    Plane state: ACTIVE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 3
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 1 :Links ok
Plane 1
    Plane state: ACTIVE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
```

```
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 3
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 1 :Links ok
Plane 2
Plane state: ACTIVE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 3
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 1 :Links ok
Plane 3
Plane state: ACTIVE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 3
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 1 :Links ok
Plane 4
Plane state: SPARE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 3
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 1 :Links ok
Plane 5
Plane state: SPARE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
```

```
        PFE 3 :Links ok
    FPC 3
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 1 :Links ok
Plane 6
  Plane state: SPARE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 3
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 1 :Links ok
Plane 7
  Plane state: SPARE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 3
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 1 :Links ok
```

#### show chassis fabric plane (MX960 with AS-MLC Modular Carrier Card)

In the following output, FPC 1 is a modular carrier card.

```
user@host>show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 5
        PFE 0 :Links ok
    FPC 8
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
```

```

        PFE 3 :Links ok
Plane 1
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 1
      PFE 0 :Links ok
    FPC 4
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
    FPC 5
      PFE 0 :Links ok
    FPC 8
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 2
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 1
      PFE 0 :Links ok
    FPC 4
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
    FPC 5
      PFE 0 :Links ok
    FPC 8
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 3
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 1
      PFE 0 :Links ok
    FPC 4
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
    FPC 5
      PFE 0 :Links ok
    FPC 8
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 4
  Plane state: SPARE
    FPC 0
```

```
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 5
        PFE 0 :Links ok
    FPC 8
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
Plane 5
Plane state: SPARE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
    FPC 4
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 5
        PFE 0 :Links ok
    FPC 8
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
```

### show chassis fabric plane (MX2010 Router)

```
user@host>show chassis fabric plane
Fabric management PLANE state
Plane 0
Plane state: ACTIVE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
    FPC 2
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 3
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 4
        PFE 0 :Links ok
    FPC 5
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 6
```

```

        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 7
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 8
        PFE 0 :Links ok
    FPC 9
        PFE 0 :Links ok
        PFE 1 :Links ok
Plane 1
  Plane state: ACTIVE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
    FPC 2
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 3
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 4
        PFE 0 :Links ok
    FPC 5
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 6
    PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 7
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 8
        PFE 0 :Links ok
    FPC 9
        PFE 0 :Links ok
        PFE 1 :Links ok
Plane 2
  Plane state: ACTIVE
    FPC 0
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 1
        PFE 0 :Links ok
    FPC 2
        PFE 0 :Links ok
        PFE 1 :Links ok
    FPC 3
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
    FPC 4
```

```
        PFE 0 :Links ok
FPC 5
        PFE 0 :Links ok
        PFE 1 :Links ok
FPC 6
        PFE 0 :Links ok
        PFE 1 :Links ok
        PFE 2 :Links ok
        PFE 3 :Links ok
FPC 7
        PFE 0 :Links ok
        PFE 1 :Links ok
FPC 8
        PFE 0 :Links ok
FPC 9
        PFE 0 :Links ok
        PFE 1 :Links ok
Plane 3
  Plane state: OFFLINE
Plane 4
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 1
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
PFE 1 :Links ok
    FPC 3
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
    FPC 4
      PFE 0 :Links ok
    FPC 5
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 6
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
    FPC 7
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 8
      PFE 0 :Links ok
    FPC 9
      PFE 0 :Links ok
      PFE 1 :Links ok
Plane 5
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 1
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
```



```
    PFE 1 :Links ok
FPC 3
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 4
    PFE 0 :Links ok
FPC 5
    PFE 0 :Links ok
    PFE 1 :Links ok
FPC 6
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 7
    PFE 0 :Links ok
    PFE 1 :Links ok
FPC 8
    PFE 0 :Links ok
FPC 9
    PFE 0 :Links ok
    PFE 1 :Links ok
Plane 6
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 1
      PFE 0 :Links ok
    FPC 2
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 3
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
    FPC 4
      PFE 0 :Links ok
    FPC 5
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 6
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
    FPC 7
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 8
      PFE 0 :Links ok
    FPC 9
      PFE 0 :Links ok
      PFE 1 :Links ok
Plane 7
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
```

```
    PFE 1 :Links ok
FPC 1
    PFE 0 :Links ok
FPC 2
    PFE 0 :Links ok
    PFE 1 :Links ok
FPC 3
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 4
    PFE 0 :Links ok
FPC 5
    PFE 0 :Links ok
    PFE 1 :Links ok
FPC 6
    PFE 0 :Links ok
PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 7
    PFE 0 :Links ok
    PFE 1 :Links ok
FPC 8
    PFE 0 :Links ok
FPC 9
    PFE 0 :Links ok
    PFE 1 :Links ok
```

#### show chassis fabric plane (MX2020 Router)

```
user@host>show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
  FPC 0
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 1
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 2
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 3
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 4
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
```

```
FPC 5
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 6
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 7
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 8
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 9
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 10
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 11
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 12
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 13
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 14
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 15
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 16
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 17
```

```

    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 18
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 19
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
Plane 1
  Plane state: ACTIVE
  FPC 0
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 1
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 2
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 3
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 4
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 5
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 6
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 7
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
  FPC 8
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
```

```
FPC 9
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 10
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 11
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 12
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 13
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 14
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 15
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 16
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 17
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 18
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 19
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
Plane 2
  Plane state: ACTIVE
  FPC 0
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
```

```
    PFE 3 :Links ok
FPC 1
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 2
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 3
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 4
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 5
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 6
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 7
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 8
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 9
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 10
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 11
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
FPC 12
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
```

```

FPC 13
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 14
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 15
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 16
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 17
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 18
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
FPC 19
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
Plane 3
...
```

#### show chassis fabric plane (MX2020 Router with MPC4E)

```

user@host > show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
  FPC 0
    PFE 0 :Links ok
    PFE 1 :Links ok
  FPC 9
    PFE 0 :Links ok
    PFE 1 :Links ok
  FPC 10
    PFE 0 :Links ok
  FPC 14
    PFE 0 :Links ok
    PFE 1 :Links ok
  FPC 19
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
Plane 1
```

```
Plane state: ACTIVE
  FPC 0
    PFE 0 :Links ok
    PFE 1 :Links ok
  FPC 9
    PFE 0 :Links ok
    PFE 1 :Links ok
  FPC 10
    PFE 0 :Links ok
  FPC 14
    PFE 0 :Links ok
    PFE 1 :Links ok
  FPC 19
    PFE 0 :Links ok
    PFE 1 :Links ok
    PFE 2 :Links ok
    PFE 3 :Links ok
Plane 2
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 9
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 10
      PFE 0 :Links ok
    FPC 14
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 19
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 3
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 9
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 10
      PFE 0 :Links ok
    FPC 14
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 19
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 4
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 9
      PFE 0 :Links ok
      PFE 1 :Links ok
```



```
FPC 10
  PFE 0 :Links ok
FPC 14
  PFE 0 :Links ok
  PFE 1 :Links ok
FPC 19
  PFE 0 :Links ok
  PFE 1 :Links ok
  PFE 2 :Links ok
  PFE 3 :Links ok
Plane 5
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 9
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 10
      PFE 0 :Links ok
    FPC 14
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 19
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 6
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 9
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 10
      PFE 0 :Links ok
    FPC 14
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 19
      PFE 0 :Links ok
      PFE 1 :Links ok
      PFE 2 :Links ok
      PFE 3 :Links ok
Plane 7
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 9
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 10
      PFE 0 :Links ok
    FPC 14
      PFE 0 :Links ok
      PFE 1 :Links ok
    FPC 19
      PFE 0 :Links ok
```

```

PFE 1 :Links ok
PFE 2 :Links ok
PFE 3 :Links ok

```

### show chassis fabric plane (TX Matrix Plus Router)

```
user@host> show chassis fabric plane
```

```
sfc0-re0:
```

Plane	State	Link errors	Destination errors	Uptime
0	Spare	NONE	NONE	
1	Online	NONE	NONE	10 hours, 16 seconds
2	Online	NONE	NONE	10 hours, 13 seconds
3	Online	NONE	NONE	10 hours, 9 seconds
4	Online	NONE	NONE	10 hours, 7 seconds

```
lcc0-re0:
```

SIB	State	Link errors	Destination errors	Uptime
0	Spare	NONE	NONE	
1	Online	NONE	NONE	10 hours, 16 seconds
2	Online	NONE	NONE	10 hours, 13 seconds
3	Online	NONE	NONE	10 hours, 9 seconds
4	Online	NONE	NONE	10 hours, 7 seconds

```
lcc2-re0:
```

SIB	State	Link errors	Destination errors	Uptime
0	Spare	NONE	NONE	
1	Online	NONE	NONE	10 hours, 16 seconds
2	Online	NONE	NONE	10 hours, 12 seconds
3	Online	NONE	NONE	10 hours, 9 seconds
4	Online	NONE	NONE	10 hours, 7 seconds

### show chassis fabric plane (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric plane
```

```
sfc0-re0:
```

Plane	State	Cable errors	Link errors	Destination errors	Uptime
0	Spare	NONE	NONE	NONE	
1	Online	NONE	NONE	NONE	5 hours, 11 minutes, 3 seconds
2	Online	NONE	NONE	NONE	8 hours, 4 minutes, 24 seconds
3	Online	NONE	NONE	NONE	8 hours, 3 minutes, 16 seconds
4	Online	NONE	NONE	NONE	8 hours, 2 minutes, 12 seconds

```
lcc2-re0:
```

SIB	State	Cable errors	Link errors	Destination errors	Uptime
0	Spare	NONE	NONE	NONE	
1	Online	NONE	NONE	NONE	5 hours, 11 minutes, 3 seconds
2	Online	NONE	NONE	NONE	8 hours, 4 minutes, 57 seconds
3	Online	NONE	NONE	NONE	8 hours, 3 minutes, 53 seconds
4	Online	NONE	NONE	NONE	8 hours, 2 minutes, 45 seconds

```
lcc4-re0:
```

```
-----
SIB   State      Cable errors  Link errors  Destination errors  Uptime
0     Spare      NONE         NONE         NONE
1     Online     NONE         NONE         NONE                5 hours, 11
minutes, 12 seconds
2     Online     NONE         NONE         NONE                8 hours, 4
minutes, 24 seconds
3     Online     NONE         NONE         NONE                8 hours, 3
minutes, 16 seconds
4     Online     NONE         NONE         NONE                8 hours, 2
minutes, 12 seconds
```

```
lcc5-re0:
```

```
-----
SIB   State      Cable errors  Link errors  Destination errors  Uptime
0     Spare      NONE         NONE         NONE
1     Online     NONE         NONE         NONE                5 hours, 11
minutes, 12 seconds
2     Online     NONE         NONE         NONE                8 hours, 4
minutes, 24 seconds
3     Online     NONE         NONE         NONE                8 hours, 3
minutes, 15 seconds
4     Online     NONE         NONE         NONE                8 hours, 2
minutes, 11 seconds
```

#### show chassis fabric plane detail (TX Matrix Plus Router)

```
user@host> show chassis fabric plane detail
sfc0-re0:
```

```
-----
Fabric Management PLANE State:
```

```
PLANE 0:   Spare
```

```
  SIB F13 0 : Spare
  SIB F13 1 : Empty
  SIB F2S 0/0 : Spare
  SIB F2S 0/2 : Spare
  SIB F2S 0/4 : Spare
  SIB F2S 0/6 : Spare
```

```
PLANE 1:   Online
```

```
  SIB F13 3 : Online
  SIB F13 4 : Empty
  SIB F2S 1/0 : Online
  SIB F2S 1/2 : Online
  SIB F2S 1/4 : Online
  SIB F2S 1/6 : Online
```

```
PLANE 2:   Online
```

```
  SIB F13 6 : Online
  SIB F13 7 : Empty
  SIB F2S 2/0 : Online
  SIB F2S 2/2 : Online
  SIB F2S 2/4 : Online
  SIB F2S 2/6 : Online
```

```
PLANE 3:   Online
```

```
  SIB F13 8 : Online
  SIB F13 9 : Online
  SIB F2S 3/0 : Online
  SIB F2S 3/2 : Online
  SIB F2S 3/4 : Online
  SIB F2S 3/6 : Online
```

```
PLANE 4:    Online
  SIB F13 11 :    Online
  SIB F13 12 :    Online
  SIB F2S 4/0 :    Online
  SIB F2S 4/2 :    Online
  SIB F2S 4/4 :    Online
  SIB F2S 4/6 :    Online
```

```
lcc0-re0:
```

```
-----
Fabric Management SIB State:
```

```
  SIB    0 :    Spare
  SIB    1 :    Online
  SIB    2 :    Online
  SIB    3 :    Online
  SIB    4 :    Online
```

```
lcc1-re0:
```

```
-----
Fabric Management SIB State:
```

```
  SIB    0 :    Spare
  SIB    1 :    Online
  SIB    2 :    Online
  SIB    3 :    Online
  SIB    4 :    Online
```

```
...
```

### show chassis fabric plane extensive (TX Matrix Plus Router )

```
user@host> show chassis fabric plane extensive
```

```
sfc0-re0:
```

```
-----
Fabric Management PLANE State:
```

```
PLANE 0:    Spare
```

```
  SIB F13 0 :    Spare
  SIB F13 1 :    Empty
  SIB F2S 0/0 :    Spare
  SIB F2S 0/2 :    Spare
  SIB F2S 0/4 :    Spare
  SIB F2S 0/6 :    Spare
  SIB F13 0 Even:
```

```
    LCC 0, SIB 0 : Links ok
```

```
    SG 0
```

```
      Port 0 : Links ok
      Port 1 : Links ok
      Port 2 : Links ok
      Port 3 : Links ok
```

```
    SG 1
```

```
      Port 0 : Links ok
      Port 1 : Links ok
      Port 2 : Links ok
      Port 3 : Links ok
```

```
    SG 2
```

```
      Port 0 : Links ok
      Port 1 : Links ok
      Port 2 : Links ok
      Port 3 : Links ok
```

```
    SG 3
```

```
      Port 0 : Links ok
      Port 1 : Links ok
      Port 2 : Links ok
```

```

Port 3      : Links ok
SIB F13 0 Odd:
LCC 1, SIB 0 : Links ok
SG 0
Port 0      : Links ok
Port 1      : Links ok
Port 2      : Links ok
Port 3      : Links ok
SG 1
Port 0      : Links ok
Port 1      : Links ok
Port 2      : Links ok
Port 3      : Links ok
SG 2
Port 0      : Links ok
Port 1      : Links ok
Port 2      : Links ok
Port 3      : Links ok
SG 3
Port 0      : Links ok
Port 1      : Links ok
Port 2      : Links ok
Port 3      : Links ok
SIB F2S 0/0: Links ok
SIB F2S 0/2: Links ok
SIB F2S 0/4: Links ok
SIB F2S 0/6: Links ok
SIB F13 1 Even:
LCC 2, SIB 0 : Unused
SG 0
Port 0      : Unused
Port 1      : Unused
Port 2      : Unused
Port 3      : Unused
SG 1
Port 0      : Unused
Port 1      : Unused
Port 2      : Unused
Port 3      : Unused
SG 2
Port 0      : Unused
Port 1      : Unused
Port 2      : Unused
Port 3      : Unused
SG 3
Port 0      : Unused
Port 1      : Unused
Port 2      : Unused
Port 3      : Unused
SIB F13 1 Odd:
LCC 3, SIB 0 : Unused
SG 0
Port 0      : Unused
Port 1      : Unused
Port 2      : Unused
Port 3      : Unused
SG 1
Port 0      : Unused
Port 1      : Unused
Port 2      : Unused
Port 3      : Unused

```

```
SG 2
  Port 0 : Unused
  Port 1 : Unused
  Port 2 : Unused
  Port 3 : Unused
SG 3
  Port 0 : Unused
  Port 1 : Unused
  Port 2 : Unused
  Port 3 : Unused
SIB F2S 0/0: Unused
SIB F2S 0/2: Unused
SIB F2S 0/4: Unused
SIB F2S 0/6: Unused
PLANE 1: Online
  SIB F13 3 : Online
  SIB F13 4 : Empty
  SIB F2S 1/0 : Online
  SIB F2S 1/2 : Online
  SIB F2S 1/4 : Online
  SIB F2S 1/6 : Online
  SIB F13 3 Even:
...
```

#### show chassis fabric plane extensive (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric plane extensive
sfc0-re0:
```

```
-----
Fabric Management PLANE State:
```

```
PLANE 0: Online
  SIB F13 0 : Empty
  SIB F13 1 : Online
  SIB F2S 0/0 : Online
  SIB F2S 0/2 : Online
  SIB F2S 0/4 : Online
  SIB F2S 0/6 : Online
  SIB F13 0
    LCC 0, SIB 0 : Unused
    PFE 0 : Unused
    PFE 1 : Unused
    PFE 2 : Unused
    PFE 3 : Unused
    PFE 4 : Unused
    PFE 5 : Unused
    PFE 6 : Unused
    PFE 7 : Unused
    PFE 8 : Unused
    PFE 9 : Unused
    PFE 10 : Unused
    PFE 11 : Unused
    PFE 12 : Unused
    PFE 13 : Unused
    PFE 14 : Unused
    PFE 15 : Unused
    LCC 1, SIB 0 : Unused
    PFE 0 : Unused
    PFE 1 : Unused
    PFE 2 : Unused
    PFE 3 : Unused
    PFE 4 : Unused
```

```

PFE 5 : Unused
PFE 6 : Unused
PFE 7 : Unused
PFE 8 : Unused
PFE 9 : Unused
PFE 10 : Unused
PFE 11 : Unused
PFE 12 : Unused
PFE 13 : Unused
PFE 14 : Unused
PFE 15 : Unused
LCC 2, SIB 0 : Unused
PFE 0 : Unused
PFE 1 : Unused
PFE 2 : Unused
PFE 3 : Unused
PFE 4 : Unused
PFE 5 : Unused
PFE 6 : Unused
PFE 7 : Unused
PFE 8 : Unused
PFE 9 : Unused
PFE 10 : Unused
...
lcc5-re0:
-----
Fabric Management SIB State:
SIB 0 : Online
LCC SIB Link State : Links ok
PFE 0 : Links ok
PFE 1 : Links ok
PFE 2 : Links ok
PFE 3 : Links ok
PFE 4 : Links ok
PFE 5 : Links ok
PFE 6 : Links ok
PFE 7 : Links ok
PFE 8 : Links ok
PFE 9 : Links ok
PFE 10 : Links ok
PFE 11 : Links ok
PFE 12 : Links ok
PFE 13 : Links ok
PFE 14 : Links ok
PFE 15 : Links ok
FPC 1
PFE 0 : Links ok
FPC 2
PFE 0 : Links ok
FPC 3
PFE 0 : Links ok
PFE 1 : Links ok
FPC 4
PFE 0 : Links ok
SIB 1 : Online
LCC SIB Link State : Links ok
PFE 0 : Links ok
PFE 1 : Links ok
PFE 2 : Links ok
PFE 3 : Links ok
PFE 4 : Links ok

```

```
PFE 5 : Links ok
PFE 6 : Links ok
PFE 7 : Links ok
PFE 8 : Links ok
PFE 9 : Links ok
PFE 10 : Links ok
PFE 11 : Links ok
PFE 12 : Links ok
PFE 13 : Links ok
PFE 14 : Links ok
PFE 15 : Links ok
FPC 1
  PFE 0 : Links ok
FPC 2
  PFE 0 : Links ok
FPC 3
  PFE 0 : Links ok
  PFE 1 : Links ok
FPC 4
  PFE 0 : Links ok
```

#### show chassis fabric plane terse (TX Matrix Plus Router)

```
user@host> show chassis fabric plane terse
sfc0-re0:
```

Plane	State	Link errors	Destination errors	Uptime
0	Spare	NONE	NONE	
1	Online	NONE	NONE	18 minutes, 37 seconds
2	Online	NONE	NONE	18 minutes, 36 seconds
3	Online	NONE	NONE	18 minutes, 33 seconds
4	Online	NONE	NONE	18 minutes, 31 seconds

```
lcc1-re0:
```

SIB	State	Link errors	Destination errors	Uptime
0	Spare	NONE	NONE	
1	Online	NONE	NONE	18 minutes, 37 seconds
2	Online	NONE	NONE	
3	Online	NONE	NONE	
4	Empty	NONE	NONE	

```
lcc2-re0:
```

SIB	State	Link errors	Destination errors	Uptime
0	Spare	NONE	NONE	
1	Online	NONE	NONE	18 minutes, 37 seconds
2	Online	NONE	NONE	18 minutes, 36 seconds
3	Online	NONE	NONE	18 minutes, 32 seconds
4	Online	NONE	NONE	18 minutes, 31 seconds

#### show chassis fabric plane terse (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fabric plane terse
```



sfc0-re0:

Plane	State	Cable errors	Link errors	Destination errors	Uptime
0	Offline	NONE	NONE	NONE	
1	Online	NONE	NONE	NONE	1 day, 18 hours, 14 minutes, 26 seconds
2	Offline	NONE	NONE	NONE	
3	Offline	NONE	NONE	NONE	
4	Offline	NONE	NONE	NONE	

lcc2-re0:

SIB	State	Cable errors	Link errors	Destination errors	Uptime
0	Offline	NONE	NONE	NONE	
1	Online	NONE	NONE	NONE	1 day, 18 hours, 17 minutes
2	Offline	NONE	NONE	NONE	
3	Offline	NONE	NONE	NONE	
4	Offline	NONE	NONE	NONE	

lcc4-re0:

SIB	State	Cable errors	Link errors	Destination errors	Uptime
0	Offline	NONE	NONE	NONE	
1	Online	NONE	NONE	NONE	1 day, 18 hours, 14 minutes, 38 seconds
2	Offline	NONE	NONE	NONE	
3	Offline	NONE	NONE	NONE	
4	Offline	NONE	NONE	NONE	

lcc5-re0:

SIB	State	Cable errors	Link errors	Destination errors	Uptime
0	Offline	NONE	NONE	NONE	
1	Online	NONE	NONE	NONE	1 day, 18 hours, 14 minutes, 34 seconds
2	Offline	NONE	NONE	NONE	
3	Offline	NONE	NONE	NONE	
4	Offline	NONE	NONE	NONE	

#### show chassis fabric plane lcc (TX Matrix Plus Router)

user@host> show chassis fabric plane lcc 7

lcc1-re0:

SIB	State	Link errors	Destination errors	Uptime
0	Spare	NONE	NONE	
1	Online	NONE	NONE	25 minutes, 17 seconds
2	Disconnected	NONE	NONE	
3	Disconnected	NONE	NONE	
4	Empty	NONE	NONE	

#### show chassis fabric plane lcc (TX Matrix Plus Router with 3D SIBs)

user@host> show chassis fabric plane lcc 2

lcc2-re0:

SIB	State	Cable errors	Link errors	Destination errors	Uptime
0	Offline	NONE	NONE	NONE	
1	Online	NONE	NONE	NONE	1 day, 18 hours, 14 minutes, 34 seconds

```

hours, 16 minutes, 44 seconds
2    Offline      NONE      NONE      NONE
3    Offline      NONE      NONE      NONE
4    Offline      NONE      NONE      NONE

```

#### show chassis fabric plane sfc (TX Matrix Plus Router)

```

user@host> show chassis fabric plane sfc 0
sfc0-re0:

```

```

-----
Plane  State          Link errors  Destination errors  Uptime
0      Spare          NONE        NONE                NONE
1      Online         NONE        NONE                27 minutes, 7 seconds
2      Online         NONE        NONE                27 minutes, 6 seconds
3      Online         NONE        NONE                27 minutes, 3 seconds
4      Online         NONE        NONE                27 minutes, 1 second

```

#### show chassis fabric plane sfc (TX Matrix Plus Router with 3D SIBs)

```

user@host> show chassis fabric plane sfc 0
sfc0-re0:

```

```

-----
Plane  State          Cable errors  Link errors  Destination errors  Uptime
0      Offline      NONE        NONE        NONE                NONE
1      Online       NONE        NONE        NONE                1 day, 18
hours, 14 minutes, 20 seconds
2      Offline      NONE        NONE        NONE                NONE
3      Offline      NONE        NONE        NONE                NONE
4      Offline      NONE        NONE        NONE                NONE

```

#### show chassis fabric plane (T1600 Router)

```

user@host> show chassis fabric plane
Plane  State          Uptime
0      Online       15 hours, 42 minutes, 9 seconds
1      Online       15 hours, 42 minutes, 9 seconds
2      Fault
3      Online       15 hours, 42 minutes, 9 seconds
4      Online       15 hours, 42 minutes, 9 seconds

```

#### show chassis fabric plane extensive (T1600 Router)

```

user@host> show chassis fabric plane extensive
Fabric Management PLANE State:
PLANE 0:   Online
  ST-SIB-L 0: Links ok
    SG 0
      Port 0   : Links ok
      Port 1   : Links ok
      Port 2   : Links ok
      Port 3   : Links ok
    SG 1
      Port 0   : Links ok
      Port 1   : Links ok
      Port 2   : Links ok
      Port 3   : Links ok
    SG 2
      Port 0   : Links ok
      Port 1   : Links ok

```

```

    Port 2    : Links ok
    Port 3    : Links ok
SG 3
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
ST-SIB-L 0
  FPC 4
    PFE 0: Links ok
    PFE 1: Links ok
  FPC 6
    PFE 0: Links ok
    PFE 1: Links ok
  FPC 7
    PFE 0: Links ok
PLANE 1:   Online
ST-SIB-L 1: Links ok
SG 0
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
SG 1
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
SG 2
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
SG 3
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
ST-SIB-L 1
  FPC 4
    PFE 0: Links ok
    PFE 1: Links ok
  FPC 6
    PFE 0: Links ok
    PFE 1: Links ok
  FPC 7
    PFE 0: Links ok
PLANE 2:   Online
ST-SIB-L 2: Links ok
SG 0
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
SG 1
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
SG 2
    Port 0    : Links ok
```

```

    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
  SG 3
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
  ST-SIB-L 2
    FPC 4
      PFE 0: Links ok
      PFE 1: Links ok
    FPC 6
      PFE 0: Links ok
      PFE 1: Links ok
    FPC 7
      PFE 0: Links ok
  PLANE 3:   Spare
  ST-SIB-L 3: Links ok
  SG 0
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
  SG 1
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
  SG 2
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
  SG 3
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
  ST-SIB-L 3
    FPC 4
      PFE 0: Links ok
      PFE 1: Links ok
    FPC 6
      PFE 0: Links ok
      PFE 1: Links ok
    FPC 7
      PFE 0: Links ok
  PLANE 4:   Online
  ST-SIB-L 4: Links ok
  SG 0
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
  SG 1
    Port 0    : Links ok
    Port 1    : Links ok
    Port 2    : Links ok
    Port 3    : Links ok
  SG 2
```

```

Port 0    : Links ok
Port 1    : Links ok
Port 2    : Links ok
Port 3    : Links ok
SG 3
Port 0    : Links ok
Port 1    : Links ok
Port 2    : Links ok
Port 3    : Links ok
ST-SIB-L 4
FPC 4
PFE 0: Links ok
PFE 1: Links ok
FPC 6
PFE 0: Links ok
PFE 1: Links ok
FPC 7
PFE 0: Links ok

```

#### show chassis fabric plane detail (T1600 Router)

```

user@host> show chassis fabric plane detail
Fabric Management PLANE State:
PLANE 0:   Online
PLANE 1:   Online
PLANE 2:   Online
PLANE 3:   Spare
PLANE 4:   Online

```

#### show chassis fabric plane (EX8200 Switch)

```

user@host> show chassis fabric plane
Fabric management PLANE state
Plane 0
Plane state: ACTIVE
Plane 1
Plane state: ACTIVE
Plane 2
Plane state: ACTIVE
Plane 3
Plane state: ACTIVE
Plane 4
Plane state: SPARE
Plane 5
Plane state: SPARE
Plane 6
Plane state: SPARE
Plane 7
Plane state: SPARE
Plane 8
Plane state: ACTIVE
Plane 9
Plane state: ACTIVE
Plane 10
Plane state: ACTIVE
Plane 11
Plane state: ACTIVE

```

## show chassis fabric plane-location

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1150</a> <a href="#">Syntax (MX Series Routers) on page 1150</a> <a href="#">Syntax (MX2010 3D Universal Edge Routers) on page 1150</a> <a href="#">Syntax (MX2020 3D Universal Edge Routers) on page 1150</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1150</a>
<b>Syntax</b>	show chassis fabric plane-location
<b>Syntax (MX Series Routers)</b>	show chassis fabric plane-location <all-members> <local> <member <i>member-id</i> >
<b>Syntax (MX2010 3D Universal Edge Routers)</b>	show chassis fabric plane-location
<b>Syntax (MX2020 3D Universal Edge Routers)</b>	show chassis fabric plane-location
<b>Syntax (TX Matrix Plus Router)</b>	show chassis fabric plane-location
<b>Release Information</b>	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.4 for EX Series switches. Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
<b>Description</b>	(M120, MX Series routers, and EX8200 switches only) Display the Control Board (CB) location of each plane. This command can be used on the master Routing Engine or the backup Routing Engine. For information about the meaning of “CBs” and “fabric plane” on the switches, see <a href="#">“EX Series Switches Hardware and CLI Terminology Mapping” on page 315</a> .  (TX Matrix Plus routers only) Display the SIB location of each fabric plane.  (PTX Series Packet Transport Routers only) Display the fabric plane location of each SIB.  (MX2010 and MX2020 Routers only) Display the fabric plane location of each Switch Fabric Board (SFB).
<b>Options</b>	<b>all-members</b> —(MX Series routers only) (Optional) Display the CB location of each fabric plane on the Routing Engines in all member routers in the Virtual Chassis configuration.

**local**—(MX Series routers only) (Optional) Display the CB location of each fabric plane on the Routing Engines in the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display the CB location of each fabric plane on the Routing Engines in the specified member in the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

**Required Privilege Level** view

**List of Sample Output** [show chassis fabric plane-location \(M120 Router\) on page 1152](#)  
[show chassis fabric plane-location \(MX240 and MX480 Routers\) on page 1152](#)  
[show chassis fabric plane-location \(MX960 Router\) on page 1152](#)  
[show chassis fabric plane-location \(MX2010 Router\) on page 1152](#)  
[show chassis fabric plane-location \(MX2020 Router\) on page 1152](#)  
[show chassis fabric plane-location \(TX Matrix Plus Router\) on page 1153](#)  
[show chassis fabric plane-location \(TX Matrix Plus Router with 3D SIBs\) on page 1153](#)  
[show chassis fabric plane-location \(EX8200 Switch\) on page 1153](#)  
[show chassis fabric plane-location \(PTX Series Packet Transport Routers\) on page 1153](#)

**Output Fields** [Table 114 on page 1151](#) lists the output fields for the **show chassis fabric plane-location** command. Output fields are listed in the approximate order in which they appear.

**Table 114: show chassis fabric plane-location Output Fields**

Field Name	Field Description
<b>Plane <i>n</i></b>	Plane number.  (PTX Series Packet Transport Routers only) Plane numbers associated with the SIB.  (MX2010 and MX2020 Routers only) Plane numbers associated with the SFB.
<b>Control Board <i>n</i></b>	Control board number.
<b>SFC ABS-SIB-F13</b>	(TX Matrix Plus routers only) Switch Interface Board (SIB) slot number on the F13 SIB.
<b>SFC ABS-SIB-F2S</b>	(TX Matrix Plus routers only) SIB slot number on the F2S SIB.
<b>LCC ST-SIB-L</b>	(TX Matrix Plus routers only) Line-card chassis (LCC) SIB slot number.
<b>SFC SIB F13</b>	(TX Matrix Plus routers with 3D SIBs only) Switch Interface Board (SIB) slot number on the F13 SIB.
<b>SFC SIB F2S</b>	(TX Matrix Plus routers with 3D SIBs only) SIB slot number on the F2S SIB.
<b>LCC SIB</b>	(TX Matrix Plus routers with 3D SIBs only) Line-card chassis (LCC) SIB slot number.

Table 114: show chassis fabric plane-location Output Fields (*continued*)

Field Name	Field Description
SIB	(PTX Series Packet Transport Routers only) SIB number.
Switch Fabric Board <i>n</i>	(MX2010 and MX2020 Routers only) SFB number.

## Sample Output

### show chassis fabric plane-location (M120 Router)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
Plane 0                Control Board 0
Plane 1                Control Board 0
Plane 2                Control Board 1
Plane 3                Control Board 1

```

### show chassis fabric plane-location (MX240 and MX480 Routers)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
Plane 0                Control Board 0
Plane 1                Control Board 0
Plane 2                Control Board 0
Plane 3                Control Board 0
Plane 4                Control Board 1
Plane 5                Control Board 1
Plane 6                Control Board 1
Plane 7                Control Board 1

```

### show chassis fabric plane-location (MX960 Router)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
Plane 0                Control Board 0
Plane 1                Control Board 0
Plane 2                Control Board 1
Plane 3                Control Board 1
Plane 4                Control Board 2
Plane 5                Control Board 2

```

### show chassis fabric plane-location (MX2010 Router)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
Plane 0                Switch Fabric Board 0
Plane 1                Switch Fabric Board 1
Plane 2                Switch Fabric Board 2
Plane 3                Switch Fabric Board 3
Plane 4                Switch Fabric Board 4
Plane 5                Switch Fabric Board 5
Plane 6                Switch Fabric Board 6
Plane 7                Switch Fabric Board 7

```

### show chassis fabric plane-location (MX2020 Router)

```

user@host> show chassis fabric plane-location

```



```

-----Fabric Plane Locations-----
Plane 0          Switch Fabric Board 0
Plane 1          Switch Fabric Board 1
Plane 2          Switch Fabric Board 2
Plane 3          Switch Fabric Board 3
Plane 4          Switch Fabric Board 4
Plane 5          Switch Fabric Board 5
Plane 6          Switch Fabric Board 6
Plane 7          Switch Fabric Board 7

```

#### show chassis fabric plane-location (TX Matrix Plus Router)

```

user@host> show chassis fabric plane-location
Fabric Plane Locations :
Plane      SFC ABS-SIB-F13      SFC ABS-SIB-F2      LCC ST-SIB-L
0          0, 1                0/0, 0/2, 0/4, 0/6      0
1          3, 4                1/0, 1/2, 1/4, 1/6      1
2          6, 7                2/0, 2/2, 2/4, 2/6      2
3          8, 9                3/0, 3/2, 3/4, 3/6      3
4          11, 12             4/0, 4/2, 4/4, 4/6      4

```

#### show chassis fabric plane-location (TX Matrix Plus Router with 3D SIBs)

```

user@host> show chassis fabric plane-location
sfc0-re0
-----Fabric Plane Locations-----
Plane      SFC SIB F13      SFC SIB F2      LCC SIB
0          0, 1                0/0, 0/2, 0/4, 0/6      0
1          3, 4                1/0, 1/2, 1/4, 1/6      1
2          6, 7                2/0, 2/2, 2/4, 2/6      2
3          8, 9                3/0, 3/2, 3/4, 3/6      3
4          11, 12             4/0, 4/2, 4/4, 4/6      4

```

#### show chassis fabric plane-location (EX8200 Switch)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
Plane 0          Control Board 0
Plane 1          Control Board 0
Plane 2          Control Board 0
Plane 3          Control Board 0
Plane 4          Control Board 1
Plane 5          Control Board 1
Plane 6          Control Board 1
Plane 7          Control Board 1
Plane 8          Control Board 2
Plane 9          Control Board 2
Plane 10         Control Board 2
Plane 11         Control Board 2

```

#### show chassis fabric plane-location (PTX Series Packet Transport Routers)

```

user@host> show chassis fabric plane-location
-----Fabric Plane Locations-----
SIB      Planes
0         0 1
1         2 3
2         4 5
3         6 7
4         8 9
5        10 11

```

6	12	13
7	14	15
8	16	17

## show chassis fabric summary

<b>Syntax</b>	show chassis fabric summary
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.4.</p> <p>Command introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
<b>Description</b>	(MX Series routers and EX8200 switches only) Display the state of all fabric planes and the elapsed uptime.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show chassis fabric summary (MX240 Router) on page 1157</a></p> <p><a href="#">show chassis fabric summary (MX480 Router) on page 1157</a></p> <p><a href="#">show chassis fabric summary (MX480 Router with MPC4E) on page 1157</a></p> <p><a href="#">show chassis fabric summary (MX960 Router) on page 1157</a></p> <p><a href="#">show chassis fabric summary (MX2010 Router) on page 1158</a></p> <p><a href="#">show chassis fabric summary (MX2020 Router) on page 1158</a></p> <p><a href="#">show chassis fabric summary (MX2020 Router with MPC4E) on page 1158</a></p> <p><a href="#">show chassis fabric summary (EX8200 Switch) on page 1158</a></p> <p><a href="#">show chassis fabric summary (PTX Series Packet Transport Router) on page 1159</a></p>
<b>Output Fields</b>	<p><a href="#">Table 115 on page 1155</a> lists the output fields for the <b>show chassis fabric summary</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 115: show chassis fabric summary Output Fields**

Field Name	Field Description
Plane	(MX Series, MX2020 and MX2010 Routers only) Plane number.

Table 115: show chassis fabric summary Output Fields (*continued*)

Field Name	Field Description
<b>State</b>	<p>(MX Series) State of the SIB or FPC:</p> <ul style="list-style-type: none"> <li>• <b>Online</b>—Switch Interface Board (SIB) is operational and running.</li> </ul> <p><b>NOTE:</b> On the Enhanced MX SCB with Trio MPC, a maximum of 4 planes are operational and running. On all the other SCBs with Trio MPC, all the planes are operational and running.</p> <ul style="list-style-type: none"> <li>• <b>Empty</b>—SIB is powered down.</li> <li>• <b>Check</b>—SIB is in the <b>Check</b> state because of the following reasons: <ul style="list-style-type: none"> <li>• SIB is not inserted properly.</li> <li>• Some destination errors are detected on the SIB. In this case, the Packet Forwarding Engine stops using the SIB to send traffic to the affected destination Packet Forwarding Engine.</li> <li>• Some link errors are detected on the channel between the SIB and a Packet Forwarding Engine. Link errors can be detected at initialization time or runtime: <ul style="list-style-type: none"> <li>• Link errors caused by a link training failure at initialization time—The Packet Forwarding Engine does not use the SIB to send traffic. The <b>show chassis fabric fpcs</b> command shows <b>Plane disabled</b> as status for this link.</li> <li>• Link errors caused by CRC errors detected at runtime—The Packet Forwarding Engine continues to use the SIB to send traffic. The <b>show chassis fabric fpcs</b> command shows <b>Link error</b> as the status for this link.</li> </ul> </li> </ul> </li> </ul> <p><b>NOTE:</b> The <b>Check</b> state does not apply to PTX Series Packet Transport Routers because there are no SIBs in the Check state.</p> <p>For information about link and destination errors, issue the <b>show chassis fabric fpcs</b> commands.</p> <ul style="list-style-type: none"> <li>• <b>Spare</b>—SIB is redundant and will move to active state if one of the working SIBs fails.</li> </ul> <p><b>NOTE:</b> <b>Spare</b> does not apply to PTX Series Packet Transport Routers because there are no spare SIBs in the device.</p> <p>(MX2010 and MX2020 Routers) State of the SFB.</p> <ul style="list-style-type: none"> <li>• <b>Online</b>—Switch Fabric Board (SFB) is operational and running.</li> <li>• <b>Offline</b>—Switch Fabric Board (SFB) is powered down.</li> <li>• <b>Check</b>—Switch Fabric Board (SFB) is in the check state.</li> </ul>
<b>Errors</b>	<p>(PTX Series only) Indicates whether there is any error on the SIB.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—No errors</li> <li>• <b>Link Errors</b>—Fabric link errors were found on the SIB RX link.</li> <li>• <b>Cell drops</b>—Fabric cell drops were found on the SIB ASIC.</li> <li>• <b>Link, Cell drops</b>—Both Link errors and cell drops were detected on at least one of the FPC's fabric links.</li> </ul> <p><b>NOTE:</b> The <b>Errors</b> column is empty only when the FPC or SIB is offline.</p>

Table 115: show chassis fabric summary Output Fields (*continued*)

Field Name	Field Description
<b>Uptime</b>	(MX Series, MX2010 and MX2020 Routers) Elapsed time the plane has been online.

## Sample Output

### show chassis fabric summary (MX240 Router)

```
user@host> show chassis fabric summary
Plane  State  Uptime
0      Online 23 hours, 26 minutes, 54 seconds
1      Online 23 hours, 26 minutes, 54 seconds
2      Check 18 hours, 33 minutes, 42 seconds
3      Online 23 hours, 26 minutes, 54 seconds
4      Spare 23 hours, 26 minutes, 54 seconds
5      Spare 23 hours, 26 minutes, 54 seconds
6      Spare 23 hours, 26 minutes, 54 seconds
7      Spare 23 hours, 26 minutes, 54 seconds
```

### show chassis fabric summary (MX480 Router)

```
user@host> show chassis fabric summary
Plane  State  Uptime
0      Online 8 hours, 45 minutes, 29 seconds
1      Online 8 hours, 45 minutes, 28 seconds
2      Online 8 hours, 45 minutes, 28 seconds
3      Online 8 hours, 45 minutes, 28 seconds
4      Spare 8 hours, 45 minutes, 28 seconds
5      Spare 8 hours, 45 minutes, 28 seconds
6      Spare 8 hours, 45 minutes, 28 seconds
7      Check 6 hours, 10 minutes, 12 seconds
```

### show chassis fabric summary (MX480 Router with MPC4E)

```
user@host > show chassis fabric summary
Plane  State  Uptime
0      Online 6 hours, 57 minutes, 44 seconds
1      Online 6 hours, 57 minutes, 40 seconds
2      Online 6 hours, 57 minutes, 39 seconds
3      Online 6 hours, 57 minutes, 34 seconds
4      Spare 6 hours, 57 minutes, 34 seconds
5      Spare 6 hours, 57 minutes, 29 seconds
6      Spare 6 hours, 57 minutes, 29 seconds
7      Spare 6 hours, 57 minutes, 24 seconds

Note:
For FPC slots with MPC Type 4 or MCC:
Fabric planes 1 and 5, 3 and 7 use shared physical links.
Those slots may run in a reduced bandwidth in case both
plane 1 and 5, or both 3 and 7 are active.
```

### show chassis fabric summary (MX960 Router)

```
user@host> show chassis fabric summary
Plane  State  Uptime
0      Online 3 hours, 7 minutes, 9 seconds
1      Online 3 hours, 7 minutes, 4 seconds
```

2	Online	3 hours, 6 minutes, 59 seconds
3	Online	3 hours, 6 minutes, 54 seconds
4	Empty	
5	Empty	

#### show chassis fabric summary (MX2010 Router)

```
user@host> show chassis fabric summary
```

Plane	State	Uptime
0	Online	1 day, 13 hours, 20 minutes, 10 seconds
1	Online	1 day, 13 hours, 19 minutes, 59 seconds
2	Online	1 day, 13 hours, 19 minutes, 49 seconds
3	Offline	
4	Online	1 day, 13 hours, 19 minutes, 28 seconds
5	Check	1 day, 13 hours, 19 minutes, 17 seconds
6	Online	1 day, 13 hours, 19 minutes, 6 seconds
7	Online	1 hour, 43 minutes, 5 seconds

#### show chassis fabric summary (MX2020 Router)

```
user@host> show chassis fabric summary
```

Plane	State	Uptime
0	Online	8 hours, 24 minutes, 1 second
1	Online	8 hours, 47 minutes, 54 seconds
2	Online	8 hours, 47 minutes, 44 seconds
3	Online	8 hours, 47 minutes, 33 seconds
4	Online	8 hours, 47 minutes, 22 seconds
5	Online	8 hours, 47 minutes, 12 seconds
6	Online	8 hours, 47 minutes, 1 second
7	Online	8 hours, 46 minutes, 50 seconds

#### show chassis fabric summary (MX2020 Router with MPC4E)

```
user@host > show chassis fabric summary
```

Plane	State	Uptime
0	Online	3 days, 6 hours, 58 minutes, 29 seconds
1	Online	3 days, 6 hours, 58 minutes, 18 seconds
2	Online	3 days, 6 hours, 58 minutes, 8 seconds
3	Online	3 days, 6 hours, 57 minutes, 57 seconds
4	Online	3 days, 6 hours, 57 minutes, 46 seconds
5	Online	3 days, 6 hours, 57 minutes, 36 seconds
6	Online	3 days, 6 hours, 57 minutes, 25 seconds
7	Online	3 days, 6 hours, 57 minutes, 14 seconds

#### show chassis fabric summary (EX8200 Switch)

```
user@host> show chassis fabric summary
```

Plane	State	Uptime
0	Online	12 days, 50 minutes, 54 seconds
1	Online	12 days, 50 minutes, 53 seconds
2	Online	12 days, 50 minutes, 53 seconds
3	Online	12 days, 50 minutes, 52 seconds
4	Spare	12 days, 50 minutes, 49 seconds
5	Spare	12 days, 50 minutes, 47 seconds
6	Spare	12 days, 50 minutes, 47 seconds
7	Spare	12 days, 50 minutes, 46 seconds
8	Online	12 days, 50 minutes, 52 seconds
9	Online	12 days, 50 minutes, 50 seconds
10	Online	12 days, 50 minutes, 50 seconds
11	Online	12 days, 50 minutes, 49 seconds

**show chassis fabric summary (PTX Series Packet Transport Router)**

```
user@host> show chassis fabric summary
```

FRU	State	Errors
SIB0	Online	None
SIB1	Online	Link Errors
SIB2	Online	None
SIB3	Online	Cell drops
SIB4	Offline	
SIB5	Online	None
SIB6	Online	Link, Cell drops
SIB7	Online	None
SIB8	Online	Link, Cell drops
FPC0	Online	None
FPC1	Online	Link Errors
FPC2	Online	None
FPC3	Offline	
FPC4	Online	None
FPC5	Online	None
FPC6	Empty	
FPC7	Empty	

## show chassis fpc

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1160</a> <a href="#">Syntax (EX Series Switches) on page 1160</a> <a href="#">Syntax (T4000 Routers) on page 1160</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 1160</a> <a href="#">Syntax (MX Series Routers and EX Series switches) on page 1160</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 1160</a> <a href="#">Syntax (QFX Series) on page 1160</a> <a href="#">Syntax (PTX Series Packet Transport Routers) on page 1160</a> <a href="#">Syntax (ACX Series Universal Access Routers) on page 1160</a>
<b>Syntax</b>	show chassis fpc <detail <slot>>   <pic-status <slot>>
<b>Syntax (EX Series Switches)</b>	show chassis fpc <detail <fpc-slot>>   <pic-status <fpc-slot>> <fpc-slot>
<b>Syntax (T4000 Routers)</b>	show chassis fpc <detail <fpc-slot>> <pic-status <fpc-slot>>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	show chassis fpc <detail <fpc-slot>>   <pic-status <fpc-slot>> <slot>
<b>Syntax (MX Series Routers and EX Series switches)</b>	show chassis fpc <detail <slot>>   <pic-status <slot>> <all-members> <local> <member <i>member-id</i> >
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)</b>	show chassis fpc <slot> detail   <detail <slot>>   <pic-status <slot>> <fpc-slot>
<b>Syntax (QFX Series)</b>	show chassis fpc <detail> <interconnect-device <i>name</i> <fpc-slot <i>fpc-slot</i> >> <node-device <i>name</i> >
<b>Syntax (PTX Series Packet Transport Routers)</b>	show chassis fpc <detail <fpc-slot>>   <pic-status <fpc-slot>> <fpc-slot>
<b>Syntax (ACX Series Universal Access Routers)</b>	show chassis fpc <detail <fpc-slot>>   <pic-status <fpc-slot>> <fpc-slot>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.



Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 11.1 for QFX Series.

Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.

Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.

Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.

**Description** Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.

**Options** **none**—Display status information for all FPCs. On a TX Matrix router, display status information for all FPCs on the attached T640 routers in the routing matrix. On a TX Matrix Plus router, display status information for all FPCs on the attached routers in the routing matrix.



**NOTE:** In EX8200 switches, line cards initialize Packet Forwarding Engine during startup. If an error occurs during hardware initialization, the FPCs with bad hardware parts power down after transferring the debug information to the Routing Engine. The Routing Engine marks the FPC offline, logs the error in system log messages (/var/log/messages), and generates an alarm to inform the user.

See the following sample output:

```
user@host> show chassis fpc
```

Utilization (%)	Temp	CPU	Utilization (%)	Memory
Slot State	(C)	Total	Interrupt	DRAM (MB) Heap
Buffer				
0 Empty				
1 Empty				
2 Empty				
3 Empty				
4 Empty				
5 Offline	---	Hard FPC error---		
6 Empty				
7 Online	26	4	0	1024 0
32				

The following sample output shows the alarm raised for the failed FPCs.

```
user@host > show chassis alarms
4 alarms currently active
```

Alarm time	Class	Description
2011-03-24 00:52:51 UTC	Major	FPC 5 Hard errors
2011-03-24 00:52:31 UTC	Major	Fan Tray Failure
2011-03-24 00:52:31 UTC	Major	Fan Tray Failure
2011-03-24 00:51:26 UTC	Minor	Loss of communication with Backup RE



**NOTE:** On T4000 routers, when you include the enhanced-mode statement at the [edit chassis network-services] hierarchy level and reboot the system, only the T4000 Type 5 FPCs present on the router become online while the remaining FPCs are offline, and FPC misconfiguration alarms are generated. The show chassis alarm command output displays FPC misconfiguration (FPC *fpc-slot* misconfig) as the reason for the generation the alarms.

The following sample output shows the FPC status after the enhanced-mode statement is configured on the T4000 router. The T4000 Type 5 FPC present in slot 5 becomes online while the remaining FPCs are offline.

```
user@host> show chassis fpc
```

	Temp	CPU Utilization (%)	Memory
Utilization (%)			
Slot State	(C)	Total	Interrupt
Buffer			DRAM (MB) Heap
0 offline	---	FPC misconfiguration---	
1 offline	---	FPC misconfiguration---	
2 offline	---	FPC misconfiguration---	
3 Empty			
4 Empty			
5 Online	66	50	0
27			2816 29

The following sample output shows FPC misconfiguration alarms.

```
user@host > show chassis alarms
3 alarms currently active
Alarm time      Class  Description
2011-03-24 00:52:51 PST Major  FPC 1 misconfig
2011-03-24 00:52:31 PST Major  FPC 2 misconfig
2011-03-24 00:52:31 PST Major  FPC 3 misconfig
```

**detail**—(Optional) Display detailed status information for all FPCs or for the FPC in the specified slot (see *fpc-slot* or *slot*).

**all-members**—(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on all members of the Virtual Chassis configuration.

**interconnect-device name**—(QFabric systems only) (Optional) Display status information for all FPCs on the Interconnect device.

**fpc-slot**—(Optional) FPC slot number:

- (TX Matrix and TX Matrix Plus router only)—On a TX Matrix router, if you specify the number of the T640 router (line-card chassis) by using the **lcc number** option (the recommended method), replace *fpc-slot* with a value from 0 through 7. Otherwise, replace *fpc-slot* with a value from 0 through 31. Likewise, on a TX Matrix Plus router, if you specify the number of the specified router (line-card chassis) by using the **lcc number** option (the recommended method), replace *fpc-slot* with

a value from 0 through 7. Otherwise, replace *fpc-slot* with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis fpc detail 1 lcc 1
user@host> show chassis fpc detail 9
```

- M120 router—Replace *fpc-slot* with a value from 0 through 5.
- MX80 router—Replace *fpc-slot* with a value from 0 through 1.
- MX104 router—Replace *fpc-slot* with a value from 0 through 2.
- MX240 router—Replace *fpc-slot* with a value from 0 through 2.
- MX480 router—Replace *fpc-slot* with a value from 0 through 5.
- MX-960 router—Replace *fpc-slot* with a value from 0 through 11.
- MX2010 router—Replace *fpc-slot-number* with a value from 0 through 9.
- MX2020 router—Replace *fpc-slot-number* with a value from 0 through 19.
- Other routers—Replace *fpc-slot* with a value from 0 through 7.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—Replace *fpc-slot* with 0.
  - EX4200 switches in a Virtual Chassis configuration—Replace *fpc-slot* with a value from 0 through 9.
  - EX6210 switches—Replace *fpc-slot* with a value from 0 through 9.
  - EX8208 switches—Replace *fpc-slot* with a value from 0 through 7.
  - EX8216 switches—Replace *fpc-slot* with a value from 0 through 15.
- QFX Series:
  - QFX3500 switches—Replace *fpc-slot* with 0.
  - QFabric systems—Replace *fpc-slot* with 0 through 31 on the Interconnect device.
- PTX Series Packet Transport Routers:
  - PTX5000 Packet Transport Router—Replace *fpc-slot* with a value from 0 through 7.
- ACX Series Universal Access Routers:
  - ACX1000 and ACX2000 Universal Access Routers—Replace *fpc-slot* with 0.

**local**—(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on the local Virtual Chassis member.

**member *member-id***—(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**node-device name**—(QFabric systems only) (Optional) Display status information for each Node device. Each Node device is equivalent to an FPC.

**pic-status**—(Optional) Display status information for all PICs or for the PIC in the specified slot (see *fpc-slot*).



**NOTE:** On T1600 routers, Type 4 FPCs with ASICs based on the SL2.0 chipset do not support the 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (10x10GE [LAN/WAN] SFPP). If you issue the `show chassis fpc` command with the `pic-status` option, the CLI displays the string “Not Supported” for 10x10GE (LAN/WAN) SFPP PICs installed on such FPCs. The following is a sample output:

```
user@host> show chassis fpc pic-status
Slot 0  Online      E2-FPC Type 1
  PIC 0  Online      1x G/E SFP, 1000 BASE
  PIC 1  Online      Adaptive Services-II
  PIC 2  Online      1x G/E IQ, 1000 BASE
  PIC 3  Online      1x G/E IQ, 1000 BASE
Slot 1  Online      FPC Type 3-ES
  PIC 0  Present     UNUSED- Not Supported
Slot 2  Online      FPC Type 4-ES
  PIC 0  Offline     4x OC-192 SONET XFP
  PIC 1  Present     10x10GE(LAN/WAN) SFPP- Not Supported
<<<<<<
Slot 4  Offline     FPC Type 1-ES
Slot 5  Offline     FPC Type 2-ES
Slot 6  Online      E2-FPC Type 3
  PIC 0  Online      1x OC-192 SONET XFP
  PIC 1  Online      4x OC-48 SONET
  PIC 2  Online      4x OC-48 SONET
  PIC 3  Online      MultiServices 500
Slot 7  Online      FPC Type 4-ES
  PIC 0  Online      4x 10GE (LAN/WAN) XFP
  PIC 1  Online      4x 10GE (LAN/WAN) XFP
```

In addition, an entry is logged in the system log messages (/var/log/messages) that the PIC is not supported. The following is a sample message logged in the system log:

```
Apr  5 08:47:36 router1 chassisd[2770]: CHASSISD_UNSUPPORTED_PIC:
  PIC 1 in FPC 2 (type 763, version 257) is not supported
```

If you see this issue, contact Juniper Networks Technical Assistance Center (JTAC) for a possible fix. For more information about this issue and a possible solution, see [PSN-2010-03-696](https://www.juniper.net/psn/2010-03-696).



**NOTE:** When there is a double-bit ECC error in a network processor's memory, the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP or Channelized E1/T1 Circuit Emulation MIC is switched to the offline state.

```
user@host> show chassis fpc pic-status
Slot 1   Online           MPC Type 2 3D Q
PIC 0    Offline          1xC0C12/4xC0C3 CH-CE- ECC error detected
```

**lcc number**—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**Required Privilege Level** view

- Related Documentation**
- [request chassis fpc on page 882](#)
  - *show chassis fpc-feb-connectivity*
  - [show chassis fabric fpcs on page 1063](#)
  - *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
  - *MX960 Flexible PIC Concentrator Description*
  - *ACX2000 and ACX2100 Routers Hardware and CLI Terminology Mapping*
  - *enhanced-mode*

- List of Sample Output**
- [show chassis fpc \(EX6210 Switch\) on page 1169](#)
  - [show chassis fpc \(M10 Router\) on page 1169](#)
  - [show chassis fpc \(M20 Router\) on page 1169](#)
  - [show chassis fpc detail \(M Series Routers\) on page 1169](#)
  - [show chassis fpc detail \(MX80 Router\) on page 1170](#)
  - [show chassis fpc \(MX104 Router\) on page 1170](#)
  - [show chassis fpc detail \(MX104 Router\) on page 1170](#)
  - [show chassis fpc pic-status \(MX104 Router\) on page 1171](#)

[show chassis fpc \(MX240 Router\) on page 1171](#)  
[show chassis fpc \(EX Series Switch\) on page 1171](#)  
[show chassis fpc detail \(EX9200 Switch\) on page 1171](#)  
[show chassis fpc \(MX480 Router\) on page 1171](#)  
[show chassis fpc \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 1172](#)  
[show chassis fpc pic-status \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 1172](#)  
[show chassis fpc pic-status \(EX Series Switch\) on page 1172](#)  
[show chassis fpc \(MX480 Router with MPC4E\) on page 1172](#)  
[show chassis fpc detail \(MX480 Router with MPC4E\) on page 1173](#)  
[show chassis fpc \(MX480 Router with MPC4E\) on page 1173](#)  
[show chassis fpc detail \(MX480 Router with MPC4E\) on page 1173](#)  
[show chassis fpc \(MX960 Router\) on page 1174](#)  
[show chassis fpc \(MX960 Router with MPC5EQ\) on page 1174](#)  
[show chassis fpc detail \(MX960 Router with MPC5EQ\) on page 1174](#)  
[show chassis fpc pic-status \(MX960 Router with MPC5EQ\) on page 1176](#)  
[show chassis fpc \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 1177](#)  
[show chassis fpc \(MX240, MX480, MX960 with Application Services Modular Line Card\) on page 1177](#)  
[show chassis fpc \(MX2010 Routers\) on page 1177](#)  
[show chassis fpc \(MX2020 Routers\) on page 1177](#)  
[show chassis fpc \(MX2020 Router with MPC4E\) on page 1178](#)  
[show chassis fpc detail \(MX2020 Router with MPC4E\) on page 1178](#)  
[show chassis fpc \(MX2020 Router with MPC5EQ and MPC6E\) on page 1179](#)  
[show chassis fpc detail \(MX2020 Router with MPC5EQ and MPC6E\) on page 1179](#)  
[show chassis fpc pic-status \(MX2020 Router with MPC5EQ and MPC6E\) on page 1181](#)  
[show chassis fpc detail \(MX Series Routers\) on page 1182](#)  
[show chassis fpc detail \(EX Series Switches\) on page 1182](#)  
[show chassis fpc \(Hardware Not Supported\) on page 1182](#)  
[show chassis fpc detail \(Hardware Not Supported\) on page 1183](#)  
[show chassis fpc pic-status on page 1183](#)  
[show chassis fpc pic-status \(M Series Routers\) on page 1183](#)  
[show chassis fpc pic-status \(M120 Router\) on page 1184](#)  
[show chassis fpc pic-status \(MX240, MX480, and MX960 Routers with Application Services Modular Line Card\) on page 1184](#)  
[show chassis fpc lcc \(TX Matrix Router\) on page 1184](#)  
[show chassis fpc pic-status \(TX Matrix Router\) on page 1184](#)  
[show chassis fpc pic-status lcc \(TX Matrix Router\) on page 1185](#)  
[show chassis fpc \(TX Matrix Plus Router\) on page 1185](#)  
[show chassis fpc lcc \(TX Matrix Plus Router\) on page 1186](#)  
[show chassis fpc detail \(TX Matrix Plus Router\) on page 1186](#)  
[show chassis fpc pic-status \(TX Matrix Plus Router\) on page 1188](#)  
[show chassis fpc \(T1600 Router\) on page 1189](#)  
[show chassis fpc detail \(T1600 Router\) on page 1189](#)  
[show chassis fpc <fpc-slot> \(EX Series Switch\) on page 1190](#)  
[show chassis fpc slot \(T1600 Router\) on page 1190](#)  
[show chassis fpc pic-status \(T1600 Router\) on page 1190](#)  
[show chassis fpc \(T4000 Router\) on page 1191](#)  
[show chassis fpc detail \(T4000 Router\) on page 1191](#)

[show chassis fpc pic-status \(T4000 Router\) on page 1192](#)  
[show chassis fpc \(QFX Series\) on page 1192](#)  
[show chassis fpc detail \(QFX3500 Switches\) on page 1192](#)  
[show chassis fpc pic-status \(QFX3500 Switches\) on page 1192](#)  
[show chassis fpc interconnect-device \(QFabric System\) on page 1192](#)  
[show chassis fpc interconnect-device \(QFabric System\) on page 1193](#)  
[show chassis fpc interconnect-device detail \(QFabric System\) on page 1193](#)  
[show chassis fpc pic-status interconnect-device \(QFabric System\) on page 1193](#)  
[show chassis fpc pic-status node-device \(QFabric System\) on page 1194](#)  
[show chassis fpc \(PTX5000 Packet Transport Router\) on page 1194](#)  
[show chassis fpc detail \(PTX5000 Packet Transport Router\) on page 1194](#)  
[show chassis fpc pic-status \(PTX5000 Packet Transport Router\) on page 1195](#)  
[show chassis fpc \(ACX2000 Universal Access Router\) on page 1195](#)  
[show chassis fpc 0 \(ACX2000 Universal Access Router\) on page 1195](#)  
[show chassis fpc detail \(ACX2000 Universal Access Router\) on page 1195](#)  
[show chassis fpc pic-status \(ACX2000 Universal Access Router\) on page 1196](#)  
[show chassis FPC 1 \(MX Routers with Media Services Blade \[MSB\]\) on page 1196](#)  
[show chassis FPC 1 detail \(MX Routers with Media Services Blade \[MSB\]\) on page 1196](#)

**Output Fields** Table 116 on page 1167 lists the output fields for the **show chassis fpc** command. Output fields are listed in the approximate order in which they appear.

**Table 116: show chassis fpc Output Fields**

Field Name	Field Description	Level of Output
<b>Slot or Slot State</b>	<p>Slot number and state. The state can be one of the following conditions:</p> <ul style="list-style-type: none"> <li>• <b>Dead</b>—Held in reset because of errors.</li> <li>• <b>Diag</b>—Slot is being ignored while the FPC is running diagnostics.</li> <li>• <b>Dormant</b>—Held in reset.</li> <li>• <b>Empty</b>—No FPC is present.</li> <li>• <b>Offline</b>—(PTX Series Packet Transport Routers only) One of the following two states is displayed: <ul style="list-style-type: none"> <li>• <b>FPC offlined due to unreachable destinations</b></li> <li>• <b>FPC Offlined due to degraded FPC action</b></li> </ul> </li> <li>• <b>Online</b>—FPC is online and running.</li> <li>• <b>Present</b>—FPC is detected by the chassis daemon but either is not supported by the current version of Junos OS or is inserted in the wrong slot. The output also states either <b>Hardware Not Supported</b> or <b>Hardware Not In Right Slot</b>. The FPC is coming up but not yet online.</li> <li>• <b>Probed</b>—Probe is complete; awaiting restart of the Packet Forwarding Engine.</li> <li>• <b>Probe-wait</b>—Waiting to be probed.</li> </ul>	all levels
<b>Logical slot</b>	Slot number.	all levels
<b>Temp (C) or Temperature</b>	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.	all levels all levels

Table 116: show chassis fpc Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Temperature (PTX Series)</b>	On PTX Series Packet Transport Routers, temperature details are provided in degrees Celsius and Fahrenheit. Output includes: <ul style="list-style-type: none"> <li>• Temperature (PMB)—Temperature of the air passing by the Processor Mezzanine Board (PMB) at the bottom of the FPC.</li> <li>• Temperature (Intake)—Temperature of the air flowing into the chassis.</li> <li>• Temperature (Exhaust)—Exhaust temperatures for multiple zones (Exhaust A and Exhaust B).</li> <li>• Temperature (TLn)—Temperature of the specified Lookup ASIC (TL) of the packet forwarding engine on the FPC.</li> <li>• Temperature (TQn)—Temperature of the specified Queuing and Memory Interface ASIC (TQ) of the packet forwarding engine on the FPC.</li> </ul>	<b>detail</b>
<b>Total CPU Utilization (%)</b>	Total percentage of CPU being used by the FPC's processor.	all levels
<b>Interrupt CPU Utilization (%)</b>	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.	none specified
<b>Memory DRAM (MB)</b>	Total DRAM, in megabytes, available to the FPC's processor.	none specified
<b>Heap Utilization (%)</b>	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak).  <b>NOTE:</b> On MX Series routers and EX Series switches in a broadband edge environment, heap utilization levels higher than 70 percent can affect unified ISSU, router stability, or scaling capability.	none specified
<b>Buffer Utilization (%)</b>	Percentage of buffer space being used by the FPC's processor for buffering internal messages.	none specified
<b>Total CPU DRAM</b>	Amount of DRAM available to the FPC's CPU.	<b>detail</b>
<b>Total RLDRAM</b>	Amount of reduced latency dynamic random access memory (RLDRAM) available to the FPC CPU.	<b>detail</b>
<b>Total DDR DRAM</b>	Amount of double data rate dynamic random access memory (DDR DRAM) available to the FPC CPU.	<b>detail</b>
<b>Total SRAM</b>	Amount of static RAM (SRAM) used by the FPC's CPU.	<b>detail</b>
<b>Total SDRAM</b>	Total amount of memory used for storing packets and notifications.	<b>detail</b>
<b>I/O Manager ASICs information</b>	I/O Manager version number, manufacturer, and part number.	<b>detail</b>
<b>Start time</b>	Time when the Routing Engine detected that the FPC was running.	<b>detail</b>



Table 116: show chassis fpc Output Fields (*continued*)

Field Name	Field Description	Level of Output
Uptime	How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running.	detail
PIC type	(pic-status output only) Type of PIC.	none specified

## Sample Output

### show chassis fpc (EX6210 Switch)

```

user@switch> show chassis fpc

```

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Utilization (%) Buffer
0	Empty						
1	Online	7	5	0	1024	0	32
2	Empty						
3	Empty						
4	Online	25	17	2	2048	0	30
5	Online	25	3	0	2048	0	24
6	Online	6	5	0	1024	0	32
7	Empty						
8	Empty						
9	Online	8	7	0	1024	0	32

### show chassis fpc (M10 Router)

```

user@host> show chassis fpc
FPC status:

```

Slot	State	Temp (C)
0	Online	27
1	Online	28

### show chassis fpc (M20 Router)

```

user@host> show chassis fpc
FPC status:

```

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Utilization (%) Buffer
0	Empty	0	0	0	0	0	0
1	Online	38	0	0	8	0	4
2	Online	35	0	0	8	0	3
3	Empty	0	0	0	0	0	0

### show chassis fpc detail (M Series Routers)

```

user@host> show chassis fpc detail 1
Slot 1 information:
State Online
Temperature 48 degrees C
Total CPU DRAM 32 MB
Total SRAM 4 MB
Total SDRAM 256 MB
I/O Manager ASICs information Version 2.0, Foundry IBM, Part number 0
I/O Manager ASICs information Version 2.0, Foundry IBM, Part number 0

```

```
Start time      2000-02-08 02:18:49 UTC
Uptime         14 hours, 41 minutes, 41 seconds
```

#### show chassis fpc detail (MX80 Router)

```
user@host> show chassis fpc detail
Slot 0 information:
  State      Online
  Temperature 47 degrees C / 116 degrees F
  Total CPU DRAM 1024 MB
  Total SRAM    331 MB
  Total SDRAM   1280 MB
  Start time   2010-02-08 12:25:33 PST
  Uptime      2 hours, 13 minutes, 19 seconds
Slot 1 information:
  State      Online
  Temperature 47 degrees C / 116 degrees F
  Total CPU DRAM 1024 MB
  Total SRAM    331 MB
  Total SDRAM   1280 MB
  Start time   2010-02-08 12:25:33 PST
  Uptime      2 hours, 13 minutes, 19 seconds
```

#### show chassis fpc (MX104 Router)

```
user@host> show chassis fpc
Temp CPU Utilization (%) Memory Utilization (%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online 32 15 5 2048 22 13
1 Online 32 15 5 2048 22 13
2 Online 32 15 5 2048 22 13
```

#### show chassis fpc detail (MX104 Router)

```
user@host> show chassis fpc detail
Slot 0 information:
  State      Online
  Temperature 32 (C)
  Total CPU DRAM 2048 MB
  Total SRAM    403 MB
  Total SDRAM   1316 MB
  Start time   2013-05-23 14:39:18 IST
  Uptime      1 hour, 20 minutes, 22 seconds
Slot 1 information:
  State      Online
  Temperature 32 (C)
  Total CPU DRAM 2048 MB
  Total SRAM    403 MB
  Total SDRAM   1316 MB
  Start time   2013-05-23 14:39:18 IST
  Uptime      1 hour, 20 minutes, 22 seconds
Slot 2 information:
  State      Online
  Temperature 32 (C)
  Total CPU DRAM 2048 MB
  Total SRAM    403 MB
  Total SDRAM   1316 MB
  Start time   2013-05-23 14:39:18 IST
  Uptime      1 hour, 20 minutes, 22 seconds
```

**show chassis fpc pic-status (MX104 Router)**

```

user@host> show chassis fpc pic-status
Slot 0   Online
Slot 1   Online
  PIC 0   Online      10x 1GE(LAN) -E SFP
  PIC 1   Online      10x 1GE(LAN) -E SFP
Slot 2   Online
  PIC 0   Online      4x 10GE(LAN) SFP+

```

**show chassis fpc (MX240 Router)**

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total Interrupt	Heap	Buffer
0	Empty				
1	Online	34	6 0	1024 18	30
2	Online	33	9 0	1024 24	30

**show chassis fpc (EX Series Switch)**

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total Interrupt	Heap	Buffer
0	Empty				
1	Online	41	13 0	2048 19	14
2	Online	42	12 0	2048 19	14

**show chassis fpc detail (EX9200 Switch)**

```

user@switch> show chassis fpc detail
Slot 2 information:
  State                               Online
  Temperature                         37
  Total CPU DRAM                      2048 MB
  Total RLDRAM                        331 MB
  Total DDR DRAM                      1536 MB
  Start time:                         2014-03-12 15:35:28 UTC
  Uptime:                             1 hour, 4 minutes, 29 seconds
  Max Power Consumption               239 Watts
Slot 3 information:
  State                               Online
  Temperature                         39
  Total CPU DRAM                      2048 MB
  Total RLDRAM                        1036 MB
  Total DDR DRAM                      6656 MB
  Start time:                         2014-03-12 15:00:18 UTC
  Uptime:                             1 hour, 39 minutes, 39 seconds
  Max Power Consumption               520 Watts

```

**show chassis fpc (MX480 Router)**

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total Interrupt	Heap	Buffer
0	Empty				
1	Online	36	9 0	1024 17	57
2	Empty				
3	Empty				
4	Empty				
5	Empty				

**show chassis fpc (MX480 Router with 100-Gigabit Ethernet CFP)**

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Interrupt	Utilization (%)	DRAM (MB)	Heap	Buffer
0	Online	33	4	0		2048	10	13
1	Online	36	7	0		2048	16	13
2	Online	29	6	0		1024	27	29
3	Online	33	0	0		0	0	0
4	Online	36	7	0		2048	19	13
5	Online	34	31	11		2048	14	13

**show chassis fpc pic-status (MX480 Router with 100-Gigabit Ethernet CFP)**

```

user@host> show chassis fpc pic-status

```

Slot 1	Online	MPC Type 3
PIC 2	Online	1X100GE CFP
Slot 2	Online	DPCE 40x 1GE R EQ
PIC 0	Online	10x 1GE(LAN) EQ
PIC 1	Online	10x 1GE(LAN) EQ
PIC 2	Online	10x 1GE(LAN) EQ
PIC 3	Online	10x 1GE(LAN) EQ
Slot 3	Online	MPC Type 3
PIC 0	Online	1X100GE CFP
PIC 2	Online	1X100GE CFP
Slot 4	Online	MPC Type 3
PIC 0	Online	1X100GE CFP
PIC 2	Online	1X100GE CFP
Slot 5	Online	MPC Type 2 3D EQ
PIC 0	Online	2x 10GE XFP
PIC 1	Online	2x 10GE XFP
PIC 2	Online	10x 1GE(LAN) SFP
PIC 3	Online	10x 1GE(LAN) SFP

**show chassis fpc pic-status (EX Series Switch)**

```

user@host> show chassis fpc pic-status

```

Slot 1	Online	EX9200 32x10G SFP
PIC 0	Online	8X10GE SFPP
PIC 1	Online	8X10GE SFPP
PIC 2	Online	8X10GE SFPP
PIC 3	Online	8X10GE SFPP
Slot 2	Online	EX9200 32x10G SFP
PIC 0	Online	8X10GE SFPP
PIC 1	Online	8X10GE SFPP
PIC 2	Online	8X10GE SFPP
PIC 3	Online	8X10GE SFPP

**show chassis fpc (MX480 Router with MPC4E)**

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Interrupt	Utilization (%)	DRAM (MB)	Heap	Buffer
0	Empty							
1	Empty							
2	Online		38	7	0	2048	19	14
3	Online		39	8	0	2048	18	14
4	Online		39	7	0	2048	17	14
5	Empty							

## show chassis fpc detail (MX480 Router with MPC4E)

```

user@host> show chassis fpc detail
Slot 2 information:
  State                Online
  Temperature           38
  Total CPU DRAM        2048 MB
  Total RLDRAM          1036 MB
  Total DDR DRAM        11264 MB
  Start time:           2013-02-18 05:06:57 PST
  Uptime:               17 hours, 41 minutes, 9 seconds
  Max Power Consumption 610 Watts
Slot 3 information:
  State                Online
  Temperature           38
  Total CPU DRAM        2048 MB
  Total RLDRAM          1036 MB
  Total DDR DRAM        11264 MB
  Start time:           2013-02-18 05:07:00 PST
  Uptime:               17 hours, 41 minutes, 6 seconds
  Max Power Consumption 610 Watts
Slot 4 information:
  State                Diagnostics
  Temperature           37
  Total CPU DRAM        0 MB
  Total RLDRAM          0 MB
  Total DDR DRAM        0 MB
  Max Power Consumption 520 Watts

```

## show chassis fpc (MX480 Router with MPC4E)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)	DRAM (MB)	Heap	Buffer
0	Empty						
1	Empty						
2	Online	38	7	0	2048	19	14
3	Online	39	8	0	2048	18	14
4	Online	39	7	0	2048	17	14
5	Empty						

## show chassis fpc detail (MX480 Router with MPC4E)

```

user@host> show chassis fpc detail
Slot 2 information:
  State                Online
  Temperature           38
  Total CPU DRAM        2048 MB
  Total RLDRAM          1036 MB
  Total DDR DRAM        11264 MB
  Start time:           2013-02-18 05:06:57 PST
  Uptime:               17 hours, 41 minutes, 9 seconds
  Max Power Consumption 610 Watts
Slot 3 information:
  State                Online
  Temperature           38
  Total CPU DRAM        2048 MB
  Total RLDRAM          1036 MB
  Total DDR DRAM        11264 MB
  Start time:           2013-02-18 05:07:00 PST
  Uptime:               17 hours, 41 minutes, 6 seconds

```

```

Max Power Consumption          610 Watts
Slot 4 information:
State                          Diagnostics
Temperature                    37
Total CPU DRAM                 0 MB
Total RLD RAM                  0 MB
Total DDR DRAM                 0 MB
Max Power Consumption          520 Watts

```

#### show chassis fpc (MX960 Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	Total	CPU Utilization (%) Interrupt	Memory DRAM (MB)	Heap	Utilization (%) Buffer
0	Empty						
1	Empty						
2	Empty						
3	Online	25	19	0	1024	15	57
4	Empty						
5	Online	26	27	0	1024	15	57
6	Empty						
7	Empty						
8	Empty						
9	Empty						
10	Empty						
11	Empty						

#### show chassis fpc (MX960 Router with MPC5EQ)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	Total	CPU Utilization (%) Interrupt	Memory DRAM (MB)	Heap	Utilization (%) Buffer
0	Online	38	16	0	3584	7	13
1	Online	31	15	0	2048	17	13
2	Empty						
3	Online	31	14	0	2048	20	13
4	Online	34	16	0	3584	7	13
5	Online	34	16	0	3584	7	13
6	Empty						
7	Online	32	9	0	2048	18	14
8	Online	36	19	0	3584	7	13
9	Online	31	9	0	2048	13	13
10	Online	35	14	0	3584	7	13
11	Online	33	11	0	2048	18	14

#### show chassis fpc detail (MX960 Router with MPC5EQ)

```

user@host> show chassis fpc detail
Slot 0 information:
State                          Online
Temperature                    38
Total CPU DRAM                 3584 MB
Total XR2                      291 MB
Total DDR DRAM                 24960 MB
Start time:                    2014-04-22 10:01:46 PDT
Uptime:                        1 hour, 23 minutes, 40 seconds
Max Power Consumption          607 Watts
Slot 1 information:
State                          Online
Temperature                    31
Total CPU DRAM                 2048 MB
Total RLD RAM                  1036 MB

```

```

Total DDR DRAM                6656 MB
Start time:                    2014-04-22 10:01:50 PDT
Uptime:                        1 hour, 23 minutes, 36 seconds
Max Power Consumption          520 Watts
Slot 3 information:
State                          Online
Temperature                     31
Total CPU DRAM                 2048 MB
Total RLD RAM                  1324 MB
Total DDR DRAM                 5120 MB
Start time:                    2014-04-22 10:01:50 PDT
Uptime:                        1 hour, 23 minutes, 36 seconds
Max Power Consumption          440 Watts
Slot 4 information:
State                          Online
Temperature                     34
Total CPU DRAM                 3584 MB
Total XR2                      291 MB
Total DDR DRAM                 24960 MB
Start time:                    2014-04-22 10:01:54 PDT
Uptime:                        1 hour, 23 minutes, 32 seconds
Max Power Consumption          607 Watts
Slot 5 information:
State                          Online
Temperature                     34
Total CPU DRAM                 3584 MB
Total XR2                      291 MB
Total DDR DRAM                 24960 MB
Start time:                    2014-04-22 10:01:56 PDT
Uptime:                        1 hour, 23 minutes, 30 seconds
Max Power Consumption          607 Watts
Slot 7 information:
State                          Online
Temperature                     32
Total CPU DRAM                 2048 MB
Total RLD RAM                  1036 MB
Total DDR DRAM                 11264 MB
Start time:                    2014-04-22 10:02:02 PDT
Uptime:                        1 hour, 23 minutes, 24 seconds
Max Power Consumption          608 Watts
Slot 8 information:
State                          Online
Temperature                     36
Total CPU DRAM                 3584 MB
Total XR2                      291 MB
Total DDR DRAM                 24960 MB
Start time:                    2014-04-22 10:02:07 PDT
Uptime:                        1 hour, 23 minutes, 19 seconds
Max Power Consumption          607 Watts
Slot 9 information:
State                          Online
Temperature                     31
Total CPU DRAM                 2048 MB
Total RLD RAM                  734 MB
Total DDR DRAM                 3108 MB
Start time:                    2014-04-22 10:02:05 PDT
Uptime:                        1 hour, 23 minutes, 21 seconds
Max Power Consumption          368 Watts
Slot 10 information:
State                          Online
Temperature                     35

```

```

Total CPU DRAM          3584 MB
Total XR2                291 MB
Total DDR DRAM          24960 MB
Start time:              2014-04-22 10:02:11 PDT
Uptime:                  1 hour, 23 minutes, 15 seconds
Max Power Consumption    607 Watts
Slot 11 information:
State                    Online
Temperature              33
Total CPU DRAM          2048 MB
Total RLDRAM             1036 MB
Total DDR DRAM          11264 MB
Start time:              2014-04-22 10:02:16 PDT
Uptime:                  1 hour, 23 minutes, 10 seconds
Max Power Consumption    608 Watts

```

### show chassis fpc pic-status(MX960 Router with MPC5EQ)

```

user@host> show chassis fpc pic-status
Slot 0  Online      MPC5E 3D Q 2CGE+4XGE
PIC 0   Online      2X10GE SFPP OTN
PIC 1   Online      1X100GE CFP2 OTN
PIC 2   Online      2X10GE SFPP OTN
PIC 3   Online      1X100GE CFP2 OTN
Slot 1  Online      MPCE Type 3 3D
PIC 0   Online      10X10GE SFPP
PIC 2   Online      1X100GE CXP
Slot 3  Online      MPC 3D 16x 10GE
PIC 0   Online      4x 10GE(LAN) SFP+
PIC 1   Online      4x 10GE(LAN) SFP+
PIC 2   Online      4x 10GE(LAN) SFP+
PIC 3   Online      4x 10GE(LAN) SFP+
Slot 4  Online      MPC5E 3D Q 2CGE+4XGE
PIC 0   Online      2X10GE SFPP OTN
PIC 1   Online      1X100GE CFP2 OTN
PIC 2   Online      2X10GE SFPP OTN
PIC 3   Online      1X100GE CFP2 OTN
Slot 5  Online      MPC5E 3D Q 2CGE+4XGE
PIC 0   Online      2X10GE SFPP OTN
PIC 1   Online      1X100GE CFP2 OTN
PIC 2   Online      2X10GE SFPP OTN
PIC 3   Online      1X100GE CFP2 OTN
Slot 7  Online      MPC4E 3D 2CGE+8XGE
PIC 0   Online      4x10GE SFPP
PIC 1   Online      1X100GE CFP
PIC 2   Online      4x10GE SFPP
PIC 3   Online      1X100GE CFP
Slot 8  Online      MPC5E 3D Q 24XGE+6XLGE
PIC 0   Offline     12X10GE SFPP OTN
PIC 1   Offline     12X10GE SFPP OTN
PIC 2   Online      3X40GE QSFPP
PIC 3   Online      3X40GE QSFPP
Slot 9  Online      MPCE Type 2 3D P
PIC 0   Online      2x 10GE XFP
PIC 1   Online      2x 10GE XFP
Slot 10 Online      MPC5E 3D Q 24XGE+6XLGE
PIC 0   Online      12X10GE SFPP
PIC 1   Online      12X10GE SFPP
PIC 2   Offline     3X40GE QSFPP
PIC 3   Offline     3X40GE QSFPP

```



```

Slot 11 Online      MPC4E 3D 2CGE+8XGE
PIC 0 Online        4x10GE SFPP
PIC 1 Online        1X100GE CFP
PIC 2 Online        4x10GE SFPP
PIC 3 Online        1X100GE CFP

```

### show chassis fpc (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host> show chassis fpc 1
      Temp CPU Utilization (%) Memory      Utilization (%)
Slot State      (C) Total  Interrupt      DRAM (MB) Heap      Buffer
1 Online          34      5          0      3072      5      13

```

### show chassis fpc (MX240, MX480, MX960 with Application Services Modular Line Card)

```

user@host> show chassis fpc 1 detail
Slot 1 information:
State                               Online
Temperature                          34
Total CPU DRAM                      3072 MB
Total RLDRAM                        259 MB
Total DDR DRAM                      4864 MB
Start time:                         2012-06-19 10:51:43 PDT
Uptime:                             16 minutes, 48 seconds
Max Power Consumption               550 Watts

```

### show chassis fpc (MX2010 Routers)

```

user@host> show chassis fpc
      Temp CPU Utilization (%) Memory      Utilization (%)
Slot State      (C) Total  Interrupt      DRAM (MB) Heap      Buffer
0 Online          34      9          0      2048      18      13
1 Online          32      9          0      2048      15      13
2 Empty
3 Empty
4 Empty
5 Empty
6 Empty
7 Empty
8 Online          31     13          0      2048      11      13
9 Online          33     10          0      2048      18      13

```

### show chassis fpc (MX2020 Routers)

```

user@host> show chassis fpc
      Temp CPU Utilization (%) Memory      Utilization (%)
Slot State      (C) Total  Interrupt      DRAM (MB) Heap      Buffer
0 Online          10     12          0      2048      18      13
1 Online           8      9          0      2048      18      13
2 Online           7      9          0      2048      18      13
3 Online           8     10          0      2048      18      13
4 Online           9     10          0      2048      18      13
5 Online           8      9          0      2048      18      13
6 Online           8     10          0      2048      18      13
7 Online           9      9          0      2048      18      13
8 Online           9     10          0      2048      18      13
9 Online          10      9          0      2048      18      13
10 Online         16      8          0      2048      18      13
11 Online         11     10          0      2048      18      13
12 Online         10     10          0      2048      18      13
13 Online         11      9          0      2048      18      13

```

14	Online	12	10	0	2048	18	13
15	Online	13	9	0	2048	18	13
16	Online	13	9	0	2048	18	13
17	Online	12	9	0	2048	18	13
18	Online	12	8	0	2048	18	13
19	Online	14	10	0	2048	18	13

#### show chassis fpc (MX2020 Router with MPC4E)

```

user@host> show chassis fpc

```

Slot	Temp	CPU Utilization (%)	Memory	Utilization (%)	DRAM (MB)	Heap	Buffer
	State	(C)	Total	Interrupt			
0	Online	33	12	2	2048	11	13
1	Empty						
2	Empty						
3	Empty						
4	Empty						
5	Empty						
6	Empty						
7	Empty						
8	Empty						
9	Online	31	10	0	2048	11	13
10	Online	32	7	0	2048	14	13
11	Empty						
12	Empty						
13	Empty						
14	Online	28	12	0	2048	15	14
15	Empty						
16	Empty						
17	Empty						
18	Empty						
19	Online	38	8	0	2048	18	13

#### show chassis fpc detail (MX2020 Router with MPC4E)

```

user@host> show chassis fpc detail

```

Slot 0 information:

State	Online
Temperature	34
Total CPU DRAM	2048 MB
Total RDRAM	806 MB
Total DDR DRAM	2632 MB
Start time:	2013-02-17 08:17:35 PST
Uptime:	1 day, 14 hours, 50 minutes, 39 seconds
Max Power Consumption	368 Watts

Slot 9 information:

State	Online
Temperature	32
Total CPU DRAM	2048 MB
Total RDRAM	806 MB
Total DDR DRAM	2632 MB
Start time:	2013-02-17 08:17:43 PST
Uptime:	1 day, 14 hours, 50 minutes, 31 seconds
Max Power Consumption	368 Watts

Slot 10 information:

State	Online
Temperature	37
Total CPU DRAM	2048 MB
Total RDRAM	1036 MB
Total DDR DRAM	6656 MB
Start time:	2013-02-17 08:17:54 PST

```

Uptime: 1 day, 14 hours, 50 minutes, 20 seconds
Max Power Consumption 520 Watts
Slot 14 information:
State Online
Temperature 32
Total CPU DRAM 2048 MB
Total RLDRAM 1036 MB
Total DDR DRAM 11264 MB
Start time: 2013-02-17 08:18:01 PST
Uptime: 1 day, 14 hours, 50 minutes, 13 seconds
Max Power Consumption 610 Watts
Slot 19 information:
State Online
Temperature 38
Total CPU DRAM 2048 MB
Total RLDRAM 1324 MB
Total DDR DRAM 5120 MB
Start time: 2013-02-17 08:18:08 PST
Uptime: 1 day, 14 hours, 50 minutes, 6 seconds
Max Power Consumption 440 Watts

```

#### show chassis fpc (MX2020 Router with MPC5EQ and MPC6E)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)
			Total Interrupt	DRAM (MB) Heap Buffer
0	Online	31	20 0	3584 7 13
1	Online	28	19 0	2048 17 13
2	Online	27	10 0	2048 18 14
3	Online	26	10 0	2048 13 13
4	Online	29	19 0	3584 7 13
5	Online	28	68 0	2048 20 13
6	Empty			
7	Empty			
8	Empty			
9	Online	36	19 0	3584 10 13
10	Online	37	26 0	3584 10 13
11	Empty			
12	Empty			
13	Empty			
14	Empty			
15	Empty			
16	Empty			
17	Online	28	43 0	3584 10 13
18	Online	29	19 0	3584 7 13
19	Online	31	19 0	3584 7 13

#### show chassis fpc detail (MX2020 Router with MPCEQ and MPC6E)

```

user@host> show chassis fpc detail
Slot 0 information:
State Online
Temperature 31
Total CPU DRAM 3584 MB
Total XR2 291 MB
Total DDR DRAM 24960 MB
Start time: 2014-04-22 23:33:19 PDT
Uptime: 6 minutes, 24 seconds
Max Power Consumption 607 Watts
Slot 1 information:

```

State	Online
Temperature	28
Total CPU DRAM	2048 MB
Total RLDRAM	1036 MB
Total DDR DRAM	6656 MB
Start time:	2014-04-22 23:33:24 PDT
Uptime:	6 minutes, 19 seconds
Max Power Consumption	520 Watts
Slot 2 information:	
State	Online
Temperature	27
Total CPU DRAM	2048 MB
Total RLDRAM	1036 MB
Total DDR DRAM	11264 MB
Start time:	2014-04-22 23:33:34 PDT
Uptime:	6 minutes, 9 seconds
Max Power Consumption	608 Watts
Slot 3 information:	
State	Online
Temperature	26
Total CPU DRAM	2048 MB
Total RLDRAM	734 MB
Total DDR DRAM	3108 MB
Start time:	2014-04-22 23:33:39 PDT
Uptime:	6 minutes, 4 seconds
Max Power Consumption	368 Watts
Slot 4 information:	
State	Online
Temperature	29
Total CPU DRAM	3584 MB
Total XR2	291 MB
Total DDR DRAM	24960 MB
Start time:	2014-04-22 23:33:51 PDT
Uptime:	5 minutes, 52 seconds
Max Power Consumption	607 Watts
Slot 5 information:	
State	Online
Temperature	28
Total CPU DRAM	2048 MB
Total RLDRAM	1324 MB
Total DDR DRAM	5120 MB
Start time:	2014-04-22 23:33:57 PDT
Uptime:	5 minutes, 46 seconds
Max Power Consumption	440 Watts
Slot 9 information:	
State	Online
Temperature	25
Total CPU DRAM	3584 MB
Total XR2	518 MB
Total DDR DRAM	49920 MB
Start time:	2014-04-22 23:31:20 PDT
Uptime:	8 minutes, 23 seconds
Max Power Consumption	1130 Watts
Slot 10 information:	
State	Online
Temperature	32
Total CPU DRAM	3584 MB
Total XR2	518 MB
Total DDR DRAM	49920 MB
Start time:	2014-04-22 23:31:25 PDT
Uptime:	8 minutes, 18 seconds

```

Max Power Consumption          1130 Watts
Slot 17 information:
State                          Online
Temperature                    25
Total CPU DRAM                 3584 MB
Total XR2                      518 MB
Total DDR DRAM                 49920 MB
Start time:                    2014-04-22 23:31:29 PDT
Uptime:                        8 minutes, 14 seconds
Max Power Consumption          1130 Watts
Slot 18 information:
State                          Online
Temperature                    29
Total CPU DRAM                 3584 MB
Total XR2                      291 MB
Total DDR DRAM                 24960 MB
Start time:                    2014-04-22 23:34:11 PDT
Uptime:                        5 minutes, 32 seconds
Max Power Consumption          607 Watts
Slot 19 information:
State                          Online
Temperature                    32
Total CPU DRAM                 3584 MB
Total XR2                      291 MB
Total DDR DRAM                 24960 MB
Start time:                    2014-04-22 23:34:20 PDT
Uptime:                        5 minutes, 23 seconds
Max Power Consumption          607 Watts

```

#### show chassis fpc pic-status (MX2020 Router with MPC5EQ and MPC6E)

```

user@host> show chassis fpc pic-status
Slot 0  Online      MPC5E 3D Q 24XGE+6XLGE
PIC 0   Online      12X10GE SFPP OTN
PIC 1   Online      12X10GE SFPP OTN
PIC 2   Offline     3X40GE QSFPP
PIC 3   Offline     3X40GE QSFPP
Slot 1  Online      MPCE Type 3 3D
PIC 0   Online      10X10GE SFPP
PIC 2   Online      1X100GE CXP
Slot 2  Online      MPC4E 3D 2CGE+8XGE
PIC 0   Online      4x10GE SFPP
PIC 1   Online      1X100GE CFP
PIC 2   Online      4x10GE SFPP
PIC 3   Online      1X100GE CFP
Slot 3  Online      MPCE Type 2 3D P
PIC 0   Online      2x 10GE XFP
PIC 1   Online      2x 10GE XFP
Slot 4  Online      MPC5E 3D Q 2CGE+4XGE
PIC 0   Online      2X10GE SFPP OTN
PIC 1   Online      1X100GE CFP2 OTN
PIC 2   Online      2X10GE SFPP OTN
PIC 3   Online      1X100GE CFP2 OTN
Slot 5  Online      MPC 3D 16x 10GE
PIC 0   Online      4x 10GE(LAN) SFP+
PIC 1   Online      4x 10GE(LAN) SFP+
PIC 2   Online      4x 10GE(LAN) SFP+
PIC 3   Online      4x 10GE(LAN) SFP+
Slot 9  Online      MPC6E 3D
PIC 0   Online      2X100GE CFP2 OTN
PIC 1   Online      2X100GE CFP2 OTN

```

```

Slot 10 Online MPC6E 3D
PIC 0 Online 24X10GE SFPP OTN
PIC 1 Online 4X100GE CXP
Slot 17 Online MPC6E 3D
PIC 0 Online 24X10GE SFPP
PIC 1 Online 4X100GE CXP
Slot 18 Online MPC5E 3D Q 24XGE+6XLGE
PIC 0 Offline 12X10GE SFPP OTN
PIC 1 Offline 12X10GE SFPP OTN
PIC 2 Online 3X40GE QSFPP
PIC 3 Online 3X40GE QSFPP
Slot 19 Online MPC5E 3D Q 24XGE+6XLGE
PIC 0 Online 12X10GE SFPP OTN
PIC 1 Offline 12X10GE SFPP OTN
PIC 2 Offline 3X40GE QSFPP
PIC 3 Online 3X40GE QSFPP

```

### show chassis fpc detail (MX Series Routers)

```

user@host> show chassis fpc detail 2
Slot 0 information:
State Online
Temperature 36 degrees C / 96 degrees F
Total CPU DRAM 1024 MB
Total RDRAM 256 MB
Total DDR DRAM 4096 MB
Start time: 2009-08-11 21:20:30 PDT
Uptime: 2 hours, 8 minutes, 50 seconds
Max Power Consumption 335 Watts

```

### show chassis fpc detail (EX Series Switches)

```

user@host> show chassis fpc detail 2
Slot 1 information:
State Online
Temperature 41
Total CPU DRAM 2048 MB
Total RDRAM 1036 MB
Total DDR DRAM 11264 MB
Start time: 2013-04-02 00:04:52 PDT
Uptime: 7 days, 9 hours, 47 minutes, 46 seconds
Max Power Consumption 610 Watts
Slot 2 information:
State Online
Temperature 41
Total CPU DRAM 2048 MB
Total RDRAM 1036 MB
Total DDR DRAM 11264 MB
Start time: 2013-04-02 00:04:56 PDT
Uptime: 7 days, 9 hours, 47 minutes, 42 seconds
Max Power Consumption 610 Watts

```

### show chassis fpc (Hardware Not Supported)

```

user@host> show chassis fpc
show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)
			Total Interrupt	DRAM (MB) Heap Buffer
0	Online	-----	CPU less FPC	-----
1	Present	-----	Hardware Not In Right Slot	-----
2	Online	0	0	0 0 0
3	Present	-----	Hardware Not Supported	-----

```

4 Empty
5 Empty
6 Online          0          0          0          0          0

```

### show chassis fpc detail (Hardware Not Supported)

```

user@host> show chassis fpc detail
Slot 0 information:
  State          Online
  Total CPU DRAM ----- CPU less FPC -----
  Start time     2006-07-07 03:21:00 UTC
  Uptime         27 minutes, 51 seconds
Slot 1 information:
  State          Present
  Reason         --- Hardware Not In Right Slot ---
Slot 2 information:
  State          Online
  Total CPU DRAM 32 MB
  Start time     2006-07-07 03:20:59 UTC
  Uptime         27 minutes, 52 seconds
Slot 3 information:
  State          Present
  Reason         --- Hardware Not Supported ---
  Total CPU DRAM 0 MB
Slot 6 information:
  State          Online
  Total CPU DRAM 32 MB
  Start time     2006-07-07 03:21:01 UTC
  Uptime         27 minutes, 50 seconds

```

### show chassis fpc pic-status

```

user@host> show chassis fpc pic-status
Slot 0 Online
  PIC 1  1x OC-12 ATM, MM
  PIC 2  1x OC-12 ATM, MM
  PIC 3  1x OC-12 ATM, MM
Slot 1 Online
  PIC 0  1x OC-48 SONET, SMIR
Slot 2 Online
  PIC 0  1x OC-192 SONET, SMSR

```

### show chassis fpc pic-status (M Series Routers)

```

user@host> show chassis fpc pic-status
Slot 1 Online      FPC Type 1
  PIC 0 Present    2x OC-3 ATM, MM- Hardware Error
  PIC 1 Online     4x OC-3 SONET, SMIR
Slot 2 Online      E-FPC Type 2
  PIC 0 Online     4x G/E, 1000 BASE-SX
  PIC 1 Online     2x G/E SFP, 1000 BASE
  PIC 3 Online     1x Tunnel
Slot 3 Online      E-FPC Type 1
  PIC 0 Online     1x G/E IQ, 1000 BASE
  PIC 2 Online     1x G/E SFP, 1000 BASE
Slot 4 Online      E-FPC Type 2
  PIC 0 Online     4x G/E SFP, 1000 BASE
  PIC 1 Online     4x G/E SFP, 1000 BASE
  PIC 2 Online     4x G/E SFP, 1000 BASE
  PIC 3 Online     4x G/E SFP, 1000 BASE

```

```
Slot 5   Online       FPC Type 2
...
```

### show chassis fpc pic-status (M120 Router)

```
user@host> show chassis fpc pic-status
Slot 1   Online       M120 CFPC 10GE
  PIC 0   Online       1x 10GE(LAN/WAN) XFP
Slot 3   Online       M120 FPC Type 2 (proto)
  PIC 0   Online       2x G/E IQ, 1000 BASE
  PIC 1   Online       4x OC-3 SONET, SMIR
  PIC 2   Online       2x G/E IQ, 1000 BASE
  PIC 3   Online       8x 1GE(LAN), IQ2
Slot 4   Online       M120 FPC Type 3 (proto)
  PIC 0   Online       10x 1GE(LAN), 1000 BASE
Slot 5   Online       M120 FPC Type 1 (proto)
  PIC 0   Present      1x G/E, 1000 BASE-LX- Not Supported
  PIC 1   Online       1x CHOC3 IQ SONET, SMLR
  PIC 2   Online       4x CHDS3 IQ
  PIC 3   Online       1x G/E SFP, 1000 BASE
```

### show chassis fpc pic-status (MX240, MX480, and MX960 Routers with Application Services Modular Line Card)

In the following output **Slot 1** and **Slot 5** are the Application Services Modular Carrier Cards (AS MCC), **PIC 0** is the Application Services Modular Storage Card (AS MSC), and **PIC 2** is the Application Services Modular Processing Card (AS MXC).

```
user@host> show chassis fpc pic-status
Slot 2   Online       MPC Type 1 3D Q
  Slot 1   Online       AS-MCC
  PIC 0   Online       AS-MSC
  PIC 2   Online       AS-MXC
Slot 4   Offline      MPC 3D 16x 10GE
Slot 5   Offline      AS-MCC
```

### show chassis fpc lcc (TX Matrix Router)

```
user@host> show chassis fpc lcc 0
lcc0-re0:
-----
Slot State      Temp CPU      Utilization (%)  Memory  Utilization (%)
      (C) Total Interrupt    DRAM (MB)   Heap    Buffer
0 Empty
1 Online        27    2         0        256      8        44
2 Online        27    3         0        256     15        44
3 Empty
4 Empty
5 Empty
6 Empty
7 Empty
```

### show chassis fpc pic-status (TX Matrix Router)

```
user@host> show chassis fpc pic-status
lcc0-re0:
-----
Slot 0   Online       FPC Type 3
  PIC 0   Online       1x OC-192 SM SR1
  PIC 1   Online       1x OC-192 SM SR2
  PIC 2   Online       1x OC-192 SM SR1
  PIC 3   Online       1x Tunnel
```



```

Slot 1  Online      FPC Type 2
PIC 0   Online      1x OC-48 SONET, SMSR
PIC 1   Online      1x OC-48 SONET, SMSR

```

```
lcc1-re0:
```

```
lcc2-re0:
```

```

Slot 1  Online      FPC Type 3
PIC 0   Online      1x OC-192 SM SR1
Slot 5  Online      FPC Type 2
PIC 0   Online      1x OC-48 SONET, SMSR
PIC 1   Online      2x G/E, 1000 BASE-LX
PIC 2   Online      2x G/E, 1000 BASE-LX
PIC 3   Online      1x OC-48 SONET, SMSR

```

```
lcc3-re0:
```

#### show chassis fpc pic-status lcc (TX Matrix Router)

```
user@host> show chassis fpc pic-status lcc 0
```

```
lcc0-re0:
```

```

Slot 0  Online      FPC Type 3
PIC 0   Online      1x OC-192 SM SR2
Slot 1  Online      FPC Type 2
PIC 0   Online      2x OC-12 ATM2 IQ, MM
PIC 1   Online      1x OC-48 SONET, SMSR
PIC 2   Online      1x OC-48 SONET, SMSR
PIC 3   Online      4x G/E, 1000 BASE-SX

```

#### show chassis fpc (TX Matrix Plus Router)

```
user@host> show chassis fpc
```

```
lcc0-re0:
```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Buffer
0	Empty					
1	Online	38	4 0	2048	3	24
2	Online	43	8 0	2048	6	24
3	Empty					
4	Online	43	6 0	2048	6	24
5	Empty					
6	Online	42	13 0	2048	6	24
7	Online	45	7 0	2048	3	24

```
lcc2-re0:
```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Buffer
0	Online	42	10 0	2048	6	24
1	Empty					
2	Online	42	11 0	2048	6	24
3	Online	40	5 0	2048	3	24
4	Online	33	26 0	1024	8	49
5	Empty					
6	Online	43	8 0	2048	6	24
7	Online	46	6 0	2048	3	24

lcc3-re0:

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Utilization (%) Buffer
0	Empty						
1	Empty						
2	Online	39	30	0	2048	7	24
3	Empty						
4	Online	41	8	0	2048	6	24
5	Online	41	12	0	2048	6	24
6	Online	40	8	0	2048	6	24
7	Online	42	4	0	2048	3	24

**show chassis fpc lcc (TX Matrix Plus Router)**

user@host&gt; show chassis fpc lcc 0

lcc0-re0:

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Utilization (%) Buffer
0	Empty						
1	Online	38	4	0	2048	3	24
2	Online	43	8	0	2048	6	24
3	Empty						
4	Online	43	6	0	2048	6	24
5	Empty						
6	Online	42	14	0	2048	6	24
7	Online	45	6	0	2048	3	24

**show chassis fpc detail (TX Matrix Plus Router)**

user@host&gt; show chassis fpc details

lcc0-re0:

Slot 1 information:

```

State                               Online
Temperature                         38 degrees C / 100 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          64 MB
Total SDRAM                         1280 MB
Start time                          2010-10-04 20:06:22 PDT
Uptime                              1 hour, 32 minutes, 51 seconds

```

Slot 2 information:

```

State                               Online
Temperature                         43 degrees C / 109 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          128 MB
Total SDRAM                         2560 MB
Start time                          2010-10-04 20:06:37 PDT
Uptime                              1 hour, 32 minutes, 36 seconds

```

Slot 4 information:

```

State                               Online
Temperature                         43 degrees C / 109 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          128 MB
Total SDRAM                         2560 MB
Start time                          2010-10-04 20:06:40 PDT
Uptime                              1 hour, 32 minutes, 33 seconds

```

Slot 6 information:

```

State                               Online
Temperature                         42 degrees C / 107 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          128 MB
Total SDRAM                         2560 MB
Start time                         2010-10-04 20:06:42 PDT
Uptime                             1 hour, 32 minutes, 31 seconds

```

## Slot 7 information:

```

State                               Online
Temperature                         45 degrees C / 113 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          64 MB
Total SDRAM                         1280 MB
Start time                         2010-10-04 20:06:43 PDT
Uptime                             1 hour, 32 minutes, 30 seconds

```

## lcc2-re0:

-----  
Slot 0 information:

```

State                               Online
Temperature                         42 degrees C / 107 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          128 MB
Total SDRAM                         2560 MB
Start time                         2010-10-04 20:06:35 PDT
Uptime                             1 hour, 32 minutes, 38 seconds

```

## Slot 2 information:

```

State                               Online
Temperature                         42 degrees C / 107 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          128 MB
Total SDRAM                         2560 MB
Start time                         2010-10-04 20:06:37 PDT
Uptime                             1 hour, 32 minutes, 36 seconds

```

## Slot 3 information:

```

State                               Online
Temperature                         40 degrees C / 104 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          64 MB
Total SDRAM                         1280 MB
Start time                         2010-10-04 20:06:28 PDT
Uptime                             1 hour, 32 minutes, 45 seconds

```

## Slot 4 information:

```

State                               Online
Temperature                         33 degrees C / 91 degrees F
Total CPU DRAM                      1024 MB
Total SRAM                          64 MB
Total SDRAM                         1280 MB
Start time                         2010-10-04 20:08:03 PDT
Uptime                             1 hour, 31 minutes, 10 seconds

```

## Slot 6 information:

```

State                               Online
Temperature                         43 degrees C / 109 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          128 MB
Total SDRAM                         2560 MB
Start time                         2010-10-04 20:06:44 PDT
Uptime                             1 hour, 32 minutes, 29 seconds

```

## Slot 7 information:

```

State                               Online
Temperature                         46 degrees C / 114 degrees F

```

```

Total CPU DRAM          2048 MB
Total SRAM              64 MB
Total SDRAM            1280 MB
Start time              2010-10-04 20:06:46 PDT
Uptime                  1 hour, 32 minutes, 27 seconds

lcc3-re0:
-----
Slot 2 information:
State                  Online
Temperature            38 degrees C / 100 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:17:31 PDT
Uptime                1 hour, 21 minutes, 42 seconds
Slot 4 information:
State                  Online
Temperature            41 degrees C / 105 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:17:34 PDT
Uptime                1 hour, 21 minutes, 39 seconds
Slot 5 information:
State                  Online
Temperature            41 degrees C / 105 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:17:36 PDT
Uptime                1 hour, 21 minutes, 37 seconds
Slot 6 information:
State                  Online
Temperature            40 degrees C / 104 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:17:39 PDT
Uptime                1 hour, 21 minutes, 34 seconds
Slot 7 information:
State                  Online
Temperature            42 degrees C / 107 degrees F
Total CPU DRAM        2048 MB
Total SRAM            64 MB
Total SDRAM           1280 MB
Start time            2010-10-04 20:17:41 PDT
Uptime                1 hour, 21 minutes, 32 seconds

```

### show chassis fpc pic-status (TX Matrix Plus Router)

```
user@host> show chassis fpc pic-status
```

```

lcc0-re0:
-----
Slot 1  Online      FPC Type 2-ES
PIC 0   Online      8x 1GE(LAN), IQ2
Slot 2  Online      FPC Type 4-ES
PIC 0   Online      4x 10GE (LAN/WAN) XFP
Slot 4  Online      FPC Type 4-ES
PIC 0   Online      4x 10GE (LAN/WAN) XFP

```

```

Slot 6  Online      FPC Type 4-ES
      PIC 0  Online      4x 10GE (LAN/WAN) XFP
      PIC 1  Online      4x 10GE (LAN/WAN) XFP
Slot 7  Online      FPC Type 3-ES
      PIC 0  Online      10x 1GE(LAN), 1000 BASE
      PIC 2  Online      1x OC-192 SM SR2
      PIC 3  Online      10x 1GE(LAN), 1000 BASE

```

lcc2-re0:

```

-----
Slot 0  Online      FPC Type 4-ES
      PIC 0  Online      4x 10GE (LAN/WAN) XFP
Slot 2  Online      FPC Type 4-ES
      PIC 0  Online      4x 10GE (LAN/WAN) XFP
      PIC 1  Online      4x 10GE (LAN/WAN) XFP
Slot 3  Online      FPC Type 2-ES
      PIC 0  Online      8x 1GE(LAN), IQ2
Slot 4  Online      FPC Type 4
      PIC 0  Online      10x10GE(LAN/WAN) SFPP
Slot 6  Online      FPC Type 4-ES
      PIC 0  Online      4x OC-192 SONET XFP
Slot 7  Online      FPC Type 3-ES
      PIC 0  Online      10x 1GE(LAN), 1000 BASE
      PIC 1  Offline     1x 10GE(LAN/WAN) IQ2E
      PIC 2  Online      1x OC-192 SM SR2
      PIC 3  Online      1x Tunnel

```

lcc3-re0:

```

-----
Slot 2  Online      FPC Type 4-ES
      PIC 0  Online      10x10GE(LAN/WAN) SFPP
Slot 4  Online      FPC Type 4-ES
      PIC 0  Online      4x OC-192 SONET XFP
Slot 5  Online      FPC Type 4-ES
      PIC 0  Online      4x OC-192 SONET XFP
      PIC 1  Online      4x 10GE (LAN/WAN) XFP
Slot 6  Online      FPC Type 4-ES
      PIC 1  Online      4x 10GE (LAN/WAN) XFP
Slot 7  Online      FPC Type 3-ES
      PIC 0  Online      10x 1GE(LAN), 1000 BASE
      PIC 1  Online      8x 1GE(TYPE3), IQ2E
      PIC 2  Online      4x OC-48 SONET

```

### show chassis fpc (T1600 Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Empty						
1	Empty						
2	Online	49	3	0	2048	3	24
3	Online	46	6	0	2048	6	24
4	Empty						
5	Online	46	5	0	2048	3	24
6	Empty						
7	Online	44	8	0	1024	7	49

### show chassis fpc detail (T1600 Router)

```

user@host> show chassis fpc detail

```

```

show chassis fpc detail
Slot 2 information:
  State                Online
  Temperature          49 degrees C / 120 degrees F
  Total CPU DRAM       2048 MB
  Total SRAM           64 MB
  Total SDRAM          1280 MB
  Start time           2010-10-04 21:12:52 PDT
  Uptime               32 minutes, 9 seconds
Slot 3 information:
  State                Online
  Temperature          47 degrees C / 116 degrees F
  Total CPU DRAM       2048 MB
  Total SRAM           128 MB
  Total SDRAM          2560 MB
  Start time           2010-10-04 21:13:06 PDT
  Uptime               31 minutes, 55 seconds
Slot 5 information:
  State                Online
  Temperature          46 degrees C / 114 degrees F
  Total CPU DRAM       2048 MB
  Total SRAM           64 MB
  Total SDRAM          1280 MB
  Start time           2010-10-04 21:12:56 PDT
  Uptime               32 minutes, 5 seconds
Slot 7 information:
  State                Online
  Temperature          44 degrees C / 111 degrees F
  Total CPU DRAM       1024 MB
  Total SRAM           64 MB
  Total SDRAM          1280 MB
  Start time           2010-10-04 21:14:34 PDT
  Uptime               30 minutes, 27 seconds

```

### show chassis fpc <fpc-slot> (EX Series Switch)

```
user@host> show chassis fpc 2
```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
2	Online	40	12 0	2048	19 14

### show chassis fpc slot (T1600 Router)

```
user@host> show chassis fpc slot 2
```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap Buffer
2	Online	49	3 0	2048	3 24

### show chassis fpc pic-status (T1600 Router)

```
user@host> show chassis fpc pic-status
```

```

Slot 2  Online  FPC Type 1-ES
PIC 0   Online  Load Type 1
PIC 1   Online  4x 1GE(LAN), IQ2E
PIC 3   Online  1x OC-12-3 SFP
Slot 3  Online  FPC Type 4-ES
PIC 0   Online  4x 10GE (LAN/WAN) XFP
PIC 1   Online  4x OC-192 SONET XFP

```

```

Slot 5   Online      FPC Type 2-ES
PIC 0    Online      Load Type 2
PIC 1    Online      8x 1GE(LAN), IQ2E
PIC 2    Online      8x 1GE(LAN), IQ2E
PIC 3    Online      1x OC-48-12-3 SFP
Slot 7   Online      FPC Type 4
PIC 0    Online      4x 10GE (LAN/WAN) XFP

```

### show chassis fpc (T4000 Router)

```
user@host> show chassis fpc
```

```

regress@stymphalian# run show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total Interrupt	Heap	Buffer
0	Online	48	15 0	2816	21 27
1	Empty				
2	Empty				
3	Online	51	15 0	2816	21 27
4	Empty				
5	Online	39	8 0	2048	6 23
6	Online	49	15 0	2816	21 27
7	Empty				

### show chassis fpc detail (T4000 Router)

```
user@host> show chassis fpc detail
```

```
Slot 0 information:
```

```

State                Online
Temperature           48 degrees C / 118 degrees F
Total CPU DRAM        2816 MB
Total SRAM            1554 MB
Total SDRAM           10752 MB
Start time            2012-02-09 22:56:25 PST
Uptime                2 hours, 40 minutes, 52 seconds

```

```
Slot 3 information:
```

```

State                Online
Temperature           51 degrees C / 123 degrees F
Total CPU DRAM        2816 MB
Total SRAM            1554 MB
Total SDRAM           10752 MB
Start time            2012-02-09 22:56:22 PST
Uptime                2 hours, 40 minutes, 55 seconds

```

```
Slot 5 information:
```

```

State                Online
Temperature           39 degrees C / 102 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2012-02-09 22:51:27 PST
Uptime                2 hours, 45 minutes, 50 seconds

```

```
Slot 6 information:
```

```

State                Online
Temperature           49 degrees C / 120 degrees F
Total CPU DRAM        2816 MB
Total SRAM            1554 MB
Total SDRAM           10752 MB
Start time            2012-02-09 22:56:29 PST
Uptime                2 hours, 40 minutes, 48 seconds

```

### show chassis fpc pic-status (T4000 Router)

```
user@host> show chassis fpc pic-status
Slot 0  Online      FPC Type 5-3D
  PIC 0  Online      12x10GE (LAN/WAN) SFPP
  PIC 1  Online      12x10GE (LAN/WAN) SFPP
Slot 3  Online      FPC Type 5-3D
  PIC 0  Online      1x100GE
  PIC 1  Online      12x10GE (LAN/WAN) SFPP
Slot 5  Online      FPC Type 4-ES
  PIC 0  Online      100GE
  PIC 1  Online      100GE CFP
Slot 6  Online      FPC Type 5-3D
  PIC 0  Online      12x10GE (LAN/WAN) SFPP
  PIC 1  Online      12x10GE (LAN/WAN) SFPP
```

### show chassis fpc (QFX Series)

```
user@switch> show chassis fpc
Temp CPU Utilization (%) Memory      Utilization (%)
Slot State              (C) Total Interrupt    DRAM (MB) Heap      Buffer
0 Online                26      2          0        2820      0        49
```

### show chassis fpc detail (QFX3500 Switches)

```
user@switch> show chassis fpc detail
Slot 0 information:
  State                      Online
  Temperature                28 degrees C / 82 degrees F
  Total CPU DRAM              2820 MB
  Total SRAM                  0 MB
  Total SDRAM                 0 MB
  Start time                  2010-09-20 01:34:13 PDT
  Uptime                      3 days, 3 hours, 31 minutes, 48 seconds
```

### show chassis fpc pic-status (QFX3500 Switches)

```
user@switch> show chassis fpc pic-status
Slot 0  Online      QFX 48x10G 4x40G Switch
  PIC 0  Online      48x 10G-SFP+
  PIC 1  Online      15x 10G-SFP+
```

### show chassis fpc interconnect-device (QFabric System)

```
user@switch> show chassis fpc interconnect-device interconnect1
FPC status:

Slot State      Temp
              (C)
0 Online        0
1 Online        0
2 Online        0
3 Online        0
4 Online        0
5 Online        0
6 Online        0
7 Online        0
8 Online        0
9 Online        0
10 Online       0
11 Online       0
12 Online       0
```



```

13 Online      0
14 Online      0
15 Online      0

```

### show chassis fpc interconnect-device (QFabric System)

```

user@switch> show chassis fpc interconnect-device interconnect1 3
FPC status:

Slot State      Temp
          (C)
3 Online        0

```

### show chassis fpc interconnect-device detail (QFabric System)

```

user@switch> show chassis fpc interconnect-device interconnect1 3 detail
Slot 3 information:
State Online
Temperature 0 degrees C / 32 degrees F
Start time 2011-08-18 10:45:04 PDT
Uptime 1 minute, 49 seconds

```

### show chassis fpc pic-status interconnect-device (QFabric System)

```

user@switch> show chassis fpc pic-status interconnect-device interconnect1
Slot 0 Online QFX 16-port QSFP+ Front Card
PIC 0 Online 16x 40G-QSFP+
PIC 1 Online 16x 40G-GE
Slot 1 Online QFX 16-port QSFP+ Front Card
PIC 0 Online 16x 40G-QSFP+
PIC 1 Online 16x 40G-GE
Slot 2 Online QFX 16-port QSFP+ Front Card
PIC 0 Online 16x 40G-QSFP+
PIC 1 Online 16x 40G-GE
Slot 3 Online QFX 16-port QSFP+ Front Card
PIC 0 Online 16x 40G-QSFP+
PIC 1 Online 16x 40G-GE
Slot 4 Online QFX 16-port QSFP+ Front Card
PIC 0 Online 16x 40G-QSFP+
PIC 1 Online 16x 40G-GE
Slot 5 Online QFX 16-port QSFP+ Front Card
PIC 0 Online 16x 40G-QSFP+
PIC 1 Online 16x 40G-GE
Slot 6 Online QFX 16-port QSFP+ Front Card
PIC 0 Online 16x 40G-QSFP+
PIC 1 Online 16x 40G-GE
Slot 7 Online QFX 16-port QSFP+ Front Card
PIC 0 Online 16x 40G-QSFP+
PIC 1 Online 16x 40G-GE
Slot 8 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE
Slot 9 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE
Slot 10 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE
Slot 11 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE
Slot 12 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE
Slot 13 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE
Slot 14 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE

```

```

Slot 15 Online      QFX Fabric Rear Card
PIC 0  Online      16x 40G-GE

```

### show chassis fpc pic-status node-device (QFabric System)

```

user@switch> show chassis fpc pic-status node-device node1
Slot node1 Online      QFX 48x10G 4x40G Switch
PIC 0  Online      48x 10G-SFP+
PIC 1  Online      4x 40G-QSFP+

```

### show chassis fpc (PTX5000 Packet Transport Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%) Total	CPU Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Utilization (%) Buffer
0	Empty						
1	Empty						
2	Online	50	6	0	2816	5	27
3	Empty						
4	Empty						
5	Online	48	9	0	2816	5	27
6	Empty						
7	Online	49	8	0	2816	5	27

### show chassis fpc detail (PTX5000 Packet Transport Router)

```

user@host> show chassis fpc detail
Slot 2 information:
State Online
Temperature 35 degrees C / 95 degrees F (PMB)
Temperature 35 degrees C / 95 degrees F (Intake)
Temperature 50 degrees C / 122 degrees F (Exhaust A)
Temperature 54 degrees C / 129 degrees F (Exhaust B)
Temperature 54 degrees C / 129 degrees F (TL0)
Temperature 52 degrees C / 125 degrees F (TQ0)
Temperature 61 degrees C / 141 degrees F (TL1)
Temperature 58 degrees C / 136 degrees F (TQ1)
Temperature 57 degrees C / 134 degrees F (TL2)
Temperature 58 degrees C / 136 degrees F (TQ2)
Temperature 62 degrees C / 143 degrees F (TL3)
Temperature 61 degrees C / 141 degrees F (TQ3)
Total CPU DRAM 2816 MB
Total SRAM 0 MB
Total SDRAM 0 MB
Start time 2012-01-12 12:05:42 PST
Uptime 3 hours, 14 minutes, 7 seconds
Slot 5 information:
State Online
Temperature 35 degrees C / 95 degrees F (PMB)
Temperature 34 degrees C / 93 degrees F (Intake)
Temperature 48 degrees C / 118 degrees F (Exhaust A)
Temperature 53 degrees C / 127 degrees F (Exhaust B)
Temperature 54 degrees C / 129 degrees F (TL0)
Temperature 52 degrees C / 125 degrees F (TQ0)
Temperature 69 degrees C / 156 degrees F (TL1)
Temperature 56 degrees C / 132 degrees F (TQ1)
Temperature 54 degrees C / 129 degrees F (TL2)
Temperature 56 degrees C / 132 degrees F (TQ2)
Temperature 59 degrees C / 138 degrees F (TL3)
Temperature 60 degrees C / 140 degrees F (TQ3)
Total CPU DRAM 2816 MB

```

```

Total SRAM                0 MB
Total SDRAM               0 MB
Start time                2012-01-12 12:05:43 PST
Uptime                   3 hours, 14 minutes, 6 seconds
Slot 7 information:
State                    Online
Temperature              35 degrees C / 95 degrees F (PMB)
Temperature              33 degrees C / 91 degrees F (Intake)
Temperature              50 degrees C / 122 degrees F (Exhaust A)
Temperature              55 degrees C / 131 degrees F (Exhaust B)
Temperature              56 degrees C / 132 degrees F (TL0)
Temperature              56 degrees C / 132 degrees F (TQ0)
Temperature              61 degrees C / 141 degrees F (TL1)
Temperature              57 degrees C / 134 degrees F (TQ1)
Temperature              55 degrees C / 131 degrees F (TL2)
Temperature              59 degrees C / 138 degrees F (TQ2)
Temperature              62 degrees C / 143 degrees F (TL3)
Temperature              62 degrees C / 143 degrees F (TQ3)
Total CPU DRAM           2816 MB
Total SRAM               0 MB
Total SDRAM              0 MB
Start time                2012-01-12 12:05:44 PST
Uptime                   3 hours, 14 minutes, 5 seconds

```

#### show chassis fpc pic-status (PTX5000 Packet Transport Router)

```

user@host> show chassis fpc pic-status
Slot 2  Online      FPC
PIC 0   Online      24x 10GE(LAN) SFP+
PIC 1   Online      24x 10GE(LAN) SFP+
Slot 5  Online      FPC
PIC 0   Online      24x 10GE(LAN) SFP+
PIC 1   Online      2x 40GE CFP
Slot 7  Online      FPC
PIC 0   Online      24x 10GE(LAN) SFP+
PIC 1   Online      2x 40GE CFP

```

#### show chassis fpc (ACX2000 Universal Access Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)
			Total Interrupt	DRAM (MB) Heap Buffer
0	Online	61	17 6	512 21 37

#### show chassis fpc 0 (ACX2000 Universal Access Router)

```

user@host> show chassis fpc 0

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)
			Total Interrupt	DRAM (MB) Heap Buffer
0	Online	61	17 6	512 21 37

#### show chassis fpc detail (ACX2000 Universal Access Router)

```

user@host> show chassis fpc detail
Slot 0 information:
State                    Online
Temperature              61 degrees C / 141 degrees F
Total CPU DRAM           512 MB
Start time                2012-05-29 02:52:06 PDT
Uptime                   27 minutes, 17 seconds

```

**show chassis fpc pic-status (ACX2000 Universal Access Router)**

```
user@host> show chassis fpc pic-status
Slot 0  Online
  PIC 0  Online      16x CHE1T1, RJ48
  PIC 1  Online      8x 1GE(LAN) RJ45
  PIC 2  Online      2x 1GE(LAN) SFP
  PIC 3  Online      2x 10GE(LAN) SFP+
```

**show chassis FPC 1 (MX Routers with Media Services Blade [MSB])**

```
user@switch> show chassis fpc 1
          Temp CPU Utilization (%)  Memory  Utilization (%)
Slot State      (C) Total Interrupt  DRAM (MB) Heap      Buffer
  1  Online      34      5          0    3072      5      13
```

**show chassis FPC 1 detail (MX Routers with Media Services Blade [MSB])**

```
user@switch> show chassis fpc 1 detail
Slot 1 information:
  State                               Online
  Temperature                         34
  Total CPU DRAM                      3072 MB
  Total RLDRAM                       259 MB
  Total DDR DRAM                     4864 MB
  Start time:                        2012-06-19 10:51:43 PDT
  Uptime:                            16 minutes, 48 seconds
  Max Power Consumption               550 Watts
```

## show chassis led

<b>List of Syntax</b>	<a href="#">show chassis led (EX Series) on page 1197</a> <a href="#">show chassis led (QFX Series) on page 1197</a>
<b>show chassis led (EX Series)</b>	<pre>show chassis led &lt;fpc-slot &lt;fpc-slot-number&gt;&gt;</pre>
<b>show chassis led (QFX Series)</b>	<pre>show chassis led &lt;fpc-slot &lt;fpc-slot-number&gt;&gt; interconnect-device name node-device name</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 10.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Display the status and colors of the chassis LEDs on the front panel of the switch. A major alarm (red) indicates a critical error condition that requires immediate action. A minor alarm (yellow) indicates a noncritical condition that requires monitoring or maintenance. A minor alarm that is left unchecked might cause interruption in service or performance degradation.</p>
<b>Options</b>	<p><b>none</b>—Display the status of the chassis status LEDs (for EX4200 switches configured as a Virtual Chassis, display the information for all Virtual Chassis members).</p> <p><b>fpc-slot &lt;fpc-slot-number&gt;</b>—(Optional) (Not on EX2200 switches) Display the information as follows:</p> <ul style="list-style-type: none"> <li>(EX3200, standalone EX4200, standalone QFX3500, and EX4500 switches) Display the status of the chassis status LEDs for either an FPC slot with no <b>fpc-slot-number</b> value specified or for the FPC slot specified by <b>fpc-slot 0</b>. <b>fpc-slot 0</b> refers to the switch itself and <b>0</b> is the only valid value for <b>fpc-slot-number</b>. Output for these options is the same as for the <b>none</b> option.</li> <li>(EX4200 switches in a Virtual Chassis with two or more members) If no <b>fpc-slot-number</b> value is specified, display the status of the chassis status LEDs for all members of the Virtual Chassis. Output for this option is the same as for the <b>none</b> option. If the <b>fpc-slot-number</b> value is specified (it equals the <b>member-id</b> value), display the status of the chassis status LEDs for the specified member.</li> <li>(EX8200 switches)—Display the status of the chassis status LEDs for the line card in the line-card slot specified by the <b>fpc-slot-number</b> value.</li> </ul> <p><b>interconnect-device name</b>—</p> <p>— (QFabric systems only) (Optional) Display the status of the chassis and interface status LEDs for the Interconnect device.</p> <p><b>node-device name</b>— (QFabric systems only) (Optional) Display the status of the chassis and interface status LEDs for the Node device.</p>

**Required Privilege Level** view

- Related Documentation**
- *Chassis Status LEDs in EX2200 Switches*
  - *Chassis Status LEDs in EX3200 Switches*
  - *Chassis Status LEDs in EX4200 Switches*
  - *Chassis Status LEDs in EX4500 Switches*
  - *Chassis Status LEDs in an EX8200 Switch*
  - *Chassis Status LEDs on a QFX3500 Device*
  - *Chassis Status LEDs in the QFX3600 and QFX3600-I Device*
  - *Management Port LEDs on a QFX3500 Device*
  - *Management Port LEDs in the QFX3600 and QFX3600-I Device*
  - *Chassis Status LEDs on a QFX3008-I Interconnect Device*
  - *Control Board LEDs on a QFX3008-I Interconnect Device*

**List of Sample Output**

[show chassis led \(EX2200 Switch\) on page 1201](#)  
[show chassis led on page 1202](#)  
[show chassis led fpc-slot 0 on page 1203](#)  
[show chassis led \(EX Series\) on page 1203](#)  
[show chassis led node-device \(QFabric System Node Device\) on page 1204](#)  
[show chassis led interconnect-device \(QFabric System - QFX3600-I Interconnect Device\) on page 1204](#)  
[show chassis led interconnect-device \(QFabric System - QFX3008-I Interconnect Device\) on page 1205](#)

**Output Fields** [Table 117 on page 1198](#) lists the output fields for the **show chassis led** command. Output fields are listed in the approximate order in which they appear.

**Table 117: show chassis led Output Fields**

Field Name	Field Description
<b>Front panel contents for slot</b>	FPC slot number of the device whose content is being displayed. The number is always 0, except for EX4200 switches in a Virtual Chassis, where it is the member ID value.
<b>Front panel contents</b> (EX8200 Switches)	
<b>Front Panel Module Information</b> (QFabric system QFX3008-I Interconnect device)	On EX8200 switches, no slot number is displayed.
<b>Front panel contents for</b> (QFabric system Node devices and QFX3600-I Interconnect devices)	On QFabric system Node devices, the name of the Node device whose content is being displayed.

Table 117: show chassis led Output Fields (*continued*)

Field Name	Field Description
<b>Alarms LED</b>	<p>(EX Series switches only) Displays status of the ALM LED:</p> <ul style="list-style-type: none"> <li>• Off—No alarm has been configured.</li> <li>• Green—No alarm has been triggered.</li> <li>• Red—Major alarm.</li> <li>• Yellow—Minor alarm</li> </ul>
<b>System LED</b>	<p>(EX Series switches only) Displays status of the SYS LED:</p> <ul style="list-style-type: none"> <li>• Off—Switch is powered off.</li> <li>• Green—Switch is operating normally.</li> <li>• Yellow—Switch is booting.</li> </ul>
<b>Master LED:</b>	<p>Displays status of the MST LED (on EX3200, EX4200, and EX8200 switches):</p> <ul style="list-style-type: none"> <li>• Green—On an EX4200 Virtual Chassis switch, indicates the switch is the master in the Virtual Chassis configuration. On other switches, indicates that the Routing Engine is operational.</li> <li>• Off <ul style="list-style-type: none"> <li>• On an EX4200 Virtual Chassis switch, indicates that this switch is not the master in the Virtual Chassis configuration.</li> <li>• On EX3200, standalone EX4200, and EX8200 switches, indicates that the Routing Engine is not operational.</li> </ul> </li> </ul>
<b>Mode LED:</b>	<p>(EX Series switches only) On an EX2200 switch only, displays the currently selected port parameter of the Status LED:</p> <ul style="list-style-type: none"> <li>• <b>ADM</b>—Administrative</li> <li>• <b>SPD</b>—Speed</li> <li>• <b>DPX</b>—Duplex</li> <li>• <b>POE</b>—Power over Ethernet</li> </ul>
<b>Status/Beacon LED</b>	<p>(QFX Series only) Displays the system status as indicated by the Status LED on the chassis. For more information, see:</p> <ul style="list-style-type: none"> <li>• <i>Chassis Status LEDs on a QFX3500 Device</i></li> <li>• <i>Chassis Status LEDs in the QFX3600 and QFX3600-I Device</i></li> </ul>
<b>LINK/SPEED LED</b>	<p>(QFX Series only) Displays the link status and speed of a management port. For more information, see:</p> <ul style="list-style-type: none"> <li>• <i>Management Port LEDs on a QFX3500 Device</i></li> <li>• <i>Management Port LEDs in the QFX3600 and QFX3600-I Device</i></li> </ul>
<b>ACTIVITY LED</b>	<p>(QFX Series only) Displays the activity status of a management port. For more information, see:</p> <ul style="list-style-type: none"> <li>• <i>Management Port LEDs on a QFX3500 Device</i></li> <li>• <i>Management Port LEDs in the QFX3600 and QFX3600-I Device</i></li> </ul>

Table 117: show chassis led Output Fields (*continued*)

Field Name	Field Description
<b>STATUS LED</b>	<p>(QFX Series only) Displays the link status of an interface as indicated by the ST LED. For more information, see:</p> <ul style="list-style-type: none"> <li>Control Board LEDs on a QFX3008-I Interconnect Device</li> <li>Access Port and Uplink Port LEDs on a QFX3500 Device</li> <li>Access Port and Uplink Port LEDs on a QFX3600 or QFX3600-I Device</li> </ul>
<b>LINK/ACTIVITY LED</b>	<p>(QFX Series only) Displays link activity or faults on an interface as indicated by the LA LED. For more information, see:</p> <ul style="list-style-type: none"> <li>Access Port and Uplink Port LEDs on a QFX3500 Device</li> <li>Access Port and Uplink Port LEDs on a QFX3600 or QFX3600-I Device</li> </ul>
<b>Status LED</b>	<p>(QFX3008-I Interconnect device only)</p> <ul style="list-style-type: none"> <li>Displays the system status as indicated by the STATUS LED on the front panel of the chassis. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</li> <li>Displays the status of a Control Board as indicated by the STATUS LED on the Control Board. For more information, see <i>Control Board LEDs on a QFX3008-I Interconnect Device</i>.</li> </ul>
<b>Power LED</b>	<p>(QFX3008-I Interconnect device only) Displays the status of system power on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
<b>Major Alarm LED</b>	<p>(QFX3008-I Interconnect device only) Displays whether a critical error condition that requires immediate action exists on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
<b>Minor Alarm LED</b>	<p>(QFX3008-I Interconnect device only) Displays whether a noncritical condition that requires monitoring or maintenance exists on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
<b>Fan 0 LED</b>	<p>(QFX3008-I Interconnect device only) Displays the status of fan trays on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
<b>Fan 1 LED</b>	
<b>Fan 2 LED</b>	
<b>Fan 3 LED</b>	
<b>Fan 4 LED</b>	
<b>Fan 5 LED</b>	
<b>Fan 6 LED</b>	
<b>Fan 7 LED</b>	
<b>Fan 8 LED</b>	



Table 117: show chassis led Output Fields (*continued*)

Field Name	Field Description
PEM 0 LED	(QFX3008-I Interconnect device only) Displays the status of power supplies on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i> .
PEM 1 LED	
PEM 2 LED	
PEM 3 LED	
PEM 4 LED	
LED info for	(QFX3008-I Interconnect device only) Displays the LED information for a Control Board.
Mastership LED	(QFX3008-I Interconnect device only) Displays status of the MASTER LED on a Control Board. For more information, see <i>Control Board LEDs on a QFX3008-I Interconnect Device</i> .
Interface	Names of the interfaces on the device.
LED (ADM/SPD/DPX/POE)	<p>(EX Series switches only) State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are:</p> <p><b>NOTE:</b> EX4500 and EX8200 switches do not have the POE port parameter.</p> <ul style="list-style-type: none"> <li>• <b>ADM</b>—Administrative</li> <li>• <b>SPD</b>—Speed</li> <li>• <b>DPX</b>—Duplex</li> <li>• <b>POE</b>—Power over Ethernet</li> </ul>

## Sample Output

### show chassis led (EX2200 Switch)

```

user@switch> show chassis led
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Amber
  System LED: Green
  Mode LED : Duplex
Interface    LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      Off
ge-0/0/1      Full Duplex
ge-0/0/2      Full Duplex
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Full Duplex
ge-0/0/6      Full Duplex
ge-0/0/7      Full Duplex
ge-0/0/8      Full Duplex
ge-0/0/9      Full Duplex
ge-0/0/10     Full Duplex
ge-0/0/11     Full Duplex

```

ge-0/0/12	Full Duplex
ge-0/0/13	Full Duplex
ge-0/0/14	Full Duplex
ge-0/0/15	Full Duplex
ge-0/0/16	Full Duplex
ge-0/0/17	Full Duplex
ge-0/0/18	Full Duplex
ge-0/0/19	Full Duplex
ge-0/0/20	Full Duplex
ge-0/0/21	Full Duplex
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Full Duplex
ge-0/0/25	Full Duplex
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Full Duplex
ge-0/0/29	Full Duplex

### show chassis led

```
user@switch> show chassis led
```

```
Front panel contents for slot: 0
```

```
-----  
LEDs status:
```

```
  Alarms LED: Off
```

```
  System LED: Green
```

```
  Master LED: Green
```

```
Interface      LED (ADM/SPD/DPX/POE)
```

```
-----  
ge-0/0/0      Off  
ge-0/0/1      Full Duplex  
ge-0/0/2      Full Duplex  
ge-0/0/3      Off  
ge-0/0/4      Off  
ge-0/0/5      Full Duplex  
ge-0/0/6      Full Duplex  
ge-0/0/7      Full Duplex  
ge-0/0/8      Full Duplex  
ge-0/0/9      Full Duplex  
ge-0/0/10     Full Duplex  
ge-0/0/11     Full Duplex  
ge-0/0/12     Full Duplex  
ge-0/0/13     Full Duplex  
ge-0/0/14     Full Duplex  
ge-0/0/15     Full Duplex  
ge-0/0/16     Full Duplex  
ge-0/0/17     Full Duplex  
ge-0/0/18     Full Duplex  
ge-0/0/19     Full Duplex  
ge-0/0/20     Full Duplex  
ge-0/0/21     Full Duplex  
ge-0/0/22     Off  
ge-0/0/23     Off  
ge-0/0/24     Full Duplex  
ge-0/0/25     Full Duplex  
ge-0/0/26     Off  
ge-0/0/27     Off  
ge-0/0/28     Full Duplex  
ge-0/0/29     Full Duplex
```

## show chassis led fpc-slot 0

```

user@switch> show chassis led fpc-slot 0
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Red
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      Off
ge-0/0/1      Off
ge-0/0/2      Off
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Off
ge-0/0/6      Off
ge-0/0/7      Off
ge-0/0/8      Off
ge-0/0/9      Off
ge-0/0/10     Off
ge-0/0/11     Off
ge-0/0/12     Off
ge-0/0/13     Off
ge-0/0/14     Off
ge-0/0/15     Off
ge-0/0/16     Off
ge-0/0/17     Off
ge-0/0/18     Off
ge-0/0/19     Off
ge-0/0/20     Off
ge-0/0/21     Off
ge-0/0/22     Off
ge-0/0/23     Off

```

## show chassis led (EX Series)

```

user@switch> show chassis led
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Amber
  Status LED: Green
  Mode LED : Duplex
Interface LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0 Off
ge-0/0/1 Full Duplex
ge-0/0/2 Full Duplex
ge-0/0/3 Off
ge-0/0/4 Off
ge-0/0/5 Full Duplex
ge-0/0/6 Full Duplex
ge-0/0/7 Full Duplex
ge-0/0/8 Full Duplex
ge-0/0/9 Full Duplex
ge-0/0/10 Full Duplex
ge-0/0/11 Full Duplex
ge-0/0/12 Full Duplex
ge-0/0/13 Full Duplex

```

```

ge-0/0/14 Full Duplex
ge-0/0/15 Full Duplex
ge-0/0/16 Full Duplex
ge-0/0/17 Full Duplex
ge-0/0/18 Full Duplex
ge-0/0/19 Full Duplex
ge-0/0/20 Full Duplex
ge-0/0/21 Full Duplex
ge-0/0/22 Off
ge-0/0/23 Off
ge-0/0/24 Full Duplex
ge-0/0/25 Full Duplex
ge-0/0/26 Off
ge-0/0/27 Off
ge-0/0/28 Full Duplex
ge-0/0/29 Full Duplex

```

### show chassis led node-device (QFabric System Node Device)

```

user@switch> show chassis led node-device node1
Front panel contents for: node1
LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	LINK/SPEED LED	ACTIVITY LED
node1:me5	Green	N/A
node1:me6	Green	N/A

Interface	STATUS LED	LINK/ACTIVITY LED
node1:xe-0/0/8	Green	Green
node1:ge-0/0/10	Green	Green
node1:ge-0/0/12	Green	Green
node1:ge-0/0/24	Green	Green
node1:ge-0/0/25	Green	Green
node1:ge-0/0/26	Green	Green
node1:ge-0/0/27	Green	Green
node1:ge-0/0/28	Green	Green
node1:ge-0/0/29	Green	Green
node1:ge-0/0/30	Green	Green
node1:ge-0/0/31	Green	Green
node1:ge-0/0/32	Green	Green
node1:ge-0/0/33	Green	Green
node1:ge-0/0/34	Green	Green
node1:ge-0/0/35	Green	Green
node1:ge-0/0/36	Green	Green
node1:ge-0/0/37	Green	Green
node1:ge-0/0/38	Green	Green
node1:ge-0/0/39	Green	Green
node1:fte-0/1/0	Green	Green Blinking
node1:fte-0/1/2	Green	Green Blinking

### show chassis led interconnect-device (QFabric System - QFX3600-I Interconnect Device)

```

user@switch> show chassis led interconnect-device IC-EG0712
Front panel contents for: FPC 0

```

```

LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	LINK/SPEED LED	ACTIVITY LED
IC-EG0712:me5	Green	N/A
IC-EG0712:me6	Green	N/A

Interface	STATUS LED	LINK/ACTIVITY LED
IC-EG0712:fte-0/1/0	Green	Green
IC-EG0712:fte-0/1/1	Green	Green Blinking
IC-EG0712:fte-0/1/2	Green	Green
IC-EG0712:fte-0/1/3	Green	Green Blinking
IC-EG0712:fte-0/1/4	Green	Green
IC-EG0712:fte-0/1/5	Green	Green Blinking
IC-EG0712:fte-0/1/6	Green	Green
IC-EG0712:fte-0/1/7	Green	Green
IC-EG0712:fte-0/1/8	Green	Green Blinking
IC-EG0712:fte-0/1/9	Green	Green Blinking
IC-EG0712:fte-0/1/10	Green	Green Blinking

### show chassis led interconnect-device (QFabric System - QFX3008-I Interconnect Device)

```

user@switch> show chassis led interconnect-device IC-EG0712
                    Front Panel Module Information
                    -----
LEDs status:
  Status LED: Green
  Power LED : Yellow Blinking
  Major Alarm LED: Red
  Minor Alarm LED: Yellow
  Fan 0 LED : Green
  Fan 1 LED : Green
  Fan 2 LED : Green
  Fan 3 LED : Green
  Fan 4 LED : Green
  Fan 5 LED : Green
  Fan 6 LED : Green
  Fan 7 LED : Green
  Fan 8 LED : Green
  Fan 9 LED : Green
  PEM 0 LED : Green
  PEM 1 LED : Green
  PEM 2 LED : Green
  PEM 3 LED : off
  PEM 4 LED : Yellow Blinking
  PEM 5 LED : off

                    LED info for: CB - 0
                    -----
LEDs status:
  Status LED: Green
  Mastership LED: Green

Interface          STATUS LED    LINK/ACTIVITY LED
-----
IC-F4899:pme0 :    Green        N/A
IC-F4899:pme1 :    off          N/A
IC-F4899:pme2 :    off          N/A
IC-F4899:pme3 :    off          N/A

                    LED info for: CB - 1
                    -----

```

## LEDs status:

Status LED: Green

Mastership LED: Amber

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:pme0 :	Green	N/A
IC-F4899:pme1 :	off	N/A
IC-F4899:pme2 :	off	N/A
IC-F4899:pme3 :	off	N/A

LED info for: FC 0 FPC - 0

## LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:fte-0/0/0	Green	N/A
IC-F4899:fte-0/0/1	Green	N/A
IC-F4899:fte-0/0/2	Green	N/A
IC-F4899:fte-0/0/3	Green	N/A
IC-F4899:fte-0/0/4	Green	N/A
IC-F4899:fte-0/0/5	Green	N/A
IC-F4899:fte-0/0/6	Green	N/A
IC-F4899:fte-0/0/7	Green	N/A
IC-F4899:fte-0/0/8	Green	N/A
IC-F4899:fte-0/0/9	Green	N/A
IC-F4899:fte-0/0/10	Green	N/A
IC-F4899:fte-0/0/11	Green	N/A
IC-F4899:fte-0/0/12	Green	N/A
IC-F4899:fte-0/0/13	Green	N/A
IC-F4899:fte-0/0/14	Green	N/A
IC-F4899:fte-0/0/15	Green	N/A

LED info for: FC 1 FPC - 1

## LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:fte-1/0/0	Green	N/A
IC-F4899:fte-1/0/1	Green	N/A

LED info for: RC 2 FPC - 10

## LEDs status:

Status LED: Green

LED info for: RC 3 FPC - 11

## LEDs status:

Status LED: Green

## show chassis location

<b>List of Syntax</b>	<a href="#">Syntax on page 1207</a> <a href="#">Syntax (TX Matrix Router) on page 1207</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1207</a> <a href="#">Syntax (MX Series Router) on page 1207</a> <a href="#">Syntax (QFX Series) on page 1207</a>
<b>Syntax</b>	show chassis location
<b>Syntax (TX Matrix Router)</b>	show chassis location <fpc   interface (by-name <i>name</i>   by-slot fpc number lcc number)   lcc number   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show chassis location <fpc   interface (by-name <i>name</i>   by-slot fpc number lcc number)   lcc number   sfc number>
<b>Syntax (MX Series Router)</b>	show chassis location <all-members> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show chassis location <interconnect-device <i>name</i> > <node-device <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the physical location of the chassis. This command can only be used on the master Routing Engine.
<b>Options</b>	<p><b>none</b>—Display all information about the physical location of the chassis. On a TX Matrix router, display all information about the physical location of the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display all information about the physical location of the TX Matrix Plus router and its attached routers.</p> <p><b>all-members</b>—(MX Series routers only) (Optional) Display the physical location of the chassis for all the member routers in the Virtual Chassis configuration.</p> <p><b>fpc</b>—(TX Matrix router and TX Matrix Plus router only) (Optional) Display the physical location of all Flexible PIC Concentrators (FPCs).</p> <p><b>interconnect-device <i>name</i></b>—(QFabric systems only) (Optional) Display the physical location of the Interconnect device.</p> <p><b>interface by-name <i>name</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) Display the physical location of a specified interface name. On a TX Matrix router, this option displays the FPC number and T640 router (line-card chassis) number associated</p>

with the specified interface. On a TX Matrix Plus router, this option displays the FPC number and router (line-card chassis) number associated with the specified interface.

**interface by-slot fpc *number* lcc *number***—(TX Matrix and TX Matrix Plus router only)

(Optional) On a TX Matrix router, display the global FPC number of an interface by specifying its local FPC number and T640 router (line-card chassis) number. On a TX Matrix Plus router, display the global FPC number of an interface by specifying its local FPC number and router (line-card chassis) number.

- The global FPC number is the FPC slot number when all the FPC slots in the routing matrix are considered: **0** through **31**. On TX Matrix Plus router with 3D SIBs, the value is **0** through **63**. The local FPC number is the FPC slot number on a particular T640 router.
- For **fpc**, replace *number* with a value from **0** through **7**.
- For **lcc**, replace *number* with a value from **0** through **7**.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the physical location of a specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display the physical location of a specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display the physical location of the chassis for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display the physical location of the chassis for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**node-device *name***—(QFabric systems only) (Optional) Display the physical location of the Node device.

**scc**—(TX Matrix routers only) (Optional) Display the physical location of the TX Matrix router (switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display the physical location of the TX Matrix Plus router (or switch-fabric chassis).



**Required Privilege Level** view

**Related Documentation** • *Displaying Chassis Physical Locations for a Routing Matrix with a TX Matrix Plus Router*

**List of Sample Output** [show chassis location on page 1209](#)  
[show chassis location fpc \(TX Matrix Router\) on page 1210](#)  
[show chassis location interface by-slot \(TX Matrix Router\) on page 1210](#)  
[show chassis location fpc \(TX Matrix Plus Router\) on page 1210](#)  
[show chassis location interface by-slot \(TX Matrix Plus Router\) on page 1210](#)  
[show chassis location \(QFX3500 Switches\) on page 1210](#)  
[show chassis location \(QFabric Systems\) on page 1210](#)

**Output Fields** Table 118 on page 1209 lists the output fields for the **show chassis location** command. Output fields are listed in the approximate order in which they appear.

**Table 118: show chassis location Output Fields**

Field Name	Field Description
country-code	Country code information.
postal-code	Postal code information.
Building	Building information.
Floor	Floor information.
Global FPC	Global FPC number. The FPC slot number, when all FPC slots in the routing matrix are considered. The range of values is 0 through 31. On TX Matrix Plus router with 3D SIBs the value is 0 through 63.
LATA	Local access transport area information.
LCC	Line-card chassis number. On a TX Matrix router, the number of a particular T640 router connected to the TX Matrix router. On a TX Matrix Plus router, the number of a particular router connected to the TX Matrix Plus router.
Local FPC	Local FPC number. On a TX Matrix router, the FPC slot number on a particular T640 router. On a TX Matrix Plus router, the FPC slot number on a particular router.

## Sample Output

**show chassis location**

```
user@host> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

**show chassis location fpc (TX Matrix Router)**

```
user@host> show chassis location fpc
Global FPC    LCC    Local FPC
    17         2        1
    21         2        5
```

**show chassis location interface by-slot (TX Matrix Router)**

```
user@host> show chassis location interface by-slot fpc 1 lcc 1
Global FPC: 9
```

**show chassis location fpc (TX Matrix Plus Router)**

```
user@host> show chassis location fpc
Global FPC    LCC    Local FPC
    0         0        0
    1         0        1
```

**show chassis location interface by-slot (TX Matrix Plus Router)**

```
user@host> show chassis location interface by-slot fpc 2 lcc 1
Global FPC: 10
```

**show chassis location (QFX3500 Switches)**

```
user@switch> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

**show chassis location (QFabric Systems)**

```
user@switch> show chassis location interconnect-device interconnect1
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

## show chassis pic

<b>List of Syntax</b>	<a href="#">Syntax on page 1211</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 1211</a> <a href="#">Syntax (MX Series Routers) on page 1211</a> <a href="#">Syntax (MX104, MX2010 and MX2020 3D Universal Edge Routers) on page 1211</a> <a href="#">Syntax (PTX Series Packet Transport Router) on page 1211</a> <a href="#">Syntax (QFX Series) on page 1211</a> <a href="#">Syntax (ACX Series Universal Access Routers) on page 1211</a>
<b>Syntax</b>	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> &lt;lcc <i>number</i>&gt;</code>
<b>Syntax (MX Series Routers)</b>	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> &lt;all-members&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt;</code>
<b>Syntax (MX104, MX2010 and MX2020 3D Universal Edge Routers)</b>	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
<b>Syntax (PTX Series Packet Transport Router)</b>	<code>show chassis pic transport fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
<b>Syntax (QFX Series)</b>	<code>show chassis pic &lt;interconnect-device <i>name</i> (fpc-slot <i>slot-number</i>   pic-slot <i>slot-number</i>)&gt; &lt;node-device <i>name</i> pic-slot <i>slot-number</i>&gt;</code>
<b>Syntax (ACX Series Universal Access Routers)</b>	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.</p>
<b>Description</b>	Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.
<b>Options</b>	<b>fpc-slot <i>slot-number</i></b> —Display information about the PIC in this particular FPC slot:

- On a TX Matrix router, if you specify the number of the T640 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, replace **slot-number** with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, replace **slot-number** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis pic fpc-slot 1 lcc 1 pic-slot 1
user@host> show chassis pic fpc-slot 9 pic-slot 1
```

- M120 routers only—Replace **slot-number** with a value from 0 through 5.
- MX80 routers only—Replace **slot-number** with a value from 0 through 1.
- MX104 routers only—Replace **slot-number** with a value from 0 through 2.
- MX240 routers only—Replace **slot-number** with a value from 0 through 2.
- MX480 routers only—Replace **slot-number** with a value from 0 through 5.
- MX960 routers only—Replace **slot-number** with a value from 0 through 11.
- MX2010 routers only—Replace **slot-number** with a value from 0 through 9.
- MX2020 routers only—Replace **slot-number** with a value from 0 through 19.
- Other routers—Replace **slot-number** with a value from 0 through 7.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—Replace **slot-number** with 0.
  - EX4200 switches in a Virtual Chassis configuration—Replace **slot-number** with a value from 0 through 9 (switch's member ID).
  - EX8208 switches—Replace **slot-number** with a value from 0 through 7 (line card).
  - EX8216 switches—Replace **slot-number** with a value from 0 through 15 (line card).
- QFX Series:
  - QFX3500 and QFX5100 standalone switches—Replace **slot-number** with 0. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.
  - QFabric systems—Replace **slot-number** with any number between 0 and 15. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.

**all-members**—(MX Series routers and EX Series switches only) (Optional) Display PIC information for all member routers in the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display PIC information for a specified Interconnect device.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display PIC information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display PIC information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the local Virtual Chassis member.

**member *member-id***—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**node-device *name***—(QFabric systems only) (Optional) Display PIC information for a specified Node device.

**pic-slot *slot-number***—Display information about the PIC in this particular PIC slot. For routers, replace *slot-number* with a value from 0 through 3. For EX3200 and EX4200 switches, replace *slot-number* with 0 for built-in network interfaces and 1 for interfaces on uplink modules. For EX8208 and EX8216 switches, replace *slot-number* with 0. For the QFX3500 standalone switch and the QFabric system, replace *slot-number* with 0 or 1.

**transport**—Display PIC information for optical transport network.

**Required Privilege Level**

view

**Related Documentation**

- [request chassis pic on page 92](#)
- *show chassis hardware*
- *Configuring the PIC Type*
- *100-Gigabit Ethernet Type 4 PIC with CFP Overview*

**List of Sample Output**    [show chassis pic fpc-slot pic-slot on page 1216](#)

[show chassis pic fpc-slot pic-slot \(PIC Offline\) on page 1217](#)  
[show chassis pic fpc-slot pic-slot \(FPC Offline\) on page 1217](#)  
[show chassis pic fpc-slot pic-slot \(FPC Not Present\) on page 1217](#)  
[show chassis pic fpc-slot pic-slot \(PIC Not Present\) on page 1217](#)  
[show chassis pic fpc-slot pic-slot \(M120 Router\) on page 1217](#)  
[show chassis pic fpc-slot pic-slot \(MX104 Router\) on page 1217](#)  
[show chassis pic fpc-slot pic-slot \(MX960 Router Bidirectional Optics\) on page 1218](#)  
[show chassis pic fpc-slot pic-slot \(MX480 Router with 100-Gigabit Ethernet MIC\) on page 1218](#)  
[show chassis pic fpc-slot pic-slot \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 1218](#)  
[show chassis pic fpc-slot pic-slot \(MX960 Router with MPC5EQ\) on page 1219](#)  
[show chassis pic fpc-slot pic-slot \(MX480 Routers with MPC4E\) on page 1219](#)  
[show chassis pic fpc-slot pic-slot \(MX480 routers with OTN Interfaces\) on page 1219](#)  
[show chassis pic fpc-slot pic-slot \(MX2010 Routers with OTN Interfaces\) on page 1219](#)  
[show chassis pic fpc-slot pic-slot \(MX2010 Routers\) on page 1220](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Routers\) on page 1220](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Routers with MPC5EQ and MPC6E\) on page 1220](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Routers with MPC6E and OTN MIC\) on page 1221](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Routers with MPC4E\) on page 1221](#)  
[show chassis pic fpc-slot pic-slot \(T1600 Router with 100-Gigabit Ethernet PIC\) on page 1221](#)  
[show chassis pic fpc-slot pic-slot lcc \(TX Matrix Router\) on page 1222](#)  
[show chassis pic fpc-slot pic-slot lcc \(TX Matrix Plus Router\) on page 1222](#)  
[show chassis pic fpc-slot pic-slot \(Next-Generation SONET/SDH SFP\) on page 1222](#)  
[show chassis pic fpc-slot pic-slot \(12-Port T1/E1\) on page 1222](#)  
[show chassis pic fpc-slot pic-slot \(4x CHOC3 SONET CE SFP\) on page 1223](#)  
[show chassis pic fpc-slot pic-slot \(SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 1223](#)  
[show chassis pic fpc-slot pic-slot \(8-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 1223](#)  
[show chassis pic fpc-slot pic-slot \(4-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 1224](#)  
[show chassis pic fpc-slot pic-slot \(1-port OC192/STM64 MIC with XFP\) on page 1224](#)  
[show chassis pic fpc-slot 1 pic-slot 2 \(8-port DS3/E3 MIC\) on page 1224](#)  
[show chassis pic fpc-slot pic-slot \(OTN\) on page 1224](#)  
[show chassis pic fpc-slot pic-slot \(QFX3500 Switch\) on page 1224](#)  
[show chassis pic fpc-slot pic-slot \(QFX5100 Standalone Switch\) on page 1225](#)  
[show chassis pic interconnect-device fpc-slot pic-slot \(QFabric Systems\) on page 1225](#)  
[show chassis pic node-device fpc-slot pic-slot \(QFabric System\) on page 1225](#)  
[show chassis pic fpc-slot pic-slot \(ACX2000 Universal Access Router\) on page 1226](#)  
[show chassis pic fpc-slot pic-slot \(MX Routers with Media Services Blade \[MSB\]\) on page 1226](#)  
[show chassis pic FPC slot PIC slot \(MX Routers with Media Services Blade \[MSB\]\) on page 1226](#)  
[show chassis pic transport fpc-slot pic-slot \(PTX Series Packet Transport Routers\) on page 1226](#)

**Output Fields** Table 119 on page 1215 lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

**Table 119: show chassis pic Output Fields**

Field Name	Field Description
<b>Type</b>	<p>PIC type.</p> <p><b>NOTE:</b> On the 1-port OC192/STM64 MICs with the SDH framing mode, the type is displayed as <b>MIC-3D-1STM64-XFP</b> and with the SONET framing mode, the type is displayed as <b>MIC-3D-1OC192-XFP</b>. By default, the 1-port OC192/STM64 MICs displays the type as <b>MIC-3D-1OC192-XFP</b>.</p>
<b>Account Layer2 Overhead</b>	(MX Series routers) Indicates whether functionality to count the Layer 2 overhead bytes in the interface statistics at the PIC level is enabled or disabled.
<b>ASIC type</b>	Type of ASIC on the PIC.
<b>State</b>	<p>Status of the PIC. State is displayed only when a PIC is in the slot.</p> <ul style="list-style-type: none"> <li>• <b>Online</b>— PIC is online and running.</li> <li>• <b>Offline</b>—PIC is powered down.</li> </ul>
<b>PIC version</b>	PIC hardware version.
<b>Uptime</b>	How long the PIC has been online.
<b>Package</b>	(Multiservices PICs only) Services package supported: <b>Layer-2</b> or <b>Layer-3</b> .
<b>Port Number</b>	Port number for the PIC.
<b>Cable Type</b>	Type of cable connected to the port: <b>LH</b> , <b>LX</b> , or <b>SX</b> .
<b>PIC Port Information (MX480 Router 100-Gigabit Ethernet CFP)</b>	<p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> <li>• Port—Port number</li> <li>• Cable type—Type of optical transceiver installed.</li> <li>• Fiber type—Type of fiber. SM is single-mode.</li> <li>• Xcvr vendor—Transceiver vendor name.</li> <li>• Xcvr vendor part number—Transceiver vendor part number.</li> <li>• Wavelength—Wavelength of the transmitted signal. Uplinks and downlinks are always 1550 nm. There is a separate fiber for each direction</li> </ul>

Table 119: show chassis pic Output Fields (*continued*)

Field Name	Field Description
<b>PIC Port Information (MX960 Router Bidirectional Optics )</b>	<p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> <li>Port—Port number</li> <li>Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. Uplink interfaces display -U. Down link interfaces display -D.</li> <li>Fiber type—Type of fiber. SM is single-mode.</li> <li>Xcvr vendor—Transceiver vendor name.</li> <li>Xcvr vendor part number—Transceiver vendor part number. <ul style="list-style-type: none"> <li>BX10-10-km bidirectional optics.</li> <li>BX40-40-km bidirectional optics.</li> <li>SFP-LX-40-km SFP optics.</li> </ul> </li> <li>Wavelength—Wavelength of the transmitted signal. Uplinks are always 1310 nm. Downlinks are either 1490 nm or 1550 nm.</li> </ul>
<b>PIC Port Information (Next-Generation SONET/SDH SFP)</b>	<p>Port-level information for the next-generation SONET/SDH SFP PIC.</p> <ul style="list-style-type: none"> <li>Port—Port number.</li> <li>Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed.</li> <li>Fiber type—Type of fiber: <b>SM</b> (single-mode) or <b>MM</b> (multimode).</li> <li>Xcvr vendor—Transceiver vendor name.</li> <li>Xcvr vendor part number—Transceiver vendor part number.</li> <li>Wavelength—Wavelength of the transmitted signal. Next-generation SONET/SDH SFPs use 1310 nm.</li> </ul>
<b>Pic port information (MX104 router)</b>	<p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> <li>Port—Port number</li> <li>Cable type—Type of optical transceiver installed.</li> <li>Fiber type—Type of fiber. SM is single-mode.</li> <li>Xcvr vendor—Transceiver vendor name.</li> <li>Xcvr vendor part number—Transceiver vendor part number.</li> <li>Wavelength—Wavelength of the transmitted signal.</li> <li>Xcvr Firmware—Firmware version of the transceiver.</li> </ul>
<b>Multirate Mode</b>	Rate-selectability status for the MIC: <b>Enabled</b> or <b>Disabled</b> .
<b>Channelization</b>	Indicates whether channelization is enabled or disabled on the DS3/E3 MIC.

## Sample Output

### show chassis pic fpc-slot pic-slot

```

user@host> show chassis pic fpc-slot 2 pic-slot 0
PIC fpc slot 2 pic slot 0 information:
Type                               10x 1GE(LAN), 1000 BASE

```



```

ASIC type           H chip
State               Online
PIC version         1.1
Uptime              1 day, 50 minutes, 58 seconds
PIC Port Information:
Port      Cable      Xcvr      Xcvr Vendor
Number    Type        Vendor Name  Part Number
0         GIGE 1000EX  FINISAR CORP.  FTRJ8519P1BNL-J3
1         GIGE 1000EX  FINISAR CORP.  FTRJ-8519-7D-JUN

```

#### show chassis pic fpc-slot pic-slot (PIC Offline)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
PIC fpc slot 1 pic slot 0 information:
State               Offline

```

#### show chassis pic fpc-slot pic-slot (FPC Offline)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC 1 is not online

```

#### show chassis pic fpc-slot pic-slot (FPC Not Present)

```

user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4 is empty

```

#### show chassis pic fpc-slot pic-slot (PIC Not Present)

```

user@host> show chassis pic fpc-slot 5 pic-slot 2
FPC 5, PIC 2 is empty

```

#### show chassis pic fpc-slot pic-slot (M120 Router)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
PC slot 3, PIC slot 0 information:
Type           2x G/E IQ, 1000 BASE
ASIC type      IQ GE 2 VLAN-TAG FPGA
State          Online
PIC version    1.16
Uptime         3 hours, 3 minutes

PIC Port Information:
Port      Cable      Xcvr      Xcvr Vendor
Number    Type        Vendor Name  Part Number
0         GIGE 1000SX  FINISAR CORP.  FTRJ8519P1BNL-J3
1         GIGE 1000SX  FINISAR CORP.  FTRJ-8519-7D-JUN

```

#### show chassis pic fpc-slot pic-slot (MX104 Router)

```

user@host> show chassis pic fpc-slot 1 pic-slot 1
FPC slot 1, PIC slot 1 information:
Type           10x 1GE(LAN) -E SFP
State          Online
PIC version    1.1
Uptime         1 hour, 30 minutes, 59 seconds

PIC port information:
Fiber      Xcvr vendor      Wave-      Xcvr
Port Cable type    type Xcvr vendor      part number  length
Firmware
3   GIGE 1000T    n/a  Methode Elec.    SP7041-M1-JN  n/a      0.0

```

6	GIGE 1000LX10	SM	FINISAR CORP.	FTLF1318P2BTL-J1	1310 nm	0.0
8	GIGE 1000T	n/a	Methode Elec.	SP7041-M1-JN	n/a	0.0
9	GIGE 1000T	n/a	Methode Elec.	SP7041-M1-JN	n/a	0.0

### show chassis pic fpc-slot pic-slot (MX960 Router Bidirectional Optics)

```

user@host> show chassis pic fpc-slot 4 pic-slot 1
FPC slot 4, PIC slot 1 information:
  Type                               10x 1GE(LAN)
  Account Layer2 Overhead            Enabled
  State                               Online
  PIC version                         0.0
  Uptime                             18 days, 5 hours, 41 minutes, 54 seconds

PIC port information:
  Port  Cable type      Fiber type  Xcvr vendor      Xcvr vendor      Wavelength
      Cable type      type      part number
0      SFP-1000BASE-BX10-D SM  SumitomoElectric SBP6H44-J3-BW-49 1490 nm
1      SFP-1000BASE-BX10-D SM  SumitomoElectric SBP6H44-J3-BW-49 1490 nm
2      SFP-1000BASE-BX10-D SM  SumitomoElectric SBP6H44-J3-BW-49 1490 nm
3      SFP-1000BASE-BX10-D SM  OCP              TRXBG1LXDBVM2-JW 1490 nm
4      SFP-1000BASE-BX10-D SM  OCP              TRXBG1LXDBVM2-JW 1490 nm
5      SFP-1000BASE-BX10-U SM  SumitomoElectric SBP6H44-J3-BW-31 1310 nm
6      SFP-1000BASE-BX10-U SM  SumitomoElectric SBP6H44-J3-BW-31 1310 nm
7      SFP-1000BASE-BX10-U SM  OCP              TRXBG1LXDBBMH-J1 1310 nm
8      SFP-1000BASE-BX10-U SM  OCP              TRXBG1LXDBBMH-J1 1310 nm
9      SFP-1000BASE-BX10-U SM  SumitomoElectric SBP6H44-J3-BW-31 1310 nm

```

### show chassis pic fpc-slot pic-slot (MX480 Router with 100-Gigabit Ethernet MIC)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                               1X100GE CFP
  State                               Online
  PIC version                         2.10
  Uptime                             4 minutes, 48 seconds

PIC port information:
  Fiber                               Xcvr vendor
  Port  Cable type      type  Xcvr vendor      part number      Wavelength
0      100GBASE LR4      SM    FINISAR CORP.    FTLC1181RDN3-J3 1310 nm

  Xcvr vendor
  firmware version
  1.8

```

### show chassis pic fpc-slot pic-slot (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                               AS-MXC
  State                               Online
  PIC version                         1.0
  Uptime                             11 hours, 18 minutes, 3 seconds

```

**show chassis pic fpc-slot pic-slot (MX960 Router with MPC5EQ)**

```

user@host> show chassis pic fpc-slot 0 pic-slot 3
FPC slot 0, PIC slot 3 information:
  Type                1X100GE CFP2 OTN
  State                Online
  PIC version          0.0
  Uptime               1 hour, 22 minutes, 42 seconds

PIC port information:

```

	Fiber	Xcvr vendor	Wave-	Xcvr
Port Cable type	type	Xcvr vendor	part number	length
Firmware				
0	10GBASE LR4	n/a	Oclaro Inc.	TRB5E20FNF-LF150 1309 nm 1.0

**show chassis pic fpc-slot pic-slot (MX480 Routers with MPC4E)**

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type                4x10GE SFPP
  State                Online
  PIC version          0.0
  Uptime               41 seconds

PIC port information:

```

	Fiber	Xcvr vendor	Wave-	Xcvr
Port Cable type	type	Xcvr vendor	part number	length
Firmware				
0	10GBASE SR	MM	OPNEXT, INC.	TRS2001EM-0014 850 nm 0.0
1	10GBASE SR	MM	OPNEXT, INC.	TRS2001EM-0014 850 nm 0.0

**show chassis pic fpc-slot pic-slot (MX480 routers with OTN Interfaces)**

```

user@host> show chassis pci fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
  Type                12X10GE SFPP OTN
  State                Online
  PIC version          0.0
  Uptime               5 hours, 28 minutes, 23 seconds

PIC port information:

```

	Fiber	Xcvr vendor	Wave-	Xcvr
Port Cable type	type	Xcvr vendor	part number	length
Firmware				
0	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BNL-J1 850 nm 0.0
1	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BCL-J1 850 nm 0.0
2	10GBASE SR	MM	OPNEXT, INC.	TRS2001EM-0014 850 nm 0.0

**show chassis pic fpc-slot pic-slot (MX2010 Routers with OTN Interfaces)**

```

user@host> show chassis pic fpc-slot 9 pic-slot 0

```

FPC slot 9, PIC slot 0 information:

```
Type                2X100GE CFP2 OTN
State                Online
PIC version          1.9
Uptime               3 hours, 56 minutes, 16 seconds
```

PIC port information:

		Fiber	Xcvr vendor		Wave-	Xcvr
Port	Cable type	type	Xcvr vendor	part number	length	
Firmware						
0	100GBASE LR4-D	SM	FUJITSU	FIM37300/222	1310 nm	1.3
1	100GBASE SR10	MM	AVAGO	AFBR-8420Z	n/a	1.0

#### show chassis pic fpc-slot pic-slot (MX2010 Routers)

```
user@host> show chassis pic fpc-slot 9 pic-slot 3
```

FPC slot 9, PIC slot 3 information:

```
Type                1X100GE CFP
Account Layer2 Overhead Enabled
State                Online
PIC version          0.0
Uptime               14 hours, 51 seconds
```

#### show chassis pic fpc-slot pic-slot (MX2020 Routers)

```
user@host> show chassis pic fpc-slot 19 pic-slot 3
```

FPC slot 19, PIC slot 3 information:

```
Type                4x 10GE(LAN) SFP+
Account Layer2 Overhead Enabled
State                Online
PIC version          0.0
Uptime               1 day, 11 hours, 26 minutes, 36 seconds
```

PIC port information:

		Fiber	Xcvr vendor		Wave-	Xcvr
Port	Cable type	type	Xcvr vendor	part number	length	
Firmware						
0	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
1	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
2	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
3	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0

#### show chassis pic fpc-slot pic-slot (MX2020 Routers with MPC5EQ and MPC6E)

```
user@host> show chassis pic fpc-slot 18 pic-slot 2
```

FPC slot 18, PIC slot 2 information:

```
Type                3X40GE QSFP
State                Online
PIC version          0.0
Uptime               6 minutes, 31 seconds
```

PIC port information:

		Fiber	Xcvr vendor		Wave-	Xcvr
Port	Cable type	type	Xcvr vendor	part number	length	

```

Firmware
 0  40GBASE SR4      MM  AVAGO          AFBR-79E4Z-D-JU2  850 nm  0.0
 1  40GBASE SR4      MM  AVAGO          AFBR-79E4Z-D-JU2  850 nm  0.0
 2  40GBASE SR4      MM  AVAGO          AFBR-79E4Z-D-JU2  850 nm  0.0

```

### show chassis pic fpc-slot pic-slot (MX2020 Routers with MPC6E and OTN MIC)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 0, PIC slot 1 information:
  Type                24X10GE SFPP OTN
  State                Online
  PIC version          1.1
  Uptime               1 hour, 33 minutes, 59 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	part number	Wave-length	Xcvr
7	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
9	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BNL-J1	850 nm	0.0
12	10GBASE LR	SM	FINISAR CORP.	FTLX1472M3BNL-J3	1310 nm	0.0
20	10GBASE ZR	SM	FINISAR CORP.	FTLX1871M3BNL-J3	1550 nm	0.0
21	10GBASE ER	SM	FINISAR CORP.	FTLX1671D3BTL-J4	1550 nm	0.0
22	10GBASE LR	SM	SOURCEPHOTONICS	SPP10SLREDFCJNP	1310 nm	0.0
23	10GBASE LR	SM	FINISAR CORP.	FTLX1471D3BNL-J1	1310 nm	0.0

### show chassis pic fpc-slot pic-slot (MX2020 Routers with MPC4E)

```

user@host> show chassis pic fpc-slot 14 pic-slot 0
FPC slot 14, PIC slot 2 information:
  Type                4x10GE SFPP
  State                Online
  PIC version          0.0
  Uptime               1 day, 14 hours, 49 minutes, 9 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	part number	Wave-length	Xcvr
0	10GBASE SR	MM	SumitomoElectric	SPP5100SR-J3	850 nm	0.0
1	10GBASE SR	MM	SumitomoElectric	SPP5100SR-J3	850 nm	0.0
3	10GBASE SR	MM	SumitomoElectric	SPP5100SR-J3	850 nm	0.0

### show chassis pic fpc-slot pic-slot (T1600 Router with 100-Gigabit Ethernet PIC)

```

user@host> run show chassis pic fpc-slot 3 pic-slot 1
FPC slot 3, PIC slot 1 information:
  Type                100GE SLOT1

```

```

ASIC type           Brooklyn 100GE FPGA
State               Online
PIC version         1.3
Uptime              10 minutes, 44 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	10GBASE LR4	SM	Opnext Inc.	TRC5E20ENFSF000F	1310 nm

### show chassis pic fpc-slot pic-slot lcc (TX Matrix Router)

```

user@host> show chassis pic fpc-slot 1 pic-slot 1 lcc 0
lcc0-re0:

```

-----

PIC fpc slot 1 pic slot 1 information:

```

Type               4x OC-3 SONET, SMIR
ASIC type          D chip
State              Online
PIC version         1.2
Uptime              5 days, 2 hours, 12 minutes, 8 seconds

```

### show chassis pic fpc-slot pic-slot lcc (TX Matrix Plus Router)

```

user@host> show chassis pic pic-slot 0 fpc-slot 8
lcc0-re0:

```

-----

FPC slot 8, PIC slot 0 information:

```

Type               1x 10GE(LAN/WAN)
State              Online
Uptime              2 hours, 46 minutes, 23 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	part number	Wavelength
0	10GBASE ZR	SM	Opnext Inc.	TRF7061BN-LF150	1550 nm
0	10GBASE ZR	SM	FINISAR CORP.	FTRX-1811-3-J2	1550 nm

### show chassis pic fpc-slot pic-slot (Next-Generation SONET/SDH SFP)

```

user@host> show chassis pic fpc-slot 4 pic-slot 0

```

FPC slot 4, PIC slot 0 information:

```

Type               4x OC-3 1x OC-12 SFP
ASIC type          D FPGA
State              Online
PIC version         1.3
Uptime              1 day, 50 minutes, 4 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC48 short reach	SM	FINISAR CORP.	FTRJ1321P18TL-J2	1310 nm
1	OC3 short reach	MM	QCP	TRPA03MM3BAS-JE	1310 nm
2	OC3 short reach	MM	QCP	TRXA03MM3BAS-JW	1310 nm
3	OC12 inter reach	SM	FINISAR CORP.	FTLF1322P18TR	1310 nm

### show chassis pic fpc-slot pic-slot (12-Port T1/E1)

```

user@host> show chassis pic fpc-slot 0 pic-slot 3

```

FPC slot 0, PIC slot 3 information:

```

Type                12x T1/E1 CE
State                Online
PIC version          1.1
CPU load average     1 percent
Interrupt load average 0 percent
Total DRAM size      128 MB
Memory buffer utilization 100 percent
Memory heap utilization 4 percent
Uptime               1 day, 22 hours, 28 minutes, 12 seconds
Internal Clock Synchronization Normal

```

#### show chassis pic fpc-slot pic-slot (4x CHOC3 SONET CE SFP)

user@host> show chassis pic fpc-slot 0 pic-slot 1

FPC slot 0, PIC slot 1 information:

```

Type                4x CHOC3 SONET CE SFP
State                Online
PIC version          1.3
CPU load average     1 percent
Interrupt load average 0 percent
Total DRAM size      128 MB
Memory buffer utilization 99 percent
Memory heap utilization 4 percent
Uptime               1 day, 22 hours, 55 minutes, 37 seconds
Internal Clock Synchronization Normal

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC3 short reach	MM	AVAGO	HFBR-57E0P-JU2	n/a
1	OC3 short reach	MM	AVAGO	HFBR-57E0P-JU2	n/a
3	OC3 long reach	SM	OPNEX INC	TRF5456AVLB314	1310 nm

#### show chassis pic fpc-slot pic-slot (SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

user@host> show chassis pic fpc-slot 0 pic-slot 0

FPC slot 0, PIC slot 0 information:

```

Type                MIC-3D-80C30C12-40C48
State                Online
PIC version          1.8
Uptime               3 days, 22 hours, 3 minutes, 50 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
1	OC12 inter reach	SM	FINISAR CORP	FTRJ1322P1BTR-J3	1310 nm
7	OC12 inter reach	SM	FINISAR CORP	FTRJ1322P1BTR-J3	1310 nm

Multirate Mode Enabled

#### show chassis pic fpc-slot pic-slot (8-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

user@host> show chassis pic fpc-slot 3 pic-slot 0

FPC slot 3, PIC slot 0 information:

```

Type                MIC-3D-8CHOC3-4CHOC12
State                Online
PIC version          1.9
Uptime               1 hour, 21 minutes, 24 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
------	------------	------------	-------------	-------------------------	------------

0	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
1	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
2	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J2	1310 nm
4	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
5	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
6	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
7	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm

#### show chassis pic fpc-slot pic-slot (4-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
```

FPC slot 5, PIC slot 0 information:

Type	MIC-3D-4CHOC3-2CHOC12
State	Online
PIC version	1.9
Uptime	1 hour, 21 minutes

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
1	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
2	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
3	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm

#### show chassis pic fpc-slot pic-slot (1-port OC192/STM64 MIC with XFP)

```
user@host> show chassis pic fpc-slot 1 pic-slot 0
```

FPC slot 1, PIC slot 0 information:

Type	MIC-3D-10C192-XFP
State	Online
PIC version	1.2
Uptime	1 day, 11 hours, 4 minutes, 6 seconds

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC192 short reach	n/a	FINISAR CORP.	FTLX1412M3BCL-J3	1310 nm

#### show chassis pic fpc-slot 1 pic-slot 2 (8-port DS3/E3 MIC)

```
user@host> show chassis pic fpc-slot 1 pic-slot 2
```

FPC slot 1, PIC slot 2 information:

Type	MIC-3D-8DS3-E3
State	Online
PIC version	1.10
Uptime	4 days, 1 hour, 29 minutes, 19 seconds
Channelization Mode	Disabled

#### show chassis pic fpc-slot pic-slot (OTN)

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
```

PIC fpc slot 5 pic slot 0 information:

Type	1x10GE(LAN),OTN
ASIC type	H chip
State	Online
PIC version	1.0
Uptime	5 minutes, 50 seconds

#### show chassis pic fpc-slot pic-slot (QFX3500 Switch)

```
user@switch> show chassis pic fpc-slot 0 pic-slot 0
```



```
FPC slot 0, PIC slot 0 information:
Type 48x 10G-SFP+ Builtin
State Online
Uptime 3 days, 3 hours, 5 minutes, 20 seconds
```

#### show chassis pic fpc-slot pic-slot (QFX5100 Standalone Switch)

```
user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
Type                               Unknown Builtin
State                              Online
Uptime                             1 day, 17 hours, 5 minutes, 9 seconds
```

#### show chassis pic interconnect-device fpc-slot pic-slot (QFabric Systems)

```
user@switch> show chassis pic interconnect-device interconnect1 fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
Type                               16x 40G-GE Builtin
State                              Online
Uptime                             2 hours, 47 minutes, 40 seconds
```

#### show chassis pic node-device fpc-slot pic-slot (QFabric System)

```
user@switch> show chassis pic node-device node1 pic-slot 0
FPC slot node1, PIC slot 0 information:
Type                               48x 10G-SFP+ Builtin
State                              Online
Uptime                             2 hours, 52 minutes, 37 seconds
```

#### PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
1	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
2	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
3	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
4	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
5	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
6	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
7	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
8	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
9	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
10	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
11	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
12	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
13	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
14	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
15	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
16	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
17	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
18	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
19	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
20	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
21	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
22	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
23	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
24	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
25	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
26	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
27	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
28	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
29	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm

30	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
31	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
32	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
33	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
34	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
35	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
36	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
37	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
38	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
39	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
40	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
41	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
42	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
43	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
44	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
45	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
46	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
47	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm

#### show chassis pic fpc-slot pic-slot (ACX2000 Universal Access Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
  Type                8x 1GE(LAN) RJ45 Built-in
  State                Online
  Uptime               6 days, 2 hours, 51 minutes, 11 seconds
```

#### show chassis pic fpc-slot pic-slot (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
  Type                AS-MSB
  State                Online
  PIC version          1.6
  Uptime               11 hours, 17 minutes, 56 seconds
```

#### show chassis pic FPC slot PIC slot (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis pic fpc-slot 1 pic-slot 2
Type                AS-MXC
State                Online
PIC version          1.0
Uptime               11 hours, 18 minutes, 3 seconds
```

#### show chassis pic transport fpc-slot pic-slot (PTX Series Packet Transport Routers)

```
user@host> show chassis pic transport fpc-slot 2 pic-slot 0
Administrative State: In Service
Operational State:   Normal
```

## show chassis routing-engine

---

**List of Syntax**    [Syntax on page 1227](#)  
                          [Syntax \(EX Series Switches\) on page 1227](#)  
                          [Syntax \(T Series routers\) on page 1227](#)  
                          [Syntax \(TX Matrix Routers\) on page 1227](#)  
                          [Syntax \(TX Matrix Plus Routers\) on page 1227](#)  
                          [Syntax \(QFX Series\) on page 1227](#)  
                          [Syntax \(MX Series Routers\) on page 1227](#)  
                          [Syntax \(MX2010 3D Universal Edge Routers\) on page 1227](#)  
                          [Syntax \(MX2020 3D Universal Edge Routers\) on page 1227](#)  
                          [Syntax \(MX104 3D Universal Edge Routers\) on page 1227](#)  
                          [Syntax \(ACX Series Universal Access Routers\) on page 1228](#)

**Syntax**    show chassis routing-engine  
                  <bios | *slot*>

**Syntax (EX Series Switches)**    show chassis routing-engine  
    <*slot*>

**Syntax (T Series routers)**    show chassis routing-engine  
    <bios | *slot*>

**Syntax (TX Matrix Routers)**    show chassis routing-engine  
    <bios | *slot*>  
    <lcc *number* | scc>

**Syntax (TX Matrix Plus Routers)**    show chassis routing-engine  
    <bios | *slot*>  
    <lcc *number* | sfc *number*>

**Syntax (QFX Series)**    show chassis routing-engine  
                                  <interconnect-device *name*>  
                                  <node-device *name*>

**Syntax (MX Series Routers)**    show chassis routing-engine  
                                  <bios | *slot*>  
                                  <all-members>  
                                  <local>  
                                  <member *member-id*>

**Syntax (MX2010 3D Universal Edge Routers)**    show chassis routing-engine  
                                  <bios | *slot*>

**Syntax (MX2020 3D Universal Edge Routers)**    show chassis routing-engine  
                                  <bios | *slot*>

**Syntax (MX104 3D Universal Edge Routers)**    show chassis routing-engine

**Syntax (ACX Series Universal Access Routers)** `show chassis routing-engine`

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
**sfc** option introduced for the TX Matrix Plus router in Junos OS Release in 9.6.  
Command introduced in Junos OS Release 11.1 for QFX Series.  
Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.  
Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.  
Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.  
Command introduced in Junos OS Release 13.2 for MX104 3D Univesral Edge Routers.

**Description** Display the status of the Routing Engine.

**Options** **none**—Display information about one or more Routing Engines. On a TX Matrix router, display information about all Routing Engines on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display information about all Routing Engines on the TX Matrix Plus router and its attached routers.

**all-members**—(MX Series routers only) (Optional) Display Routing Engine information for all members of the Virtual Chassis configuration.

**bios**—(Optional) Display the (BIOS) firmware version.

**interconnect-device *number***—(QFabric systems only) (Optional) Display Routing Engine information for a specified Interconnect device.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Routing Engine information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display Routing Engine information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display Routing Engine information for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display Routing Engine information for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-device *number***—(QFabric systems only) (Optional) Display Routing Engine information for a specified Node device.

**scc**—(TX Matrix routers only) (Optional) Display Routing Engine information for the TX Matrix router (switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display Routing Engine information for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

**slot**—(Systems with multiple Routing Engines) (Optional) Display information for an individual Routing Engine. Replace *slot* with 0 or 1. For QFX3500 switches, there is only one Routing Engine, so you do not need to specify the slot number.

**Required Privilege Level** view

**Related Documentation**

- [request chassis routing-engine master on page 96](#)
- *Configuring Routing Engine Redundancy*
- *Switching the Global Master and Backup Roles in a Virtual Chassis Configuration*

**List of Sample Output**

- [show chassis routing-engine \(M5 Router\) on page 1231](#)
- [show chassis routing-engine \(M10 Router\) on page 1232](#)
- [show chassis routing-engine \(M20 Router\) on page 1232](#)
- [show chassis routing-engine \(M40 Router\) on page 1233](#)
- [show chassis routing-engine \(M120 Router\) on page 1233](#)
- [show chassis routing-engine \(M160 Router\) on page 1234](#)
- [show chassis routing-engine \(MX104 Router\) on page 1234](#)
- [show chassis routing-engine \(MX240 Router\) on page 1235](#)
- [show chassis routing-engine \(MX480 Router\) on page 1236](#)
- [show chassis routing-engine \(MX960 Router\) on page 1236](#)
- [show chassis routing-engine \(MX2010 Router\) on page 1236](#)
- [show chassis routing-engine \(MX2020 Router\) on page 1237](#)
- [show chassis routing-engine \(T320 router\) on page 1238](#)
- [show chassis routing-engine \(T640 router\) on page 1239](#)
- [show chassis routing-engine \(T1600 router\) on page 1239](#)
- [show chassis routing-engine \(T4000 router\) on page 1240](#)
- [show chassis routing-engine \(TX Matrix Router\) on page 1241](#)
- [show chassis routing-engine lcc \(TX Matrix Router\) on page 1242](#)
- [show chassis routing-engine bios \(TX Matrix Router\) on page 1242](#)
- [show chassis routing-engine \(TX Matrix Plus Router\) on page 1243](#)
- [show chassis routing-engine lcc \(TX Matrix Plus Router\) on page 1244](#)
- [show chassis routing-engine bios \(TX Matrix Plus Router\) on page 1245](#)
- [show chassis routing-engine \(QFX Series\) on page 1245](#)
- [show chassis routing-engine interconnect-device \(QFabric systems\) on page 1245](#)
- [show chassis routing-engine \(PTX Series Packet Transport Switch\) on page 1246](#)

[show chassis routing-engine \(EX9200 Switch\) on page 1247](#)

[show chassis routing-engine \(ACX2000 Universal Access Router\) on page 1247](#)

[show chassis routing-engine \(ACX1000 Universal Access Router\) on page 1248](#)

**Output Fields** [Table 120 on page 1230](#) lists the output fields for the **show chassis routing-engine** command. Output fields are listed in the approximate order in which they appear.

**Table 120: show chassis routing-engine Output Fields**

Field Name	Field Description
<b>Slot</b>	(Systems with single and multiple Routing Engines) Slot number.
<b>Current state</b>	(Systems with multiple Routing Engines) Current state of the Routing Engine: <b>Master</b> , <b>Backup</b> , or <b>Disabled</b> .
<b>Election priority</b>	(Systems with multiple Routing Engines) Election priority for the Routing Engine: <b>Master</b> or <b>Backup</b> .
<b>Temperature</b>	Temperature of the air flowing past the Routing Engine.
<b>CPU Temperature</b>	Temperature of the CPU.
<b>DRAM</b>	Total DRAM available to the Routing Engine's processor.  Starting with Junos OS Release 12.3R1, the DRAM field displays both available memory and installed memory.
<b>Memory utilization</b>	Percentage of Routing Engine memory being used.
<b>CPU utilization</b>	Information about the Routing Engine's CPU utilization: <ul style="list-style-type: none"> <li>• <b>User</b>—Percentage of CPU time being used by user processes.</li> <li>• <b>Background</b>—Percentage of CPU time being used by background processes.</li> <li>• <b>Kernel</b>—Percentage of CPU time being used by kernel processes.</li> <li>• <b>Interrupt</b>—Percentage of CPU time being used by interrupts.</li> <li>• <b>Idle</b>—Percentage of CPU time that is idle.</li> </ul>
<b>Model</b>	Routing Engine model number.
<b>Serial ID</b>	(Systems with multiple Routing Engines) Identification number of the Routing Engine in this slot.
<b>Start time</b>	Time at which the Routing Engine started running.
<b>Uptime</b>	How long the Routing Engine has been running.
<b>Routing Engine BIOS Version</b>	BIOS version being run by the Routing Engine.

Table 120: show chassis routing-engine Output Fields (*continued*)

Field Name	Field Description
Last reboot reason	<p>Reason for last reboot, including:</p> <ul style="list-style-type: none"> <li><b>power cycle/failure</b>—Halt of the Routing Engine using the <b>halt</b> command, powering down using the power button on the chassis or any other method (such as removal of the control board or Routing Engine), and then powering back the Routing Engine. A halt of the operating system also occurs if you enter the <b>request system halt</b> command. You can enter this command to halt the system operations on the chassis or specific Routing Engines. To restart the software, press any key on the keyboard.</li> <li><b>watchdog</b>—Reboot due to a hardware watchdog. A watchdog is a hardware monitoring process that examines the health and performance of the router to enable the device to recover from failures. A watchdog checks for problems at certain intervals, and reboots the routing engine if a problem is encountered.</li> <li><b>reset-button reset</b>—(Not available on the J Series router or EX Series switch) Reboot due to pressing of the reset button on the Routing Engine.</li> <li><b>power-button hard power off</b>—Reboot due to pressing of the power button on the chassis. A powering down of the software also occurs if you enter the <b>request system power-off</b> command. You can enter this command to power down the chassis or specific Routing Engines; you can then restart the software.</li> <li><b>misc hardware reason</b>—Reboot due to miscellaneous hardware reasons.</li> <li><b>thermal shutdown</b>—Reboot due to the router or switch reaching a critical temperature at which point it is unsafe to continue operations.</li> <li><b>hard disk failure</b>—Reboot due to a hard disk or solid-state drive (SSD) failure.</li> <li><b>reset from debugger</b>—Reboot due to reset from the debugger.</li> <li><b>chassis control reset</b>—Restart the chassis process that manages PICs, FPCs, and other hardware components. The chassis control module that runs the Routing Engine performs management and monitoring functions, and it provides a single access point for operational and maintenance functions. A reset of the chassis management process occurs when you enter the <b>restart chassis-control</b> command.</li> <li><b>bios auto recovery reset</b>—Reboot due to a BIOS auto-recovery reset.</li> <li><b>could not be determined</b>—Reboot due to an undetermined reason.</li> <li><b>Router rebooted after a normal shutdown</b>—Reboot due to a normal shutdown. This reason is displayed if the Routing Engine is powered down by pushing and holding the online/offline button on the Routing Engine faceplate for 30 seconds, and then powered back. A reboot of the software also occurs if you enter the <b>request system reboot</b> command. You can enter this command to reboot the chassis or specific Routing Engines.</li> </ul>
Load averages	Routing Engine load averages for the last 1, 5, and 15 minutes.

## Sample Output

### show chassis routing-engine (M5 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature                25 degrees C / 77 degrees F
  DRAM                       768 MB
  Memory utilization         21 percent
  CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                  0 percent

```

Idle	100 percent
Model	RE-2.0
Serial ID	31000007349bf701
Start time	2003-12-04 09:42:17 PST
Uptime	26 days, 1 hour, 12 minutes, 27 seconds
Last reboot reason	Router rebooted after a normal shutdown
Load averages:	1 minute 5 minute 15 minute
	0.00 0.01 0.00

#### show chassis routing-engine (M10 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
  Temperature          25 degrees C / 77 degrees F
  DRAM                 768 MB
  Memory utilization    21 percent
  CPU utilization:
    User               0 percent
    Background         0 percent
    Kernel             0 percent
    Interrupt          0 percent
    Idle               100 percent
  Model                RE-2.0
  Serial ID            31000007349bf701
  Start time           2003-12-04 09:42:17 PST
  Uptime               26 days, 1 hour, 12 minutes, 27 seconds
  Last reboot reason    Router rebooted after a normal shutdown
  Load averages:      1 minute 5 minute 15 minute
                      0.00 0.01 0.00
```

#### show chassis routing-engine (M20 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state        Master
  Election priority     Master (default)
  Temperature          29 degrees C / 84 degrees F
  DRAM                 768 MB
  Memory utilization    20 percent
  CPU utilization:
    User               1 percent
    Background         0 percent
    Kernel             2 percent
    Interrupt          0 percent
    Idle               97 percent
  Model                RE-2.0
  Serial ID            58000007348d9a01
  Start time           2003-12-30 07:05:47 PST
  Uptime               3 hours, 41 minutes, 14 seconds
  Last reboot reason    Router rebooted after a normal shutdown
  Load averages:      1 minute 5 minute 15 minute
                      0.00 0.02 0.00

Routing Engine status:
Slot 1:
  Current state        Backup
  Election priority     Backup (default)
  Temperature          29 degrees C / 84 degrees F
  DRAM                 768 MB
  Memory utilization    0 percent
  CPU utilization:
```



```

User                0 percent
Background          0 percent
Kernel              1 percent
Interrupt           0 percent
Idle                99 percent
Model               RE-2.0
Serial ID           d800000734745701
Start time          2003-06-17 16:37:33 PDT
Uptime              195 days, 18 hours, 47 minutes, 9 seconds
Last reboot reason   Router rebooted after a normal shutdown

```

### show chassis routing-engine (M40 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature        25 degrees C / 77 degrees F
  DRAM               768 MB
  Memory utilization  21 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel            0 percent
    Interrupt         0 percent
    Idle              100 percent
  Model              RE-2.0
  Serial ID           31000007349bf701
  Start time          2003-12-04 09:42:17 PST
  Uptime              26 days, 1 hour, 12 minutes, 27 seconds
  Last reboot reason   Router rebooted after a normal shutdown
  Load averages:      1 minute   5 minute  15 minute
                      0.00        0.01    0.00

```

### show chassis routing-engine (M120 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority   Master (default)
  Temperature        46 degrees C / 114 degrees F
  CPU temperature     44 degrees C / 111 degrees F
  DRAM               2048 MB
  Memory utilization  18 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel            5 percent
    Interrupt         0 percent
    Idle              95 percent
  Model              RE-A-1000
  Serial ID           1000621154
  Start time          2006-10-31 17:10:05 PST
  Uptime              14 minutes, 31 seconds
  Last reboot reason   Router rebooted after a normal shutdown
  Load averages:      1 minute   5 minute  15 minute
                      0.02        0.07    0.07

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority   Backup (default)
  Temperature        45 degrees C / 113 degrees F

```

```
CPU temperature      42 degrees C / 107 degrees F
DRAM                2048 MB
Memory utilization   15 percent
CPU utilization:
  User               0 percent
  Background         0 percent
  Kernel             0 percent
  Interrupt          0 percent
  Idle               100 percent
Model               RE-A-1000
Serial ID            1000621151
Start time           2006-10-31 17:10:04 PST
Uptime               14 minutes, 30 seconds
Last reboot reason   Router rebooted after a normal shutdown
```

### show chassis routing-engine (M160 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority   Master (default)
  Temperature         43 degrees C / 109 degrees F
  DRAM               2048 MB
  Memory utilization  11 percent
  CPU utilization:
    User              1 percent
    Background        0 percent
    Kernel             2 percent
    Interrupt         0 percent
    Idle              97 percent
  Model              RE-3.0
  Serial ID           210865700403
  Start time          2003-12-23 12:25:55 PST
  Uptime              6 days, 22 hours, 33 minutes, 24 seconds
  Last reboot reason  Router rebooted after a normal shutdown
  Load averages:     1 minute   5 minute   15 minute
                     0.24       0.13       0.04

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority   Backup (default)
  Temperature         40 degrees C / 104 degrees F
  DRAM               2048 MB
  Memory utilization  9 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel             0 percent
    Interrupt         0 percent
    Idle              100 percent
  Model              RE-3.0
  Serial ID           210865700332
  Start time          2003-12-23 12:25:55 PST
  Uptime              6 days, 22 hours, 33 minutes, 21 seconds
  Last reboot reason  Router rebooted after a normal shutdown
```

### show chassis routing-engine (MX104 Router)

```
user@host> show chassis routing-engine
```

```

Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             32 degrees C / 89 degrees F
  CPU temperature         42 degrees C / 107 degrees F
  DRAM                   3840 MB (3840 MB installed)
  Memory utilization      18 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                3 percent
    Interrupt             2 percent
    Idle                  94 percent
  Model                   RE-MX-104
  Serial ID               CAAR5925
  Start time              2013-06-05 13:17:08 IST
  Uptime                  1 hour, 15 minutes, 8 seconds
  Last reboot reason      0x200:normal shutdown
  Load averages:         1 minute   5 minute   15 minute
                        0.87       0.90       0.41

Routing Engine status:
Slot 1:
  Current state           Backup
  Election priority       Backup (default)
  Temperature             32 degrees C / 89 degrees F
  CPU temperature         38 degrees C / 100 degrees F
  DRAM                   3840 MB (3840 MB installed)
  Memory utilization      13 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                1 percent
    Interrupt             2 percent
    Idle                  97 percent
  Model                   RE-MX-104
  Serial ID               CAAM6369
  Start time              2013-06-05 13:07:37 IST
  Uptime                  1 hour, 24 minutes, 34 seconds
  Last reboot reason      0x200:normal shutdown
  Load averages:         1 minute   5 minute   15 minute
                        0.19       0.15       0.06

```

### show chassis routing-engine (MX240 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Backup
  Election priority       Master (default)
  Temperature             40 degrees C / 104 degrees F
  CPU temperature         47 degrees C / 116 degrees F
  DRAM                   3584 MB
  Memory utilization      7 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                0 percent
    Interrupt             0 percent
    Idle                  100 percent
  Model                   RE-S-2000

```

Serial ID	1000703522
Start time	2007-12-19 10:35:40 PST
Uptime	16 days, 3 hours, 15 minutes, 23 seconds
Last reboot reason	Router rebooted after a normal shutdown

#### show chassis routing-engine (MX480 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             41 degrees C / 105 degrees F
  CPU temperature         38 degrees C / 100 degrees F
  DRAM                    2048 MB
  Memory utilization      13 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                2 percent
    Interrupt             0 percent
    Idle                  98 percent
  Model                   RE-S-1300
  Serial ID               1000697044
  Start time              2008-01-04 06:46:08 PST
  Uptime                  8 hours, 17 minutes, 16 seconds
  Last reboot reason      Router rebooted after a normal shutdown
```

#### show chassis routing-engine (MX960 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             37 degrees C / 98 degrees F
  CPU temperature         37 degrees C / 98 degrees F
  DRAM                    2048 MB
  Memory utilization      18 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                4 percent
    Interrupt             0 percent
    Idle                  96 percent
  Model                   RE-S-1300
  Serial ID               1000617944
  Start time              2006-10-26 12:37:13 PDT
  Uptime                  6 days, 4 hours, 59 minutes, 40 seconds
  Last reboot reason      Router rebooted after a normal shutdown
  Load averages:         1 minute   5 minute   15 minute
                        0.16       0.08       0.02
```

#### show chassis routing-engine (MX2010 Router)

```
user@host> show chassis routing-engine

Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             3 degrees C / 37 degrees F
```

```

CPU temperature      3 degrees C / 37 degrees F
DRAM                17152 MB
Memory utilization   13 percent
CPU utilization:
  User              0 percent
  Background        0 percent
  Kernel            4 percent
  Interrupt         2 percent
  Idle              95 percent
Model               RE-S-1800x4
Serial ID           9009099704
Start time          2012-10-02 14:33:32 PDT
Uptime              14 hours, 39 minutes, 39 seconds
Last reboot reason   Router rebooted after a normal shutdown.
Load averages:      1 minute  5 minute  15 minute
                    0.06      0.05      0.01

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority   Backup (default)
  Temperature        1 degrees C / 33 degrees F
  CPU temperature     2 degrees C / 35 degrees F
  DRAM               17152 MB
  Memory utilization  11 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel            0 percent
    Interrupt         0 percent
    Idle              100 percent
  Model              RE-S-1800x4
  Serial ID           9009099706
  Start time          2012-10-02 10:36:06 PDT
  Uptime              18 hours, 36 minutes, 57 seconds
  Last reboot reason   Router rebooted after a normal shutdown.
  Load averages:      1 minute  5 minute  15 minute
                    0.01      0.00      0.00

```

### show chassis routing-engine (MX2020 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority   Master (default)
  Temperature        6 degrees C / 42 degrees F
  CPU temperature     6 degrees C / 42 degrees F
  DRAM               17152 MB
  Memory utilization  14 percent
  CPU utilization:
    User              1 percent
    Background        0 percent
    Kernel            7 percent
    Interrupt         2 percent
    Idle              91 percent
  Model              RE-S-1800x4
  Serial ID           9009089704
  Start time          2012-10-02 11:05:24 PDT
  Uptime              2 days, 15 hours, 49 minutes, 13 seconds
  Last reboot reason   Router rebooted after a normal shutdown.
  Load averages:      1 minute  5 minute  15 minute

```

```

                                0.10      0.05      0.01
Routing Engine status:
Slot 1:
  Current state                Backup
  Election priority            Backup (default)
  Temperature                  7 degrees C / 44 degrees F
  CPU temperature              5 degrees C / 41 degrees F
  DRAM                        17152 MB
  Memory utilization           12 percent
  CPU utilization:
    User                       0 percent
    Background                 0 percent
    Kernel                     0 percent
    Interrupt                  0 percent
    Idle                       99 percent
  Model                       RE-S-1800x4
  Serial ID                    9009094138
  Start time                  2012-10-02 11:09:57 PDT
  Uptime                      2 days, 15 hours, 44 minutes, 27 seconds
  Last reboot reason          Router rebooted after a normal shutdown.
  Load averages:              1 minute   5 minute  15 minute
                                0.00      0.00      0.00

```

#### show chassis routing-engine (T320 router)

```

user@host> show chassis routing-engine
Slot 0:
  Current state                Master
  Election priority            Master (default)
  Temperature                  51 degrees C / 123 degrees F
  CPU temperature              55 degrees C / 131 degrees F
  DRAM                        3584 MB
  Memory utilization           11 percent
  CPU utilization:
    User                       0 percent
    Background                 0 percent
    Kernel                     2 percent
    Interrupt                  0 percent
    Idle                       97 percent
  Model                       RE-A-2000
  Serial ID                    9009010618
  Start time                  2012-10-10 01:24:05 PDT
  Uptime                      5 days, 10 hours, 49 minutes, 23 seconds
  Last reboot reason          0x1:power cycle/failure
  Load averages:              1 minute   5 minute  15 minute
                                0.00      0.05      0.04

Routing Engine status:
Slot 1:
  Current state                Backup
  Election priority            Backup (default)
  Temperature                  45 degrees C / 113 degrees F
  CPU temperature              48 degrees C / 118 degrees F
  DRAM                        3584 MB
  Memory utilization           9 percent
  CPU utilization:
    User                       0 percent
    Background                 0 percent
    Kernel                     0 percent
    Interrupt                  0 percent
    Idle                       100 percent
  Model                       RE-A-2000

```

```

Serial ID          9009003642
Start time         2012-10-10 01:24:04 PDT
Uptime            5 days, 10 hours, 49 minutes, 28 seconds
Last reboot reason 0x1:power cycle/failure

```

### show chassis routing-engine (T640 router)

```
user@host> show chassis routing-engine
```

```
Routing Engine status:
```

```
Slot 0:
```

```

Current state      Master
Election priority  Master (default)
Temperature        50 degrees C / 122 degrees F
CPU temperature    58 degrees C / 136 degrees F
DRAM              3584 MB
Memory utilization 14 percent
CPU utilization:
  User            1 percent
  Background      0 percent
  Kernel          4 percent
  Interrupt       1 percent
  Idle            95 percent
Model            RE-A-2000
Serial ID        1000686556
Start time       2012-10-10 01:24:02 PDT
Uptime          5 days, 10 hours, 50 minutes, 27 seconds
Last reboot reason 0x1:power cycle/failure
Load averages:   1 minute   5 minute   15 minute
                  1.24      0.33      0.12

```

```
Routing Engine status:
```

```
Slot 1:
```

```

Current state      Backup
Election priority  Backup (default)
Temperature        44 degrees C / 111 degrees F
CPU temperature    49 degrees C / 120 degrees F
DRAM              3584 MB
Memory utilization 12 percent
CPU utilization:
  User            0 percent
  Background      0 percent
  Kernel          0 percent
  Interrupt       1 percent
  Idle            99 percent
Model            RE-A-2000
Serial ID        1000702739
Start time       2012-10-10 01:24:02 PDT
Uptime          5 days, 10 hours, 50 minutes, 26 seconds
Last reboot reason 0x1:power cycle/failure

```

### show chassis routing-engine (T1600 router)

```
user@host> show chassis routing-engine
```

```
Routing Engine status:
```

```
Slot 0:
```

```

Current state      Master
Election priority  Master (default)
Temperature        48 degrees C / 118 degrees F
CPU temperature    58 degrees C / 136 degrees F
DRAM              3584 MB
Memory utilization 13 percent
CPU utilization:

```

```

User                0 percent
Background          0 percent
Kernel              3 percent
Interrupt            1 percent
Idle                96 percent
Model               RE-A-2000
Serial ID           1000704521
Start time           2012-10-10 01:23:41 PDT
Uptime              5 days, 10 hours, 46 minutes, 56 seconds
Last reboot reason   0x1:power cycle/failure
Load averages:       1 minute   5 minute   15 minute
                     0.05       0.03       0.01

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority   Backup (default)
  Temperature         44 degrees C / 111 degrees F
  CPU temperature     48 degrees C / 118 degrees F
  DRAM                3584 MB
  Memory utilization  12 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel             0 percent
    Interrupt          0 percent
    Idle              100 percent
  Model              RE-A-2000
  Serial ID           9009006579
  Start time           2012-10-10 01:23:42 PDT
  Uptime              5 days, 10 hours, 46 minutes, 54 seconds
  Last reboot reason   0x1:power cycle/failure

```

#### show chassis routing-engine (T4000 router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority   Master (default)
  Temperature         33 degrees C / 91 degrees F
  CPU temperature     50 degrees C / 122 degrees F
  DRAM                8960 MB
  Memory utilization  18 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel             4 percent
    Interrupt          1 percent
    Idle              95 percent
  Model              RE-DUO-1800
  Serial ID           P737F-002248
  Start time           2012-02-09 22:49:53 PST
  Uptime              2 hours, 21 minutes, 35 seconds
  Last reboot reason   Router rebooted after a normal shutdown.
  Load averages:       1 minute   5 minute   15 minute
                     0.00       0.04       0.00

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority   Backup (default)
  Temperature         32 degrees C / 89 degrees F

```



```

CPU temperature      46 degrees C / 114 degrees F
DRAM                8960 MB
Memory utilization   24 percent
CPU utilization:
  User               0 percent
  Background         0 percent
  Kernel             0 percent
  Interrupt          0 percent
  Idle              99 percent
Model               RE-DUO-1800
Serial ID           P737F-002653
Start time          2012-02-08 20:12:51 PST
Uptime              1 day, 4 hours, 58 minutes, 28 seconds
Last reboot reason   Router rebooted after a normal shutdown.

```

### show chassis routing-engine (TX Matrix Router)

```

user@host> show chassis routing-engine
scc-re0:

```

#### Routing Engine status:

##### Slot 0:

```

Current state      Master
Election priority   Master (default)
Temperature        34 degrees C / 93 degrees F
CPU temperature     33 degrees C / 91 degrees F
DRAM               2048 MB
Memory utilization  12 percent
CPU utilization:
  User             0 percent
  Background       0 percent
  Kernel           2 percent
  Interrupt        0 percent
  Idle            98 percent
Model             RE-4.0
Serial ID          P11123900153
Start time         2004-08-05 18:42:05 PDT
Uptime             9 days, 22 hours, 49 minutes, 50 seconds
Last reboot reason Router rebooted after a normal shutdown
Load averages:     1 minute   5 minute   15 minute
                   0.00      0.08      0.07

```

#### lcc0-re0:

#### Routing Engine status:

##### Slot 0:

```

Current state      Master
Election priority   Master (default)
Temperature        33 degrees C / 91 degrees F
CPU temperature     30 degrees C / 86 degrees F
DRAM               2048 MB
Memory utilization  12 percent
CPU utilization:
  User             0 percent
  Background       0 percent
  Kernel           1 percent
  Interrupt        0 percent
  Idle            98 percent
Model             RE-3.0
Serial ID          210865700363
Start time         2004-08-05 18:42:05 PDT

```

```

Uptime                9 days, 22 hours, 48 minutes, 20 seconds
Last reboot reason    Router rebooted after a normal shutdown
Load averages:        1 minute   5 minute   15 minute
                       0.00      0.02      0.00

```

```
lcc2-re0:
```

```
-----
Routing Engine status:
```

```
Slot 0:
```

```

Current state          Master
Election priority      Master (default)
Temperature            34 degrees C / 93 degrees F
CPU temperature        35 degrees C / 95 degrees F
DRAM                  2048 MB
Memory utilization     12 percent
CPU utilization:
  User                 0 percent
  Background           0 percent
  Kernel               2 percent
  Interrupt            0 percent
  Idle                 98 percent
Model                  RE-4.0
Serial ID              P11123900126
Start time             2004-08-05 18:42:05 PDT
Uptime                9 days, 22 hours, 49 minutes, 4 seconds
Last reboot reason    Router rebooted after a normal shutdown
Load averages:        1 minute   5 minute   15 minute
                       0.01      0.01      0.0

```

#### show chassis routing-engine lcc (TX Matrix Router)

```
user@host> show chassis routing-engine 0 lcc 0
```

```
lcc0-re0:
```

```
-----
Routing Engine status:
```

```
Slot 0:
```

```

Current state          Master
Election priority      Master (default)
Temperature            33 degrees C / 91 degrees F
CPU temperature        30 degrees C / 86 degrees F
DRAM                  2048 MB
Memory utilization     12 percent
CPU utilization:
  User                 0 percent
  Background           0 percent
  Kernel               1 percent
  Interrupt            0 percent
  Idle                 98 percent
Model                  RE-3.0
Serial ID              210865700363
Start time             2004-08-05 18:42:05 PDT
Uptime                7 days, 22 hours, 49 minutes, 6 seconds
Last reboot reason    Router rebooted after a normal shutdown
Load averages:        1 minute   5 minute   15 minute
                       0.00      0.00      0.00

```

#### show chassis routing-engine bios (TX Matrix Router)

```
user@host> show chassis routing-engine bios
```

```
scc-re0:
```

```
Routing Engine BIOS Version: V1.0.0
1cc0-re0:
```

```
-----
Routing Engine BIOS Version: V1.0.17
1cc2-re0:
```

```
-----
Routing Engine BIOS Version: V1.0.0
```

### show chassis routing-engine (TX Matrix Plus Router)

```
user@host> show chassis routing-engine
sfc0-re0:
```

```
-----
Routing Engine status:
```

Slot 0:

Current state	Master
Election priority	Master (default)
Temperature	27 degrees C / 80 degrees F
CPU temperature	42 degrees C / 107 degrees F
DRAM	3327 MB
Memory utilization	12 percent
CPU utilization:	
User	0 percent
Background	0 percent
Kernel	2 percent
Interrupt	0 percent
Idle	98 percent
Model	RE-TXP-SFC
Serial ID	737A-1024
Start time	2009-05-11 17:39:49 PDT
Uptime	3 hours, 45 minutes, 25 seconds
Last reboot reason	Router rebooted after a normal shutdown.
Load averages:	1 minute    5 minute    15 minute
	0.00            0.00            0.00

Routing Engine status:

Slot 1:

Current state	Backup
Election priority	Backup (default)
Temperature	29 degrees C / 84 degrees F
CPU temperature	43 degrees C / 109 degrees F
DRAM	3327 MB
Memory utilization	11 percent
CPU utilization:	
User	0 percent
Background	0 percent
Kernel	0 percent
Interrupt	0 percent
Idle	100 percent
Model	RE-TXP-SFC
Serial ID	737A-1024
Start time	2009-05-11 17:08:54 PDT
Uptime	4 hours, 16 minutes, 52 seconds
Last reboot reason	0x1:power cycle/failure

```
1cc0-re0:
```

```
-----
Routing Engine status:
```

Slot 0:

Current state	Master
Election priority	Master (default)
Temperature	30 degrees C / 86 degrees F

```

CPU temperature          43 degrees C / 109 degrees F
DRAM                    3327 MB
Memory utilization       9 percent
CPU utilization:
  User                   0 percent
  Background             0 percent
  Kernel                 2 percent
  Interrupt              0 percent
  Idle                   98 percent
Model                   RE-TXP-LCC
Serial ID                737F-1024
Start time              2009-05-11 17:40:32 PDT
Uptime                  3 hours, 44 minutes, 51 seconds
Last reboot reason      Router rebooted after a normal shutdown.
Load averages:          1 minute   5 minute   15 minute
                        0.00       0.00       0.00

Routing Engine status:
Slot 1:
  Current state          Backup
  Election priority      Backup (default)
  Temperature            30 degrees C / 86 degrees F
  CPU temperature        43 degrees C / 109 degrees F
  DRAM                   3327 MB
  Memory utilization     9 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               0 percent
    Interrupt            0 percent
    Idle                 100 percent
  Model                  RE-TXP-LCC
  Serial ID              737F-1024
  Start time             2009-05-06 17:31:32 PDT
  Uptime                  5 days, 3 hours, 54 minutes, 19 seconds
  Last reboot reason      Router rebooted after a normal shutdown.

```

#### show chassis routing-engine lcc (TX Matrix Plus Router)

```

user@host> show chassis routing-engine 0 lcc 0
1cc0-re0:
-----
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            30 degrees C / 86 degrees F
  CPU temperature        43 degrees C / 109 degrees F
  DRAM                   3327 MB
  Memory utilization     9 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               2 percent
    Interrupt            0 percent
    Idle                 98 percent
  Model                  RE-TXP-LCC
  Serial ID              737F-1024
  Start time             2009-05-11 17:40:32 PDT
  Uptime                  3 hours, 45 minutes, 26 seconds
  Last reboot reason      Router rebooted after a normal shutdown.
  Load averages:        1 minute   5 minute   15 minute

```

```

                                0.00      0.00      0.00
Routing Engine status:
Slot 1:
  Current state                Backup
  Election priority            Backup (default)
  Temperature                  30 degrees C / 86 degrees F
  CPU temperature              43 degrees C / 109 degrees F
  DRAM                        3327 MB
  Memory utilization           9 percent
  CPU utilization:
    User                       0 percent
    Background                 0 percent
    Kernel                     0 percent
    Interrupt                  0 percent
    Idle                       100 percent
  Model                        RE-TXP-LCC
  Serial ID                    737F-1024
  Start time                   2009-05-06 17:31:32 PDT
  Uptime                       5 days, 3 hours, 54 minutes, 59 seconds
  Last reboot reason           Router rebooted after a normal shutdown.

```

#### show chassis routing-engine bios (TX Matrix Plus Router)

```

user@host> show chassis routing-engine bios
sfc0-re0:

```

```

-----
Routing Engine BIOS Version: V0.0.Z

```

```

lcc0-re0:

```

```

-----
Routing Engine BIOS Version: V0.0.N

```

#### show chassis routing-engine (QFX Series)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state Master
  Election priority Master (default)
  DRAM 2820 MB
  Memory utilization 49 percent
  CPU utilization:
    User 1 percent
    Background 0 percent
    Kernel 1 percent
    Interrupt 0 percent
    Idle 97 percent
  Model QFX3500-48S4Q
  Serial ID S/N ED3709
  Uptime 3 days, 4 hours, 29 minutes, 42 seconds
  Last reboot reason 0x200:chassis control reset
  Load averages: 1 minute 5 minute 15 minute
0.37 0.26 0.19

```

#### show chassis routing engine interconnect-device (QFabric systems)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state                Master
  Election priority            Master (default)
  Temperature                  48 degrees C / 118 degrees F

```

```

DRAM                                3312 MB
Memory utilization                    63 percent
CPU utilization:
  User                               14 percent
  Background                         0 percent
  Kernel                             5 percent
  Interrupt                           0 percent
  Idle                               81 percent
Model                                RE-QFXC08-CB4S
Serial ID                            BUILTIN
Start time                           2011-07-06 13:26:15 UTC
Uptime                               11 hours, 24 minutes, 57 seconds
Last reboot reason                    0x4:reset-button reset
Load averages:                       1 minute   5 minute   15 minute
                                      2.62       2.31       2.28

Routing Engine status:
Slot 1:
  Current state                       Backup
  Election priority                   Backup (default)
  Temperature                         39 degrees C / 102 degrees F
  DRAM                                3312 MB
  Memory utilization                  59 percent
  CPU utilization:
    User                              9 percent
    Background                        0 percent
    Kernel                            1 percent
    Interrupt                          0 percent
    Idle                              91 percent
  Model                               RE-QFXC08-CB4S
  Serial ID                           BUILTIN
  Start time                           2011-07-06 13:24:58 UTC
  Uptime                               11 hours, 26 minutes, 18 seconds
  Last reboot reason                  0x4:reset-button reset

```

### show chassis routing-engine (PTX Series Packet Transport Switch)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state                       Master
  Election priority                   Master (default)
  Temperature                         60 degrees C / 140 degrees F
  CPU temperature                     76 degrees C / 168 degrees F
  DRAM                                17152 MB
  Memory utilization                  11 percent
  CPU utilization:
    User                              0 percent
    Background                        0 percent
    Kernel                            4 percent
    Interrupt                          0 percent
    Idle                              95 percent
  Model                               RE-DUO-2600
  Serial ID                           P737A-002231
  Start time                           2011-12-21 16:54:37 PST
  Uptime                               25 minutes, 44 seconds
  Last reboot reason                  Router rebooted after a normal shutdown.
  Load averages:                     1 minute   5 minute   15 minute
                                      0.01       0.02       0.06

Routing Engine status:
Slot 1:

```

```

Current state           Backup
Election priority       Backup (default)
Temperature             50 degrees C / 122 degrees F
CPU temperature         64 degrees C / 147 degrees F
DRAM                   17152 MB
Memory utilization      10 percent
CPU utilization:
  User                  0 percent
  Background            0 percent
  Kernel                0 percent
  Interrupt             0 percent
  Idle                  99 percent
Model                  RE-DU0-2600
Serial ID               P737A-002438
Start time              2011-12-21 16:52:26 PST
Uptime                  27 minutes, 49 seconds
Last reboot reason      Router rebooted after a normal shutdown.

```

### show chassis routing-engine (EX9200 Switch)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             35 degrees C / 95 degrees F
  CPU temperature         33 degrees C / 91 degrees F
  DRAM                   8157 MB
  Installed Memory        8192 MB
  Memory utilization      18 percent
CPU utilization:
  User                    1 percent
  Background              0 percent
  Kernel                  4 percent
  Interrupt               1 percent
  Idle                    94 percent
Model                    RE-S-EX9200-1800X4
Serial ID                 9009119555
Start time                2014-03-12 14:58:05 UTC
Uptime                    1 hour, 41 minutes, 51 seconds
Last reboot reason        Router rebooted after a normal shutdown.
Load averages:            1 minute  5 minute 15 minute
                        0.02      0.02    0.00

Routing Engine status:
Slot 1:
  Current state           Backup
  Election priority       Backup (default)

```

[...Output truncated...]

### show chassis routing-engine (ACX2000 Universal Access Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature             53 degrees C / 127 degrees F
  DRAM                   1536 MB
  Memory utilization      25 percent
CPU utilization:
  User                    0 percent
  Background              0 percent
  Kernel                  0 percent

```

Interrupt	1 percent
Idle	99 percent
Model	RE-ACX-2000
Start time	2012-05-09 00:57:07 PDT
Uptime	5 days, 3 hours, 16 minutes, 15 seconds
Last reboot reason	Router rebooted after a normal shutdown.
Load averages:	1 minute 5 minute 15 minute
	0.00 0.03 0.05

#### show chassis routing-engine (ACX1000 Universal Access Router)

user@host> show chassis routing-engine

##### Routing Engine status:

Temperature	36 degrees C / 96 degrees F
DRAM	768 MB
Memory utilization	50 percent
CPU utilization:	
User	3 percent
Background	0 percent
Kernel	6 percent
Interrupt	0 percent
Idle	91 percent
Model	RE-ACX-1000
Start time	2012-05-10 07:12:23 PDT
Uptime	4 days, 10 hours, 46 minutes, 53 seconds
Last reboot reason	Router rebooted after a normal shutdown.
Load averages:	1 minute 5 minute 15 minute
	0.00 0.00 0.00



## show log

<b>List of Syntax</b>	<a href="#">Syntax on page 1249</a> <a href="#">Syntax (QFabric System) on page 1249</a> <a href="#">Syntax (TX Matrix Routers) on page 1249</a>
<b>Syntax</b>	<pre>show log &lt;filename   user &lt;username&gt;&gt;</pre>
<b>Syntax (QFabric System)</b>	<pre>show log filename &lt;device-type (device-id   device-alias)&gt;</pre>
<b>Syntax (TX Matrix Routers)</b>	<pre>show log &lt;all-lcc   lcc number   scc&gt; &lt;filename   user &lt;username&gt;&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <i>device-type (device-id   device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series.</p>
<b>Description</b>	List log files, display log file contents, or display information about users who have logged in to the router or switch.
<b>Options</b>	<p><b>none</b>—List all log files.</p> <p><b>&lt;all-lcc   lcc number   scc&gt;</b>—(TX Matrix routers only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace <i>number</i> with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).</p> <p><b>device-type</b>—(QFabric system only) (Optional) Display log messages for only one of the following device types:</p> <ul style="list-style-type: none"> <li>• <b>director-device</b>—Display logs for Director devices.</li> <li>• <b>infrastructure-device</b>—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).</li> <li>• <b>interconnect-device</b>—Display logs for Interconnect devices.</li> <li>• <b>node-device</b>—Display logs for Node devices.</li> </ul>



**NOTE:** If you specify the *device-type* optional parameter, you must also specify either the *device-id* or *device-alias* optional parameter.

**(device-id | device-alias)**—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

**filename**—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.



**NOTE:** The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of messages.

**user <username>**—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include *username*, display logging information about the specified user.

**Required Privilege Level** trace

**List of Sample Output** [show log on page 1250](#)  
[show log filename on page 1250](#)  
[show log filename \(QFabric System\) on page 1251](#)  
[show log user on page 1251](#)

## Sample Output

**show log**

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin      211663 Oct  1 19:44 dcd
-rw-r--r-- 1 root bin      999947 Oct  1 19:41 dcd.0
-rw-r--r-- 1 root bin      999994 Oct  1 17:48 dcd.1
-rw-r--r-- 1 root bin      238815 Oct  1 19:44 rpd
-rw-r--r-- 1 root bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r-- 1 root bin      1061095 Oct  1 12:13 rpd.1
-rw-r--r-- 1 root bin      1052026 Oct  1 06:08 rpd.2
-rw-r--r-- 1 root bin      1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin      1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin      1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin      1056350 Sep 30 07:04 rpd.6
-rw-r--r-- 1 root bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin       19656 Oct  1 19:37 wtmp
```

**show log filename**

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT rcv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT rcv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT rcv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT rcv len 144 V9 seq 151 op chnge Type ifdev devindex 44
```

```

Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

```

### show log filename (QFabric System)

```

user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)

```

### show log user

```

user@host> show log user
darius  mg2546                Thu Oct  1 19:37   still logged in
darius  mg2529                Thu Oct  1 19:08 - 19:36 (00:28)
darius  mg2518                Thu Oct  1 18:53 - 18:58 (00:04)
root    mg1575                Wed Sep 30 18:39 - 18:41 (00:02)
root    ttyp2      jun.site.per Wed Sep 30 18:39 - 18:41 (00:02)
alex    ttyp1      192.168.1.2   Wed Sep 30 01:03 - 01:22 (00:19)

```

## show pfe next-hop

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1252</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 1252</a>
<b>Syntax</b>	<code>show pfe next-hop</code> <code>&lt;interface <i>interface-name</i>&gt;</code>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	<code>show pfe next-hop</code> <code>&lt;fpc <i>slot</i>&gt;</code> <code>&lt;interface <i>interface-name</i>&gt;</code> <code>&lt;lcc <i>number</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display Packet Forwarding Engine next-hop information.
<b>Options</b>	<p><b>none</b>—Display all Packet Forwarding Engine next-hop information.</p> <p><b>fpc <i>slot</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.</p> <ul style="list-style-type: none"><li>On a TX Matrix router, if you specify the number of a T640 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot</i></b> with a value from <b>0</b> through <b>7</b>. Otherwise, replace <b><i>slot</i></b> with a value from <b>0</b> through <b>31</b>.</li><li>On a TX Matrix Plus router, if you specify the number of a T1600 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot</i></b> with a value from <b>0</b> through <b>7</b>. Otherwise, replace <b><i>slot</i></b> with a value from <b>0</b> through <b>31</b>.</li><li>On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot</i></b> with a value from <b>0</b> through <b>7</b>. Otherwise, replace <b><i>slot</i></b> with a value from <b>0</b> through <b>63</b>.</li></ul> <p>For example, the following commands have the same result:</p> <pre>user@host&gt; show pfe next-hop fpc 1 lcc 1 user@host&gt; show pfe next-hop fpc 9</pre> <p><b>interface <i>interface-name</i></b>—(Optional) Display the Packet Forwarding Engine next-hop interface.</p> <p><b>lcc <i>number</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine next-hop interface for a specific T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine next-hop interface for the router (or line-card chassis) that is connected to a TX Matrix Plus router.</p>

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**Required Privilege Level** admin

**Related Documentation**

- *Routing Matrix with TXP-T1600 Configuration*
- *Routing Matrix with TXP-T1600-3D Configuration*
- *Routing Matrix with TXP-T4000-3D Configuration*
- *Routing Matrix with a TXP-Mixed-LCC-3D Configuration*

**List of Sample Output**

[show pfe next-hop on page 1254](#)  
[show pfe next-hop fpc \(TX Matrix Router\) on page 1254](#)  
[show pfe next-hop fpc \(TX Matrix Plus Router\) on page 1255](#)

**Output Fields** Table 121 on page 1254 lists the output fields for the **show pfe next hop** command. Output fields are listed in the approximate order in which they appear.

**Table 121: show pfe next-hop Output Fields**

Field Name	Field Description
ID	The next-hop ID for the entry.
Type	The next-hop type for the entry.
Interface	The interface to which the next-hop entry is assigned.
Protocol	The protocol type for the next-hop entry.
Encap	Encapsulation type for the next-hop entry.
Next Hop Addr	Next-hop address for the next-hop entry.
MTU	MTU value for the nexthop entry.

## Sample Output

### show pfe next-hop

```

user@host> show pfe next-hop
Nexthop Info:
  ID      Type      Interface      Protocol      Encap      Next Hop Addr      MTU
  ----      -      -      -      -      -      -
  4         Mcast      -              IPv4          -          0.0.0.0             0
  5         Bcast      -              IPv4          -          -                   0
  7         Discard     -              IPv4          -          -                   0
  8         MDiscard    -              IPv4          -          -                   0
  9         Reject      -              IPv4          -          -                   0
  13        Local      -              IPv4          -          192.168.4.60        0
  14        Resolve    fxp0.0         IPv4          Unspecified   -                   0
  17        Local      -              IPv4          -          127.0.0.1           0
  18        Unicast     fxp0.0         IPv4          Unspecified   192.168.4.254       0
  21        Local      -              IPv4          -          11.1.0.1            0
  22        Unicast     at-0/1/0.0     IPv4          ATM SNAP      11.1.0.2            4482
  ...

```

### show pfe next-hop fpc (TX Matrix Router)

```

user@host> show pfe next-hop fpc 1
Slot 1
Nexthop Info:
  ID      Type      Interface      Next Hop Addr      Protocol      Encap      MTU
  ----      -      -      -      -      -      -
  5         Mcast      -              default            IPv4          -          0
  6         Bcast      -              -                  IPv4          -          0
  8         Discard     -              -                  IPv4          -          0
  9         MDiscard    -              -                  IPv4          -          0
  13        Mcast      -              default            IPV6          -          0
  17        MDiscard    -              -                  IPV6          -          0
  18        Reject      -              -                  IPV6          -          0
  24        Discard     -              -                  None          -          0

```

```

68      Local - 192.168.66.113 IPv4 - 0
69      Resolve fxp0.0 - IPv4 Unspecified 0
70      Unicast fxp0.0 192.168.71.254 IPv4 Unspecified 0
256     Local - 10.71.71.1 IPv4 - 0
257     Local - 127.0.0.1 IPv4 - 0
258     Mcast.local..1 default IPv4 Unspecified 0
259     Bcast.local..1 - IPv4 Unspecified 0
261     Discard.local..1 - IPv4 Unspecified 0
262     MDiscard.local..1 - IPv4 Unspecified 0
269     Mcast.local..1 default IPV6 Unspecified 0
271     Discard.local..1 - IPV6 Unspecified 0
...

```

### show pfe next-hop fpc (TX Matrix Plus Router)

```
user@host> show pfe next-hop fpc 0
```

Slot 0

ID	Type	Interface	Next Hop Addr	Protocol	Encap	MTU
31	Mcast	-	default	IPv4	-	0
32	Bcast	-	-	IPv4	-	0
34	Discard	-	-	IPv4	-	0
35	MDiscard	-	-	IPv4	-	0
36	Reject	-	-	IPv4	-	0
39	Mcast	-	default	IPv6	-	0
42	Discard	-	-	IPv6	-	0
43	MDiscard	-	-	IPv6	-	0
44	Reject	-	-	IPv6	-	0
49	Receive	-	-	MPLS	-	0
50	Discard	-	-	MPLS	-	0
111	Mcast	.local..1	default	IPv4	Unspecified	0
112	Bcast	.local..1	-	IPv4	Unspecified	0
114	Discard	.local..1	-	IPv4	Unspecified	0
115	MDiscard	.local..1	-	IPv4	Unspecified	0
116	Reject	.local..1	-	IPv4	Unspecified	0
119	Mcast	.local..1	default	IPv6	Unspecified	0
122	Discard	.local..1	-	IPv6	Unspecified	0
123	MDiscard	.local..1	-	IPv6	Unspecified	0
124	Reject	.local..1	-	IPv6	Unspecified	0
191	Mcast	.local..2	default	IPv4	Unspecified	0
192	Bcast	.local..2	-	IPv4	Unspecified	0
194	Discard	.local..2	-	IPv4	Unspecified	0
195	MDiscard	.local..2	-	IPv4	Unspecified	0
196	Reject	.local..2	-	IPv4	Unspecified	0
322	Local	-	10.1.0.5	IPv4	-	0
323	Resolve	bcm0.0	-	IPv4	Unspecified	0
326	Local	-	129.0.0.5	IPv4	-	0
327	Resolve	bcm0.0	-	IPv4	Unspecified	0
328	Local	-	fe80::201:ff:fe01:5	IPv6	-	0
329	Receive	bcm0.0	ff02::1:ff01:5	IPv6	Unspecified	0
330	Receive	bcm0.0	fe80::	IPv6	Unspecified	0
331	Resolve	bcm0.0	-	IPv6	Unspecified	0
332	Local	-	fec0::a:1:0:5	IPv6	-	0
333	Receive	bcm0.0	ff02::1:ff00:5	IPv6	Unspecified	0
334	Receive	bcm0.0	fec0::	IPv6	Unspecified	0
335	Resolve	bcm0.0	-	IPv6	Unspecified	0
348	Local	-	192.168.178.4	IPv4	-	0
349	Resolve	em0.0	-	IPv4	Unspecified	0

350	Unicast	em0.0	192.168.178.126	IPv4	Unspecified	0
357	Local	-	fe80::201:1ff:fe01:5	IPv6	-	0
512	Local	-	10.255.178.11	IPv4	-	0
513	Local	-	127.0.0.1	IPv4	-	0
515	Local	-	abcd::10:255:178:11	IPv6	-	0
516	Local	-	fe80::200:ff:fe00:0	IPv6	-	0
517	Local	-	127.0.0.1	IPv4	-	0
518	Mcast	.local..3	default	IPv4	Unspecified	0
519	Bcast	.local..3	-	IPv4	Unspecified	0
521	Discard	.local..3	-	IPv4	Unspecified	0
522	MDiscard	.local..3	-	IPv4	Unspecified	0
523	Reject	.local..3	-	IPv4	Unspecified	0
531	Mcast	.local..3	default	IPv6	Unspecified	0
533	Discard	.local..3	-	IPv6	Unspecified	0
534	MDiscard	.local..3	-	IPv6	Unspecified	0
535	Reject	.local..3	-	IPv6	Unspecified	0
539	Mgroup	-	-	IPv4	-	0
540	Bcast	ge-15/0/3.0	-	IPv4	Ethernet	0
541	Receive	ge-15/0/3.0	14.2.1.0	IPv4	Ethernet	0
542	Local	-	14.2.1.1	IPv4	-	0
543	Resolve	ge-15/0/3.0	-	IPv4	Ethernet	0
544	Bcast	ge-31/0/4.0	-	IPv4	Ethernet	0
545	Receive	ge-31/0/4.0	14.1.1.0	IPv4	Ethernet	0
546	Local	-	14.1.1.1	IPv4	-	0
547	Resolve	ge-31/0/4.0	-	IPv4	Ethernet	0
548	Unicast	ge-31/0/4.0	14.1.1.2	IPv4	Ethernet	0
549	Unicast	ge-15/0/3.0	14.2.1.2	IPv4	Ethernet	0
550	Bcast	ae1.0	-	IPv4	Ethernet	0
551	Receive	ae1.0	11.1.1.0	IPv4	Ethernet	0
552	Local	-	11.1.1.1	IPv4	-	0
553	Resolve	ae1.0	-	IPv4	Ethernet	0
554	Aggreg.	ae1.0	-	IPv4	Ethernet	0
555	Unicast	ge-23/0/8.0	11.1.1.2	IPv4	Ethernet	0
556	Unicast	ge-7/0/9.0	11.1.1.2	IPv4	Ethernet	0
557	Aggreg.	ae1.0	-	MPLS	Ethernet	0
558	Unicast	ge-23/0/8.0	-	MPLS	Ethernet	0
559	Unicast	ge-7/0/9.0	-	MPLS	Ethernet	0
560	Aggreg.	ae1.0	-	MPLS	Ethernet	0
561	Unicast	ge-23/0/8.0	-	MPLS	Ethernet	0
562	Unicast	ge-7/0/9.0	-	MPLS	Ethernet	0



## show pfe route

<b>List of Syntax</b>	<a href="#">Syntax on page 1257</a> <a href="#">Syntax (EX Series Switches) on page 1257</a> <a href="#">Syntax (QFX Series) on page 1257</a> <a href="#">Syntax (MX Series) on page 1257</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 1257</a>
<b>Syntax</b>	<pre>show pfe route &lt;&lt;inet6   ip   iso&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;mpls&gt; &lt;summary&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show pfe route &lt;&lt;inet6   ip&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;mpls&gt; &lt;summary&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show pfe route &lt;&lt;inet6   ip&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt; &lt;hw (host   lpm   multicast)&gt;&gt; &lt;&lt;clnp&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;mpls&gt; &lt;summary&gt; &lt;hw&gt;</pre>
<b>Syntax (MX Series)</b>	<pre>show pfe route &lt;&lt;inet6   ip&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;dhcp&gt; &lt;mpls&gt; &lt;summary&gt;</pre>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	<pre>show pfe route &lt;fpc slot&gt; &lt;&lt;inet6   ip   iso&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;lcc number&gt; &lt;mpls&gt; &lt;summary&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.3 for the MX Series.</p> <p>Command option <b>hw</b> introduced in Junos OS Release 14.1X53-D10 for the QFX Series.</p>
<b>Description</b>	<p>Display the routes in the Packet Forwarding Engine forwarding table. The Packet Forwarding Engine forwards packets between input and output interfaces.</p>



**NOTE:** The Routing Engine maintains a master copy of the forwarding table. It copies the forwarding table to the Packet Forwarding Engine, which is the part of the router or switch responsible for forwarding packets. To display the routes in the Routing Engine forwarding table, use the **show route forwarding table** command. For more information, see the [CLI Explorer](#).

**Options** **none**—Display all Packet Forwarding Engine forwarding table information.

**clnp**—(Optional) Show International Standards Organization (ISO) connectionless-mode network protocol (CLNP) route table information.

**dhcp**—(Optional) Display Packet Forwarding Engine DHCP-Snooping route table information.

**fpc slot**—(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.

- On a TX Matrix router, if you specify the number of a T640 router by using the **lcc number** option (the recommended method), replace **slot** with a value from **0** through **7**. Otherwise, replace **slot** with a value from **0** through **31**.
- On a TX Matrix Plus router, if you specify the number of a T1600 router by using the **lcc number** option (the recommended method), replace **slot** with a value from **0** through **7**. Otherwise, replace **slot** with a value from **0** through **31**.
- On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the **lcc number** option (the recommended method), replace **slot** with a value from **0** through **7**. Otherwise, replace **slot** with a value from **0** through **63**.

For example, the following commands have the same result:

```
user@host> show pfe route fpc 1 lcc 1
user@host> show pfe route fpc 9
```

**host**—(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display host routes installed in the on-chip hardware table.

**hw**—(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display routes installed in the on-chip hardware table (as opposed to displaying routes from the routing table and the PFE forwarding table before they are installed in the hardware).

**index index**—(Optional) Display table index.

**inet6**—(Optional) Display Packet Forwarding Engine IPv6 routes.

**ip**—(Optional) Display Packet Forwarding Engine IPv4 routes.

**iso**—(Optional) Display ISO version routing tables.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, the slot number of the T640 router (or line-card chassis) that houses the FPC. On a TX Matrix Plus router, the slot number of the router (line-card chassis) that houses the FPC.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**lpm**—(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display longest prefix match (LPM) routes installed in the on-chip hardware table.

**mpls**—(Optional) Display Packet Forwarding Engine MPLS information.

**multicast**—(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display multicast routes installed in the on-chip hardware table.

**prefix *prefix***—(Optional) IPv4 or IPv6 prefix for which to show table entries.

**summary**—(Optional) Display summary of Packet Forwarding Engine information.

**table <*table-name*>**—(Optional) Display table information.

**Required Privilege Level**

admin

**Related Documentation**

- *Routing Matrix with TXP-T1600 Configuration*
- *Routing Matrix with TXP-T1600-3D Configuration*
- *Routing Matrix with TXP-T4000-3D Configuration*
- *Routing Matrix with a TXP-Mixed-LCC-3D Configuration*

**List of Sample Output**

[show pfe route ip on page 1261](#)  
[show pfe route iso on page 1261](#)  
[show pfe route lcc summary \(TX Matrix Router\) on page 1261](#)  
[show pfe route lcc summary \(TX Matrix Plus Router\) on page 1263](#)  
[show pfe route summary \(MX Series Router\) on page 1264](#)  
[show pfe route summary hw \(QFX Series\) on page 1265](#)

[show pfe route ip hw host \(QFX Series\) on page 1265](#)

**Output Fields** [Table 122 on page 1260](#) lists the output fields for the **show pfe route** command. Output fields are listed in the approximate order in which they appear.

**Table 122: show pfe route Output Fields**

Field Name	Field Description
<b>Destination</b>	Destination address for the entry.
<b>NH IP Addr</b>	Next-hop IP address for the entry.
<b>Type</b>	Next-hop type for the entry
<b>NH ID</b>	Next-hop ID for the entry
<b>Encap</b>	Encapsulation type for the next-hop entry.
<b>Interface</b>	Interface to which the next-hop entry is assigned.

[Table 123 on page 1260](#) lists the output fields for the QFX Series **show pfe route** hardware table (**hw**) commands. Output fields are listed in the approximate order in which they appear.

**Table 123: QFX Series show pfe route Hardware Table Output Fields**

Field Name	Field Description
<b>Max</b>	Maximum routing entries per route type.
<b>Used</b>	Number of routing entries consumed per route type.
<b>Free</b>	Number of unused routing entries per route type.
<b>% Free</b>	Percentage of unused routing entries per route type.
<b>Rtt</b>	Internal routing engine index number of the route table.
<b>VRF</b>	Internal hardware index number for the corresponding route table.
<b>Destination</b>	Destination address for the entry.
<b>Type</b>	( <b>show pfe route summary hw</b> )—Route type for the entry: IPv4 or IPv6 route, and host, LPM, or multicast route.  ( <b>show pfe route (ip   inet6) hw</b> )—Next-hop type for the entry.
<b>NH ID</b>	Next-hop ID for the entry
<b>Interface</b>	Interface to which the next-hop entry is assigned.

Table 123: QFX Series show pfe route Hardware Table Output Fields (*continued*)

Field Name	Field Description
HW NH-ID	Internal hardware index number of the next-hop.
Src-MAC-Address	Source MAC address.
Port	Port number.
Dst-MAC-Address	Destination MAC address.
VLAN	ID of the multicast group VLAN.
GROUP	Internal hardware index number of the multicast group next-hop.
CLASS	Internal class number of the multicast group.

## Sample Output

### show pfe route ip

```
user@host> show pfe route ip
```

```
IPv4 Route Table 0, default.0, 0x0:
```

Destination	NH IP Addr	Type	NH ID	Interface
default		Discard	8	
127.0.0.1	127.0.0.1	Local	256	
172.16/12	192.168.71.254	Unicast	68	fxp0.0
192.168.0/18	192.168.71.254	Unicast	68	fxp0.0
192.168.40/22	192.168.71.254	Unicast	68	fxp0.0
192.168.64/18	192.168.71.254	Unicast	68	fxp0.0
192.168.64/21		Resolve	67	fxp0.0
192.168.71.249	192.168.71.249	Local	66	
192.168.220.0/30		Resolve	303	fe-0/0/0.0
192.168.220.0	192.168.220.0	Receive	301	fe-0/0/0.0
224.0.0.1		Mcast	5	
255.255.255.255		Bcast	6	

```
...
```

### show pfe route iso

```
user@host# show pfe route iso
```

```
CLNS Route Table 0, CLNP.0, 0x0:
```

Destination	Type	NH ID	Interface
default	Reject	60	
47.0005.80ff.f800.0000.0108.0001.0102.5508.2159/152	Local	514	Local
49.0001.00a0.c96b.c491/72	Local	536	

### show pfe route lcc summary (TX Matrix Router)

```
user@host> show pfe route lcc 2 summary
```

## Slot 0

## IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	43	3081
1	4	281

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	68

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	9	717
1	5	389

## Slot 1

## IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	43	3081
1	4	281

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	68

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	9	717
1	5	389

## Slot 16

## IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	41	2938
1	4	281

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	68

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	9	717
1	5	389

## Slot 17

```
IPv4 Route Tables:
Index      Routes      Size(b)
-----
Default    41         2938
1          4          281
```

```
MPLS Route Tables:
Index      Routes      Size(b)
-----
Default    1           68
```

```
IPv6 Route Tables:
Index      Routes      Size(b)
-----
Default    9           717
1          5           389
```

#### show pfe route lcc summary (TX Matrix Plus Router)

```
user@host> show pfe route lcc 2 summary
```

Slot 0

```
IPv4 Route Tables:
Index      Routes      Size(b)
-----
Default    25         2266
1          9          815
2          6          545
3          5          453
4         15         1371
5          5          453
6         13         1187
```

```
MPLS Route Tables:
Index      Routes      Size(b)
-----
Default    1           88
4          5          452
```

```
IPv6 Route Tables:
Index      Routes      Size(b)
-----
Default    7           697
1         13         1305
3          4          385
4          4          385
5          4          385
6         18         1833
```

Slot 6

```
IPv4 Route Tables:
Index      Routes      Size(b)
-----
Default    25         2266
1          9          815
```

2	6	545
3	5	453
4	15	1371
5	5	453
6	13	1187

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	88
4	5	452

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	7	697
1	13	1305
3	4	385
4	4	385
5	4	385
6	18	1833
...		

**show pfe route summary (MX Series Router)**

```
user@host> show pfe route summary
```

Slot 0

## DHCP-Snooping Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	144

## IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	25	2266
1	9	815
2	6	545
3	5	453
4	15	1371
5	5	453
6	13	1187

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	88
4	5	452

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	7	697
1	13	1305
3	4	385
4	4	385
5	4	385
6	18	1833



...

**show pfe route summary hw (QFX Series)**

```

user@switch> show pfe route summary hw
Slot 0
Unit: 0
Profile active: 12-profile-three
Type          Max      Used      Free      % free
-----
IPv4 Host      8192     103       8073     98.55
IPv4 LPM       16384     9       16369     99.91
IPv4 Mcast     4096      2       4037     98.56

IPv6 Host      4096      6       4037     98.56
IPv6 LPM(< 64) 8192      3       8185     99.91
IPv6 LPM(> 64) 256       1    255     99.61
IPv6 Mcast     2048      0       2019     98.58

```

**show pfe route ip hw host (QFX Series)**

```

user@switch> show pfe route ip host hw
Slot 0
Unit: 0
IPv4 Host entries present: 103
Rtt  VRF  Destination          Type      NH-ID  Interface
      HW NH-ID  Src-MAC-Address      Port  Dst-MAC-Address
-----
4    3    255.255.255.255      Bcast    1695   .local.    .4
ifl 550 100003  00:00:00:01:02:03  127  00:00:00:01:02:03
0    1    200.1.1.42          Unicast   1743   et-0/1/1   .0
ifl 559 100268  84:18:88:de:96:fd  53    00:00:00:21:12:23
0    1    200.1.1.56          Unicast   1743   et-0/1/1   .0
ifl 559 100268  84:18:88:de:96:fd  53    00:00:00:21:12:23
0    1    200.1.1.61          Unicast   1743   et-0/1/1   .0
ifl 559 100268  84:18:88:de:96:fd  53    00:00:00:21:12:23
0    1    11.1.1.2            Unicast   1743   et-0/1/1   .0
ifl 559 100268  84:18:88:de:96:fd  53    00:00:00:21:12:23
0    1    200.1.1.73          Unicast   1743   et-0/1/1   .0
ifl 559 100268  84:18:88:de:96:fd  53    00:00:00:21:12:23
0    1    200.1.1.76          Unicast   1743   et-0/1/1   .0
ifl 559 100268  84:18:88:de:96:fd  53    00:00:00:21:12:23
0    1    200.1.1.18          Unicast   1743   et-0/1/1   .0
ifl 559 100268  84:18:88:de:96:fd  53    00:00:00:21:12:23
0    1    200.1.1.5           Unicast   1743   et-0/1/1   .0
ifl 559 100268  84:18:88:de:96:fd  53    00:00:00:21:12:23
0    1    200.1.1.23          Unicast   1743   et-0/1/1   .0
ifl 559 100268  84:18:88:de:96:fd  53    00:00:00:21:12:23
0    1    101.1.1.255         Bcast    1664   ae0        .0
ifl 544 100003  00:00:00:01:02:03  127  00:00:00:01:02:03
0    1    200.1.1.40          Unicast   1743   et-0/1/1   .0
ifl 559 100268  84:18:88:de:96:fd  53    00:00:00:21:12:23
0    1    200.1.1.58          Unicast   1743   et-0/1/1   .0
ifl 559 100268  84:18:88:de:96:fd  53    00:00:00:21:12:23. . .
. . .

```

## show pfe terse

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1266</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Router) on page 1266</a> <a href="#">Syntax (MX Series Router) on page 1266</a>
<b>Syntax</b>	show pfe terse
<b>Syntax (TX Matrix and TX Matrix Plus Router)</b>	show pfe terse <lcc <i>number</i>   scc> <sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show pfe terse <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display Packet Forwarding Engine status information.
<b>Options</b>	<b>none</b> —Display brief information about the Packet Forwarding Engine.  <b>all-members</b> —(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for all members in the Virtual Chassis configuration.  <b>lcc <i>number</i></b> —(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine information for a T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine information for the router (or line-card chassis) that is connected to a TX Matrix Plus router. Replace <i>number</i> with the following values depending on the LCC configuration: <ul style="list-style-type: none"><li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li><li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li><li>• 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li><li>• 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li></ul> <b>local</b> —(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**scc**—(TX Matrix routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix Plus router (or switch-fabric chassis).

**Required Privilege Level** admin

**List of Sample Output** [show pfe terse \(TX Matrix Router\) on page 1267](#)  
[show pfe terse \(TX Matrix Plus Router\) on page 1267](#)  
[show pfe terse sfc \(TX Matrix Plus Router\) on page 1267](#)

## Sample Output

### show pfe terse (TX Matrix Router)

```
user@host> show pfe terse
Slot Type Slot State Flags Uptime
0 SFM Present Online 0x0bf 01:25:42
2 SFM Present Online 0x0bf 01:25:40
0 FPC Present Online 0x102 01:25:57
1 FPC Present Online 0x102 01:25:55
2 FPC Present Online 0x102 01:25:53
```

### show pfe terse (TX Matrix Plus Router)

```
user@host> show pfe terse
sfc0-re0:
-----
Slot Type Slot State Uptime
0 LCC Present Online 2d 05:26

lcc0-re0:
-----
Slot Type Slot State Uptime
0 GFPC Present Online 2d 05:25
1 GFPC Present Online 2d 05:25
```

### show pfe terse sfc (TX Matrix Plus Router)

```
user@host> show pfe terse sfc 0
sfc0-re0:
-----
Slot Type Slot State Uptime
0 LCC Present Online 2d 05:25
```

## show system alarms

<b>Syntax</b>	show system alarms
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Display active system alarms.
<b>Options</b>	This command has no options.
<b>Additional Information</b>	<p>System alarms are preset. They include a <i>configuration</i> alarm that appears when no rescue configuration alarm is set and a <i>license</i> alarm that appears when a software feature is configured and no valid license is configured for the feature. On EX6200 switches, an alarm can be triggered by an internal link error. For more information about system alarms, see the <i>Junos OS Administration Library for Routing Devices</i>.</p> <p>In Junos OS release 11.1 and later, alarms for fans also show the slot number of the malfunctioning fans in the CLI output.</p> <p>Starting with Junos OS Release 13.2, you can view degraded fabric alarms on a routing matrix based on TX Matrix Plus router with 3D SIBs. The alarm indicates that the source FPC is running with a degraded fabric condition. This alarm is an early warning of a possible fabric black-hole condition. When the degraded fabric alarm is raised on the source FPC, you can take remedial action to avoid a fabric black-hole condition. The degraded fabric alarm is raised on the source FPC if both the following conditions are met:</p> <ul style="list-style-type: none"> <li>• The active Packet Forwarding Engine destinations are reachable on one or no active switching planes.</li> <li>• At least one of the inactive switching planes has a fault that causes the destination Packet Forwarding Engine to become unreachable.</li> </ul>
<b>Required Privilege Level</b>	admin
<b>List of Sample Output</b>	<p><a href="#">show system alarms on page 1269</a></p> <p><a href="#">show system alarms (Fan Tray) on page 1269</a></p> <p><a href="#">show system alarms (QFX Series) on page 1269</a></p> <p><a href="#">show system alarms (EX6200) on page 1269</a></p> <p><a href="#">show system alarms (TX Matrix Plus router with 3D SIBs) on page 1269</a></p>
<b>Output Fields</b>	Table 124 on page 1268 lists the output fields for the <b>show system alarms</b> command. Output fields are listed in the approximate order in which they appear.

**Table 124: show system alarms Output Fields**

Field Name	Field Description
Alarm time	Date and time the alarm was first recorded.

Table 124: show system alarms Output Fields (*continued*)

Field Name	Field Description
<b>Class</b>	Severity class for this alarm: <b>Minor</b> or <b>Major</b> .
<b>Description</b>	Information about the alarm.

## Sample Output

### show system alarms

```

user@host> show system alarms
2 alarms currently active
Alarm time          Class    Description
2005-02-24 17:29:34 UTC  Minor    IPsec VPN tunneling usage requires a
license
2005-02-24 17:29:34 UTC  Minor    Rescue configuration is not sent

```

### show system alarms (Fan Tray)

```

user@host> show system alarms
4 alarms currently active
Alarm time          Class    Description
2010-11-11 20:27:38 UTC  Major    Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC  Minor    Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 0 Failure

```

### show system alarms (QFX Series)

```

user@switch> show system alarms
2 alarms currently active
Alarm time Class Description
2005-02-24 17:29:34 UTC Minor Rescue configuration is not sent

```

### show system alarms (EX6200)

```

user@switch> show system alarms
2 alarms currently active
Alarm time          Class    Description
2013-04-05 16:51:41 PDT  Major    FPC 8 internal link errors detected
2013-04-04 18:05:35 PDT  Minor    Rescue configuration is not set

```

### show system alarms (TX Matrix Plus router with 3D SIBs)

```

user@router> show system alarms

sfc0-re0:
-----
2 alarms currently active
Alarm time          Class    Description
2013-05-08 18:13:58 UTC  Major    LCC 0 Major Errors
2013-05-08 17:48:46 UTC  Major    LCC 7 Major Errors

lcc0-re1:
-----
1 alarm currently active
Alarm time          Class    Description

```

2013-05-08 18:19:24 UTC Major FPC 1 degraded fabric condition detected

lcc7-re0:

-----  
1 alarm currently active

Alarm time	Class	Description
------------	-------	-------------

2013-05-08 18:19:24 UTC	Major	FPC 7 degraded fabric condition detected
-------------------------	-------	--

## show system audit

<b>List of Syntax</b>	<a href="#">Syntax on page 1271</a> <a href="#">Syntax (EX Series Switch and MX Series Router) on page 1271</a> <a href="#">Syntax (TX Matrix Router) on page 1271</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1271</a> <a href="#">Syntax (QFX Series) on page 1271</a>
<b>Syntax</b>	show system audit <root-only>
<b>Syntax (EX Series Switch and MX Series Router)</b>	show system audit <all-members> <local> <member <i>member-id</i> > <root-only>
<b>Syntax (TX Matrix Router)</b>	show system audit <all-lcc   lcc <i>number</i>   scc> <root-only>
<b>Syntax (TX Matrix Plus Router)</b>	show system audit <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> > <root-only>
<b>Syntax (QFX Series)</b>	show system audit <infrastructure <i>name</i>   interconnect-device <i>name</i>   node-group <i>name</i>   root-only>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Display the state and checksum values for file systems.
<b>Options</b>	<p><b>none</b>—Display the state and checksum values for all file systems.</p> <p><b>all-chassis</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display file system MD5 hash and permissions information for all of the chassis.</p> <p><b>all-lcc</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display file system MD5 hash and permissions information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display file system MD5 hash and permissions information for all T1600 or T4000 routers connected to the TX Matrix Plus router.</p> <p><b>all-members</b>—(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on all members of the Virtual Chassis configuration.</p> <p><b>lcc <i>number</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display file system MD5 hash and permissions information for a specific T640 router</p>

that is connected to the TX Matrix router. On a TX Matrix Plus router, display file system MD5 hash and permissions information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**infrastructure *name***—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for a fabric control Routing Engine or a fabric control Routing Engine.

**interconnect-device *name***—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for the Interconnect device.

**local**—(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on the local Virtual Chassis member.

**member *member-id***—(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for the Node group

**root-only**—(Optional) Check only the root (/) file system. On a QFabric system, you can check the root (/) file system on the infrastructure (fabric manager Routing Engine and fabric control Routing Engine), Interconnect device, or Node group.

**scc**—(TX Matrix routers only) (Optional) Display file system MD5 hash and permissions information for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display file system MD5 hash and permissions information for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

**Additional Information** To redirect the output to a file, issue the following command:

***ssh device-name 'show system audit root-only' > output-file***



If you save the output of the **show system audit root-only** command to a file, you can compare it to subsequent output from the command to determine whether anything has changed.

By default, when you issue the **show system audit** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level	admin
List of Sample Output	<a href="#">show system audit root-only on page 1273</a> <a href="#">show system audit lcc (TX Matrix Router) on page 1274</a> <a href="#">show system audit lcc (TX Matrix Plus Router) on page 1275</a> <a href="#">show system audit root-only (QFX3500 Switch) on page 1277</a>

## Sample Output

### show system audit root-only

```

user@host> show system audit root-only
#          user: root
#          machine: my-host
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.          type=dir nlink=23 size=1024 time=950252640.0
.cshrc     uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
           md5digest=f414e06fea6bd646244b98e13d6e6226
.kernel.jkernel.backup \
           mode=0744 size=1934552 time=944688902.0 \
           md5digest=2c343cf0bd9fea8f04f78604feed7aa4
.profile   uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
           md5digest=55a1e3c6c67789c9d3a1cce1ea39f670
COPYRIGHT  uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \
           md5digest=7df8bc77dcee71382ea73eb0ec6a9243
boot.config mode=0644 size=3 time=945902618.0 \
           md5digest=93d722493ed38477338a1405d7dcbb40
boot.help  uid=3 gid=7 mode=0444 size=411 time=939182876.0 \
           md5digest=9b7126385734bcae753f4179ab59d8e5
compat     type=link mode=0777 size=11 time=915149058.0 \
           link=/usr/compat
kernel     mode=0444 size=1947607 time=950230892.0 \
           md5digest=1a2a8aff2fec678a918ba0d6bf063980
kernel.avr uid=1112 size=1947642 time=950252597.0 \
           md5digest=82e1637682d58ec28964dfce7fccb62e
kernel.config \
           mode=0644 size=0 time=915149058.0 \
           md5digest=d41d8cd98f00b204e9800998ecf8427e
sys        type=link mode=0777 size=11 time=915149029.0 \
           link=/usr/src/sys

```

## show system audit lcc (TX Matrix Router)

```

user@host> show system audit lcc 2
lcc2-re0:
-----
#       user: root
#       machine: rodin-lcc2
#       tree: /
#       date: Mon Sep 13 11:55:33 2004

# .
/set type=file uid=0 gid=0 mode=0555 nlink=1 flags=none
.      type=dir nlink=20 size=512 time=1094982121.0
  COPYRIGHT mode=0644 size=4735 time=986012708.0 \
    md5digest=78396df1404ad742e6eb1be28f0cd63b
  kernel type=link mode=0700 size=17 time=1090266262.0 \
    link=/packages/jkernel

# ./altconfig
altconfig type=dir nlink=2 size=512 time=1089801320.0
# ./altconfig
..

# ./altroot
altroot type=dir nlink=2 size=512 time=1089801320.0
# ./altroot
..

# ./b
b type=dir mode=0755 nlink=2 size=512 time=1093961429.0
# ./b
..

# ./bin
/set type=file uid=0 gid=0 mode=0700 nlink=1 flags=none
bin type=dir mode=0755 nlink=2 size=512 time=1089843059.0
  [ type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/test
  cat type=link size=27 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/cat
  chmod type=link size=29 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/chmod
  cp type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/cp
  csh type=link size=27 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/csh
  date type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/date
  dd type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/dd
  df type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/df
  echo type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/echo
  ed type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/ed
  expr type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/expr
  hostname type=link size=32 time=1090266270.0 \

```

```

link=/packages/mnt/jbase/bin/hostname
kill      type=link size=28 time=1090266270.0 \
link=/packages/mnt/jbase/bin/kill
ln         type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/ln
ls         type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/ls
mkdir     type=link size=29 time=1090266270.0 \
link=/packages/mnt/jbase/bin/mkdir
mv         type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/mv
ps         type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/ps
pwd        type=link size=27 time=1090266270.0 \
link=/packages/mnt/jbase/bin/pwd
rcp        type=link size=27 time=1090266270.0 \
link=/packages/mnt/jbase/bin/rcp
red        type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/red
rm         type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/rm
rmdir     type=link size=29 time=1090266270.0 \
link=/packages/mnt/jbase/bin/rmdir
sh         type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/sh
sleep     type=link size=29 time=1090266270.0 \
link=/packages/mnt/jbase/bin/sleep
stty      type=link size=28 time=1090266270.0 \
link=/packages/mnt/jbase/bin/stty
sync      type=link size=28 time=1090266270.0 \
link=/packages/mnt/jbase/bin/sync
tcsh      type=link size=27 time=1090266270.0 \
link=/packages/mnt/jbase/bin/csh
test      type=link size=28 time=1090266270.0 \
link=/packages/mnt/jbase/bin/test
# ./bin
..

# ./boot
/set type=file uid=0 gid=0 mode=0444 nlink=1 flags=none
boot      type=dir mode=0555 nlink=3 size=512 time=1095069935.0
boot0     size=512 time=1094978286.0 \
md5digest=6f780822dd4ae482a20462b66e542cca
boot1     mode=0555 size=512 time=1094978294.0 \
md5digest=8d112b09df342cd0b60fdb9bdcde8e07
boot2     mode=0555 size=7680 time=1094978294.0 \
md5digest=28eb58c4068c6b85717e1484f9e028e4
cdboot    mode=0555 size=165888 time=1094978298.0 \
md5digest=1474c6b800dfc82ba552d7c36116d07d
kgzldr.o  size=5996 time=1094982121.0 \
md5digest=c53dc948eb07e2ea4eb0413e4c4634a3
loader    mode=0555 size=163840 time=1094978298.0 \
md5digest=82d9dc2d31033476bfb61bb7264c4fed
loader.4th size=9237 time=986013631.0 \
md5digest=43144391465ad50267d31e0a320be1de
...

```

#### show system audit lcc (TX Matrix Plus Router)

```
user@host> show system audit all-chassis
```

```

sfc0-re0:
-----
#       user: root
#       machine: finalfive
#       tree: /
#       date: Mon May 18 00:13:16 2009

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1 flags=none
.      type=dir nlink=23 size=512 time=1242347096.0
  COPYRIGHT mode=0644 size=6196 time=1168587741.0 \
    md5digest=bbad415e1c29bbdd9b383537100412c
    kernel type=link size=17 time=1242347011.0 link=/packages/jkernel
    staging type=link mode=0777 size=8 time=1242346935.0 link=/var/tmp

# ./snap
.snap type=dir mode=0775 nlink=2 size=512 time=1242346922.0
# ./snap
..

# ./altconfig
altconfig type=dir mode=0500 nlink=2 size=512 time=1242319843.0
# ./altconfig
..

# ./altroot
altroot type=dir mode=0500 nlink=2 size=512 time=1242319843.0
# ./altroot
..

# ./bin
bin type=dir nlink=2 size=512 time=1242346944.0
  \133 type=link size=28 time=1242346942.0 \
    link=/packages/mnt/jbase/bin/test
  cat type=link size=27 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/cat
  chflags type=link size=31 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/chflags
  chmod type=link size=29 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/chmod
  cp type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/cp
  csh type=link size=27 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/csh
  date type=link size=28 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/date
  dd type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/dd
  df type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/df
  echo type=link size=28 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/echo
  ed type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/ed
  expr type=link size=28 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/expr
  hostname type=link size=32 time=1242346941.0 \

```

```

kill      link=/packages/mnt/jbase/bin/hostname
          type=link size=28 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/kill
ln        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ln
ls        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ls
mkdir     type=link size=29 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/mkdir
mv        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/mv
pax       type=link size=27 time=1242346944.0 \
          link=/packages/mnt/jbase/bin/pax
ps        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ps
pwd       type=link size=27 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/pwd
rcp       type=link size=27 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rcp
red       type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/red
rm        type=link size=26 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rm
rmdir     type=link size=29 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rmdir
sh        type=link size=26 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sh
sleep     type=link size=29 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sleep
stty      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/stty
sync      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sync
tcsh      type=link size=27 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/csh
test      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/test
# ./bin
...

```

#### show system audit root-only (QFX3500 Switch)

```

user@switch> show system audit root-only
#          user: root
#          machine: my-host
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.          type=dir nlink=23 size=1024 time=950252640.0
.cshrc     uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
          md5digest=f414e06fea6bd646244b98e13d6e6226
.kernel.jkernel.backup \
          mode=0744 size=1934552 time=944688902.0 \
          md5digest=2c343cf0bd9fea8f04f78604feed7aa4
.profile   uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
          md5digest=55a1e3c6c67789c9d3a1cce1ea39f670
COPYRIGHT  uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \
          md5digest=7df8bc77dcee71382ea73eb0ec6a9243
boot.config mode=0644 size=3 time=945902618.0 \

```

```

boot.help      md5digest=93d722493ed38477338a1405d7dcbb40
                uid=3 gid=7 mode=0444 size=411 time=939182876.0 \
                md5digest=9b7126385734bcae753f4179ab59d8e5
compat         type=link mode=0777 size=11 time=915149058.0 \
                link=/usr/compat
kernel         mode=0444 size=1947607 time=950230892.0 \
                md5digest=1a2a8aff2fec678a918ba0d6bf063980
kernel.avr     uid=1112 size=1947642 time=950252597.0 \
                md5digest=82e1637682d58ec28964dfee7fccb62e
kernel.config \
                mode=0644 size=0 time=915149058.0 \
                md5digest=d41d8cd98f00b204e9800998ecf8427e
sys            type=link mode=0777 size=11 time=915149029.0 \
                link=usr/src/sys
    
```

## show system buffers

<b>List of Syntax</b>	<a href="#">Syntax on page 1279</a> <a href="#">Syntax (EX Series) on page 1279</a> <a href="#">Syntax (TX Matrix Router) on page 1279</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1279</a> <a href="#">Syntax (MX Series Router) on page 1279</a> <a href="#">Syntax (QFX Series) on page 1279</a>
<b>Syntax</b>	show system buffers
<b>Syntax (EX Series)</b>	show system buffers <all-members> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	show system buffers <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system buffers <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system buffers <all-members> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show system buffers <infrastructure <i>name</i>   interconnect-device <i>name</i>   node-group <i>name</i>   root-only (infrastructure <i>name</i>   interconnect-device <i>name</i>   node-group <i>name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about the buffer pool that the Routing Engine uses for local traffic. Local traffic is the routing and management traffic that is exchanged between the Routing Engine and the Packet Forwarding Engine within the router or switch, as well as the routing and management traffic from IP (that is, from OSPF, BGP, SNMP, ping operations, and so on).
<b>Options</b>	<b>none</b> —Show all buffer statistics.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show buffer statistics for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, show buffer statistics for all routers connected to the TX Matrix Plus router.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show buffer statistics for all of the chassis.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for all members of the Virtual Chassis configuration.

**infrastructure *name***—(QFabric systems only) (Optional) Show buffer statistics for a fabric control Routing Engine or a fabric control Routing Engine.

**interconnect-device *name***—(QFabric systems only) (Optional) Show buffer statistics for the Interconnect device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show buffer statistics for a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, show buffer statistics for a specific router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Show buffer statistics for the Node group

**sfc**—(TX Matrix Plus routers only) (Optional) Show buffer statistics for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system buffers** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

A special type of memory buffer called a *cluster* is 2 KB in size. For more information, see *The Design and Implementation of the 4.4BSD Operation System* by McKusic, Bostic, Karels, and Quarterman.



Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li></ul>
List of Sample Output	<a href="#">show system buffers on page 1282</a> <a href="#">show system buffers scc (TX Matrix Router) on page 1283</a> <a href="#">show system buffers sfc (TX Matrix Plus Router) on page 1283</a> <a href="#">show system buffers all-chassis (TX Matrix Plus Router) on page 1283</a> <a href="#">show system buffers node-group (QFabric System) on page 1284</a>
Output Fields	<a href="#">Table 125 on page 1282</a> describes the output fields for the <b>show system buffers</b> command. Output fields are listed in the approximate order in which they appear.

Table 125: show system buffers Output Fields

Field Name	Field Description
<b>mbufs in use</b>	Memory buffers (mbufs) are 128-byte buffers that are used for various purposes inside the kernel. Each memory buffer has a type, and the output itemizes the amount allocated for each type. Types with no memory buffers allocated are not displayed.
<b>mbufs allocated to packet headers</b>	Number of memory buffers currently holding packet headers
<b>mbufs allocated to control blocks</b>	Number of memory buffers currently holding the state for sockets.
<b>mbufs allocated to send data</b>	Number of memory buffers currently holding socket send data.
<b>mbufs allocated to pfe refill data</b>	Number of memory buffers currently holding Packet Forwarding Engine refill data.
<b>mbufs allocated to fxp data</b>	Number of memory buffers currently holding fxp data.
<b>mbufs allocated to socket names and addresses</b>	Number of memory buffers currently holding addresses for sockets.
<b>mbuf clusters in use</b>	Allocation statistics for memory buffer clusters.
<b>allocated to network</b>	Total amount of memory in use by the networking and interprocess communication (IPC) code.
<b>requests for memory denied</b>	Number of times a memory allocation request within the IPC and networking code failed.
<b>requests for memory delayed</b>	Number of times a memory allocation request within the IPC and networking code was postponed.
<b>calls to protocol drain routines</b>	Number of times a memory allocation request within the IPC and networking code triggered a memory reclamation attempt.

## Sample Output

### show system buffers

```

user@host> show system buffers
397/893/1290 mbufs in use (current/cache/total)
395/331/726/30000 mbuf clusters in use (current/cache/total/max)
384/256 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
889K/885K/1774K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/5/1024 sfbufs in use (current/peak/max)

```

```

0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

```

#### show system buffers scc (TX Matrix Router)

```

user@host> show system buffers scc
213 mbufs in use:
    11 mbufs allocated to packet headers
    26 mbufs allocated to socket names and addresses
    2 mbufs allocated to socket options
    17 mbufs allocated to socket send data
    2 mbufs allocated to pfe data
    155 mbufs allocated to fxp data (rx)
    511 mbufs allocated to <mbuf type 86>
    256 mbufs allocated to <mbuf type 92>
924/1162 mbuf clusters in use
2788 Kbytes allocated to network (75% in use)
0 requests for memory denied
0 requests for memory delayed
0 calls to protocol drain routines

```

#### show system buffers sfc (TX Matrix Plus Router)

```

user@host> show system buffers sfc 0

sfc0-re0:
-----
4363/2807/7170 mbufs in use (current/cache/total)
4358/1968/6326/30000 mbuf clusters in use (current/cache/total/max)
256/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
9806K/4637K/14444K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/10/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

```

#### show system buffers all-chassis (TX Matrix Plus Router)

```

user@host> show system buffers all-chassis

sfc0-re0:
-----
4363/2807/7170 mbufs in use (current/cache/total)
4358/1968/6326/30000 mbuf clusters in use (current/cache/total/max)
256/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
9806K/4637K/14444K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/10/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed

```

```
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc0-re0:
```

```
-----
772/2558/3330 mbufs in use (current/cache/total)
772/598/1370/30000 mbuf clusters in use (current/cache/total/max)
768/512 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1737K/1835K/3572K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc1-re0:
```

```
-----
773/2437/3210 mbufs in use (current/cache/total)
773/453/1226/30000 mbuf clusters in use (current/cache/total/max)
768/384 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1739K/1515K/3254K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/7/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc2-re0:
```

```
-----
816/2514/3330 mbufs in use (current/cache/total)
816/554/1370/30000 mbuf clusters in use (current/cache/total/max)
768/512 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1836K/1736K/3572K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
```

#### show system buffers node-group (QFabric System)

```
user@switch> show system buffers node-group node1
node-group node1:
```

```
-----
2/2698/2700 mbufs in use (current/cache/total)
2/1520/1522/30000 mbuf clusters in use (current/cache/total/max)
0/1280 mbuf+clusters out of packet secondary zone in use (current/cache)
```

```

0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
4K/3714K/3719K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/6/6656 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

```

re0:

```

-----
516/639/1155 mbufs in use (current/cache/total)
515/147/662/30000 mbuf clusters in use (current/cache/total/max)
512/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1159K/453K/1612K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

```

re1:

```

-----
519/771/1290 mbufs in use (current/cache/total)
518/176/694/30000 mbuf clusters in use (current/cache/total/max)
512/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1165K/544K/1710K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines

```

## show system connections

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1286</a> <a href="#">Syntax (EX Series) on page 1286</a> <a href="#">Syntax (TX Matrix Router) on page 1286</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1286</a> <a href="#">Syntax (MX Series Router) on page 1286</a> <a href="#">Syntax (QFX Series) on page 1286</a>
<b>Syntax</b>	<code>show system connections</code> <code>&lt;extensive&gt;</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt;</code> <code>&lt;inet   inet6&gt;</code> <code>&lt;show-routing-instances&gt;</code>
<b>Syntax (EX Series)</b>	<code>show system connections</code> <code>&lt;extensive&gt;</code> <code>&lt;all-members&gt;</code> <code>&lt;inet   inet6&gt;</code> <code>&lt;local&gt;</code> <code>&lt;member <i>member-id</i>&gt;</code> <code>&lt;show-routing-instances&gt;</code>
<b>Syntax (TX Matrix Router)</b>	<code>show system connections</code> <code>&lt;extensive&gt;</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt;</code> <code>&lt;inet   inet6&gt;</code> <code>&lt;show-routing-instances&gt;</code>
<b>Syntax (TX Matrix Plus Router)</b>	<code>show system connections</code> <code>&lt;extensive&gt;</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt;</code> <code>&lt;inet   inet6&gt;</code> <code>&lt;show-routing-instances&gt;</code>
<b>Syntax (MX Series Router)</b>	<code>show system connections</code> <code>&lt;extensive&gt;</code> <code>&lt;all-members&gt;</code> <code>&lt;inet   inet6&gt;</code> <code>&lt;local&gt;</code> <code>&lt;member <i>member-id</i>&gt;</code> <code>&lt;show-routing-instances&gt;</code>
<b>Syntax (QFX Series)</b>	<code>show system connections</code> <code>&lt;extensive&gt;</code> <code>&lt;inet&gt;</code> <code>&lt;infrastructure <i>name</i>&gt;</code> <code>&lt;interconnect-device <i>name</i>&gt;</code> <code>&lt;node-group <i>name</i>&gt;</code> <code>&lt;show-routing-instances&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.

**sfc** option introduced for the TX Matrix Plus router in Junos OS Release 9.6.  
Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display information about the active IP sockets on the Routing Engine. Use this command to verify which servers are active on a system and what connections are currently in progress.

**Options** **none**—Display information about all active IP sockets on the Routing Engine.

**extensive**—(Optional) Display exhaustive system process information, which, for TCP connections, includes the TCP control block. This option is useful for debugging TCP connections.

**all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system connection activity for all the routers in the chassis.

**all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system connection activity for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system connection activity for all connected T1600 or T4000 LCCs

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for all members of the Virtual Chassis configuration.

**inet | inet6**—(Optional) Display IPv4 connections or IPv6 connections, respectively.

**infrastructure name**—(QFabric systems only) (Optional) Display system connection activity for the fabric control Routing Engines or fabric manager Routing Engines.

**interconnect-device name**—(QFabric systems only) (Optional) Display system connection activity for the Interconnect device.

**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system connection activity for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system connection activity for a specific router that is connected to the TX Matrix Plus router. Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display system connection activity for the Node group.

**scc**—(TX Matrix routers only) (Optional) Display system connection activity for the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix routers only) (Optional) Display system connection activity for the TX Matrix Plus router.

**show-routing-instances**—(Optional) Display routing instances.

**Additional Information** By default, when you issue the **show system connections** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation** • [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [show system connections on page 1289](#)  
[show system connections extensive on page 1289](#)  
[show system connections lcc \(TX Matrix Router\) on page 1291](#)  
[show system connections show-routing-instances on page 1291](#)  
[show system connections \(TX Matrix Plus Router\) on page 1292](#)  
[show system connections sfc \(TX Matrix Plus Router\) on page 1295](#)  
[show system connections show-routing-instances \(TX Matrix Plus Router\) on page 1297](#)  
[show system connections \(QFX3500 Switch\) on page 1302](#)

**Output Fields** [Table 126 on page 1288](#) describes the output fields for the **show system connections** command. Output fields are listed in the approximate order in which they appear.

**Table 126: show system connections Output Fields**

Field Name	Field Description
<b>Proto</b>	Protocol of the socket: IP, TCP, or UDP for IPv4 or IPv6.
<b>Recv-Q</b>	Number of input packets received by the protocol and waiting to be processed by the application.
<b>Send-Q</b>	Number of output packets sent by the application and waiting to be processed by the protocol.



Table 126: show system connections Output Fields (*continued*)

Field Name	Field Description
Local Address	Local address and port of the socket, separated by a period. An asterisk (*) indicates that the bound address is the wildcard address. Server sockets typically have the wildcard address and a well-known port bound to them.
Foreign Address	Foreign address and port of the socket, separated by a period. An asterisk (*) indicates that the address or port is a wildcard.
Routing Instance	(Displayed only when the <b>show-routing-instance</b> option is used.) Routing instances associated with active IP sockets on the Routing Engine.
(state)	For TCP, the protocol state of the socket.

## Sample Output

### show system connections

```

user@host> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address          Foreign Address         (state)
tcp      0      2 192.168.4.16.513       208.197.169.254.894    ESTABLISHED
tcp      0      0 192.168.4.16.513       208.197.169.195.945    ESTABLISHED
tcp      0      0 *.23                   *.*                     LISTEN
tcp      0      0 *.22                   *.*                     LISTEN
tcp      0      0 *.513                  *.*                     LISTEN
tcp00 *.514             *.*                     LISTEN
tcp 0 0*.21                   *.*                     LISTEN
tcp00 *.79              *.*                     LISTEN
tcp 00 *.1023                *.*                     LISTEN
tcp 00 *.111                 *.*                     LISTEN
udp00192.168.4.16.1634   208.197.169.249.2049
udp00192.168.4.16.1627   208.197.169.254.2049
udp00192.168.4.16.1371   208.197.169.195.2049
udp00*.*                 *.*
udp00*.9999              *.*
udp00 *.161             *.*
udp00192.168.4.16.1039   192.168.4.16.1023
udp00192.168.4.16.1038   192.168.4.16.1023
udp 00 192.168.4.16.1037     192.168.4.16.1023
udp00192.168.4.16.1036   192.168.4.16.1023
udp00*.1022              *.*
udp00*.1023              *.*
udp00*.111               *.*
udp00*.                  *.*

```

### show system connections extensive

```

user@host> show system connections extensive

Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address          Foreign Address         (state)
tcp4      0      6 192.168.187.15.23

```

```

172.27.133.138.3013 ESTABLISHED
sndsbcc: 6 sndsbmbcnt: 256 sndsbmbmax: 272000
sndsblowat: 2048 sndsbhiwat: 34000
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 533120
rcvsblowat: 1 rcvsbhiwat: 66640
proc id: 0 proc name:
iss: 2566994072 sndup: 2566994491
snduna: 2566994491 sndnxt: 2566994494 sndwnd: 64094
sndmax: 2566994494 sndcwnd: 6589 sndsssthresh: 2720
irs: 236981199 rcvup: 236981325
rcvnxt: 236981327 rcvadv: 237046862 rcvwnd: 66640
rtt: 140058623 srtt: 15519 rttv: 908
rxtcur: 1200 rxtshift: 0 rtseq: 2566994491
rttmin: 1000 mss: 1360
flags: SACK_PERMIT [0x2000200]
tcp4 0 0 10.255.165.93.179
10.255.165.203.65141 ESTABLISHED
sndsbcc: 0 sndsbmbcnt: 0 sndsbmbmax: 131072
sndsblowat: 2048 sndsbhiwat: 16384
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 131072
rcvsblowat: 1 rcvsbhiwat: 16384
proc id: 0 proc name:
iss: 2555995917 sndup: 2555995917
snduna: 2555995917 sndnxt: 2555995917 sndwnd: 16384
sndmax: 2555995917 sndcwnd: 1000 sndsssthresh: 1073725440
irs: 2123825753 rcvup: 2123860681
rcvnxt: 2123860681 rcvadv: 2123877065 rcvwnd: 16384
rtt: 0 srtt: 3309 rttv: 72
rxtcur: 1200 rxtshift: 0 rtseq: 2555995898
rttmin: 1000 mss: 500
flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x3e0]
tcp4 0 0 10.255.165.203.65141
10.255.165.93.179 ESTABLISHED
sndsbcc: 0 sndsbmbcnt: 0 sndsbmbmax: 131072
sndsblowat: 2048 sndsbhiwat: 16384
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 131072
rcvsblowat: 1 rcvsbhiwat: 16384
proc id: 5022 proc name: rpd
iss: 2123825753 sndup: 2123860662
snduna: 2123860681 sndnxt: 2123860681 sndwnd: 16384
sndmax: 2123860681 sndcwnd: 1000 sndsssthresh: 1073725440
irs: 2555995917 rcvup: 2555995917
rcvnxt: 2555995917 rcvadv: 2556012301 rcvwnd: 16384
rtt: 0 srtt: 3279 rttv: 22
rxtcur: 1200 rxtshift: 0 rtseq: 2123860662
rttmin: 1000 mss: 500
flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x100003e0]
tcp4 0 0 10.255.165.203.179
10.255.165.113.52404 ESTABLISHED
sndsbcc: 0 sndsbmbcnt: 0 sndsbmbmax: 131072
sndsblowat: 2048 sndsbhiwat: 16384
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 131072
rcvsblowat: 1 rcvsbhiwat: 16384
proc id: 0 proc name:
iss: 1109297190 sndup: 1109332099
snduna: 1109332118 sndnxt: 1109332118 sndwnd: 16384
sndmax: 1109332118 sndcwnd: 1000 sndsssthresh: 1073725440
irs: 1476831634 rcvup: 1476866449
rcvnxt: 1476866449 rcvadv: 1476882833 rcvwnd: 16384
rtt: 0 srtt: 3235 rttv: 18
rxtcur: 1200 rxtshift: 0 rtseq: 1109332099

```

```

rttmin:      1000  mss:      500
flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x3e0]

```

### show system connections lcc (TX Matrix Router)

```
user@host> show system connections lcc 2
```

```
lcc2-re0:
```

```
-----
Active Internet connections (including servers)
```

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	(state)
tcp4	0	0	192.168.66.131.1342	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.2059	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.4571	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.2496	192.168.66.130.23	ESTABLISHED
tcp4	0	0	*.3221	*.*	LISTEN
tcp4	0	0	*.23	*.*	LISTEN
tcp4	0	0	*.22	*.*	LISTEN
tcp4	0	0	*.514	*.*	LISTEN
tcp4	0	0	*.513	*.*	LISTEN
tcp4	0	0	*.21	*.*	LISTEN
tcp4	0	0	*.79	*.*	LISTEN
tcp4	0	0	*.6234	*.*	LISTEN
udp4	0	0	*.514	*.*	
udp4	0	0	*.6333	*.*	

### show system connections show-routing-instances

```
user@host> show system connections show-routing-instances
```

```
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address Foreign Address Routing Instance
(state)
```

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	Routing Instance
tcp4	0	0	192.168.69.204.23	172.17.28.19.4267	default
			ESTABLISHED		
tcp4	0	0	192.168.69.204.58540	10.209.7.138.23	default
			ESTABLISHED		
tcp4	0	0	192.168.69.204.23	172.17.28.19.1098	default
			ESTABLISHED		
tcp4	0	0	192.168.7.1.57668	192.168.9.1.179	default
			ESTABLISHED		
tcp4	0	0	192.168.7.1.179	192.168.8.1.49209	default
			ESTABLISHED		
tcp4	0	0	128.0.0.1.6234	128.0.3.17.1024	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	128.0.0.4.9000	128.0.0.4.59103	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	128.0.0.4.59103	128.0.0.4.9000	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	*.32012	*.*	
__juniper_private1__			LISTEN		
tcp4	0	0	*.9000	*.*	
__juniper_private1__			LISTEN		
tcp4	0	0	*.33007	*.*	
__juniper_private2__			LISTEN		
tcp46	0	0	*.179	*.*	default
			LISTEN		
tcp4	0	0	*.179	*.*	default
			LISTEN		
tcp4	0	0	*.6154	*.*	
__juniper_private1__			LISTEN		
tcp4	0	0	*.6153	*.*	

```

__juniper_private1__ LISTEN
tcp4      0      0 *.7000          *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.6152          *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.6156          *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.33005         *.*
__juniper_private2__ LISTEN
tcp4      0      0 *.31343         *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.31341         *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.32003         *.*
__juniper_private2__ LISTEN
tcp4      0      0 *.666           *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.38            *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.3221          *.*
LISTEN                                           default

```

#### show system connections (TX Matrix Plus Router)

```

user@host> show system connections
sfc0-re0:

```

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      3 192.168.178.11.23
172.17.28.19.3565        ESTABLISHED
tcp4      0      0 192.168.178.11.23
172.17.28.204.62719      ESTABLISHED
tcp4      0      0 192.168.178.11.23
192.168.69.199.51255     ESTABLISHED
tcp4      0      0 192.168.178.11.23
172.24.26.227.42860     ESTABLISHED
tcp4      0      0 *.6156            *.*
LISTEN
tcp4      0      0 162.0.0.4.32012
162.0.0.5.58935         ESTABLISHED
tcp4      0      0 *.32012           *.*
LISTEN
tcp4      0      0 *.33007           *.*
LISTEN
tcp4      0      0 *.666             *.*
LISTEN
tcp4      0      0 162.0.0.4.6161
162.0.0.5.62026         ESTABLISHED
tcp4      0      0 *.33005           *.*
LISTEN
tcp4      0      0 162.0.0.4.9000
162.0.0.4.51611         ESTABLISHED
tcp4      0      0 162.0.0.4.51611
162.0.0.4.9000         ESTABLISHED
tcp4      0      0 *.6151            *.*
LISTEN
tcp4      0      0 *.6154            *.*
LISTEN
tcp4      0      0 *.6153            *.*

```

```

tcp4      0      0 *.31343      LISTEN      *. *
tcp4      0      0 *.31341      LISTEN      *. *
tcp4      0      0 *.9000       LISTEN      *. *
tcp4      0      0 *.6152       LISTEN      *. *
tcp4      0      0 *.32003      LISTEN      *. *
tcp4      0      0 *.33009      LISTEN      *. *
tcp4      0      0 *.3221       LISTEN      *. *
tcp4      0      0 *.23         LISTEN      *. *
tcp4      0      0 *.22         LISTEN      *. *
tcp4      0      0 *.514        LISTEN      *. *
tcp4      0      0 *.513        LISTEN      *. *
tcp4      0      0 *.21         LISTEN      *. *
tcp4      0      0 *.79         LISTEN      *. *
tcp4      0      0 *.514        LISTEN      *. *
tcp4      0      0 *.513        LISTEN      *. *
tcp4      0      0 *.6234       LISTEN      *. *
udp4      0      0 127.0.0.1.123 LISTEN      *. *
udp4      0      0 10.255.178.11.123 LISTEN      *. *
udp4      0      0 *.123        LISTEN      *. *
udp46     0      0 *.514        LISTEN      *. *
udp4      0      0 *.514        LISTEN      *. *
udp46     0      0 *.62027      LISTEN      *. *
udp4      0      0 *.59363      LISTEN      *. *
udp4      0      0 *.31342      LISTEN      *. *
udp46     0      0 *.161        LISTEN      *. *
udp4      0      0 *.161        LISTEN      *. *
udp4      0      0 *.31340      LISTEN      *. *
udp4      0      0 *.31340      LISTEN      *. *
udp46     0      0 *.49152      LISTEN      *. *
udp46     0      0 *.4784       LISTEN      *. *
udp46     0      0 *.3784       LISTEN      *. *
udp4      0      0 *.49152      LISTEN      *. *
udp4      0      0 *.4784       LISTEN      *. *
udp4      0      0 *.3784       LISTEN      *. *
udp4      0      0 *.6333       LISTEN      *. *
ip4       0      0 *. *         LISTEN      *. *
ip4       0      0 *. *         LISTEN      *. *

```

```
lcc0-re0:
```

```
-----
Active Internet connections (including servers)
```

```

Proto Recv-Q Send-Q Local Address                               Foreign Address
              (state)
tcp4      0      0 192.168.178.3.23
```

```

172.24.26.227.50399
tcp4      0      0 *.6234          ESTABLISHED      *.*
          0      0 *.7000          LISTEN            *.*
          0      0 *.9000          LISTEN            *.*
          0      0 *.33009         LISTEN            *.*
          0      0 *.3221          LISTEN            *.*
          0      0 *.23            LISTEN            *.*
          0      0 *.22            LISTEN            *.*
          0      0 *.514           LISTEN            *.*
          0      0 *.513           LISTEN            *.*
          0      0 *.21            LISTEN            *.*
          0      0 *.79            LISTEN            *.*
          0      0 *.514           LISTEN            *.*
          0      0 *.513           LISTEN            *.*
          0      0 *.514           LISTEN            *.*
          0      0 *.514           LISTEN            *.*
          0      0 *.59924         LISTEN            *.*
          0      0 *.59412         LISTEN            *.*
          0      0 *.161           LISTEN            *.*
          0      0 *.161           LISTEN            *.*
          0      0 *.31342         LISTEN            *.*
          0      0 *.6333          LISTEN            *.*

```

lcc1-re0:

-----  
Active Internet connections (including servers)

Proto	Recv-Q	Send-Q	Local Address (state)	Foreign Address
tcp4	0	0	*.6234 LISTEN	*.*
tcp4	0	0	*.7000 LISTEN	*.*
tcp4	0	0	*.9000 LISTEN	*.*
tcp4	0	0	*.3221 LISTEN	*.*
tcp4	0	0	*.23 LISTEN	*.*
tcp4	0	0	*.22 LISTEN	*.*
tcp4	0	0	*.514 LISTEN	*.*
tcp4	0	0	*.513 LISTEN	*.*
tcp4	0	0	*.21 LISTEN	*.*
tcp4	0	0	*.79 LISTEN	*.*

```

tcp4      0      0 *.514          *.*
          LISTEN
tcp4      0      0 *.513          *.*
          LISTEN
tcp4      0      0 *.33009        *.*
          LISTEN
udp46     0      0 *.514          *.*
udp4      0      0 *.514          *.*
udp46     0      0 *.59924        *.*
udp4      0      0 *.59412        *.*
udp4      0      0 *.31342        *.*
udp46     0      0 *.161          *.*
udp4      0      0 *.161          *.*
udp4      0      0 *.6333         *.*

```

lcc2-re0:

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      0 *.6234          *.*
          LISTEN
tcp4      0      0 *.7000          *.*
          LISTEN
tcp4      0      0 *.9000          *.*
          LISTEN
tcp4      0      0 *.33009        *.*
          LISTEN
tcp4      0      0 *.3221          *.*
          LISTEN
tcp4      0      0 *.23           *.*
          LISTEN
tcp4      0      0 *.22           *.*
          LISTEN
tcp4      0      0 *.514          *.*
...

```

#### show system connections sfc (TX Matrix Plus Router)

```

user@host> show system connections sfc 0
sfc0-re0:

```

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      0 162.0.0.4.514        132.0.0.4.952
          TIME_WAIT
tcp4      0      0 162.0.0.4.514        131.0.0.4.694
          TIME_WAIT
tcp4      0      0 162.0.0.4.514        130.0.0.4.860
          TIME_WAIT
tcp4      0      0 162.0.0.4.514        129.0.0.4.716
          TIME_WAIT
tcp4      0      0 162.0.0.4.996        132.0.0.4.514
          TIME_WAIT
tcp4      0      0 162.0.0.4.798        131.0.0.4.514
          TIME_WAIT
tcp4      0      0 162.0.0.4.995        130.0.0.4.514
          TIME_WAIT
tcp4      0      0 162.0.0.4.895        129.0.0.4.514
          TIME_WAIT

```

tcp4	0	0	192.168.178.11.21		
172.17.28.204.64662				TIME_WAIT	
tcp4	0	0	192.168.178.11.21		
172.17.28.204.51612				TIME_WAIT	
tcp4	0	0	*.6156		*.*
			LISTEN		
tcp4	0	0	*.9000		*.*
			LISTEN		
tcp4	0	0	*.666		*.*
			LISTEN		
tcp4	0	2	192.168.178.11.23		
172.17.28.19.3565				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
172.17.28.204.62719				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
192.168.69.199.51255				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
172.24.26.227.42860				ESTABLISHED	
tcp4	0	0	162.0.0.4.32012		162.0.0.5.58935
			ESTABLISHED		
tcp4	0	0	*.32012		*.*
			LISTEN		
tcp4	0	0	*.33007		*.*
			LISTEN		
tcp4	0	1432	162.0.0.4.6161		162.0.0.5.62026
			ESTABLISHED		
tcp4	0	0	*.33005		*.*
			LISTEN		
tcp4	0	0	162.0.0.4.9000		162.0.0.4.51611
			FIN_WAIT_2		
tcp4	0	0	162.0.0.4.51611		162.0.0.4.9000
			CLOSE_WAIT		
tcp4	0	0	*.6151		*.*
			LISTEN		
tcp4	0	0	*.6154		*.*
			LISTEN		
tcp4	0	0	*.6153		*.*
			LISTEN		
tcp4	0	0	*.31343		*.*
			LISTEN		
tcp4	0	0	*.31341		*.*
			LISTEN		
tcp4	0	0	*.6152		*.*
			LISTEN		
tcp4	0	0	*.32003		*.*
			LISTEN		
tcp4	0	0	*.33009		*.*
			LISTEN		
tcp4	0	0	*.3221		*.*
			LISTEN		
tcp4	0	0	*.23		*.*
			LISTEN		
tcp4	0	0	*.22		*.*
			LISTEN		
tcp4	0	0	*.514		*.*
			LISTEN		
tcp4	0	0	*.513		*.*
			LISTEN		
tcp4	0	0	*.21		*.*
			LISTEN		
tcp4	0	0	*.79		*.*



```

                                LISTEN
tcp4      0      0 *.514                                *.*
                                LISTEN
tcp4      0      0 *.513                                *.*
                                LISTEN
tcp4      0      0 *.6234                               *.*
                                LISTEN
udp4      0      0 127.0.0.1.123                       *.*
udp4      0      0 10.255.178.11.123                   *.*
udp4      0      0 *.123                                *.*
udp46     0      0 *.514                                *.*
udp4      0      0 *.514                                *.*
udp46     0      0 *.50895                              *.*
udp4      0      0 *.50794                              *.*
udp4      0      0 *.31342                              *.*
udp46     0      0 *.161                                *.*
udp4      0      0 *.161                                *.*
udp4      0      0 *.31340                              *.*
udp4      0      0 *.31340                              *.*
udp46     0      0 *.49152                              *.*
udp46     0      0 *.4784                               *.*
udp46     0      0 *.3784                               *.*
udp4      0      0 *.49152                              *.*
udp4      0      0 *.4784                               *.*
udp4      0      0 *.3784                               *.*
udp4      0      0 *.6333                               *.*
ip4       104    0 *.*                                  *.*
ip4       0      0 *.*                                  *.*
ip4       0      0 *.*                                  *.*

```

#### show system connections show-routing-instances (TX Matrix Plus Router)

```

user@host> show system connections show-routing-instances
sfc0-re0:
-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Foreign Address
-----
                                Routing Instance      (state)
tcp4      0      0 *.6156                   __juniper_private1__    LISTEN      *.*
tcp4      0      0 *.9000                   __juniper_private1__    LISTEN      *.*
tcp4      0      0 *.666                    __juniper_private1__    LISTEN      *.*
tcp4      0      2 192.168.178.11.23        default                  ESTABLISHED 172.17.28.19.3565
tcp4      0      0 192.168.178.11.23        default                  ESTABLISHED 172.17.28.204.62719
tcp4      0      0 192.168.178.11.23        default                  ESTABLISHED 192.168.69.199.51255
tcp4      0      0 192.168.178.11.23        default                  ESTABLISHED 172.24.26.227.42860
tcp4      0      0 162.0.0.4.32012          __juniper_private1__    ESTABLISHED 162.0.0.5.58935
tcp4      0      0 *.32012                  __juniper_private1__    LISTEN      *.*
tcp4      0      0 *.33007                  __juniper_private2__    LISTEN      *.*
tcp4      0      0 162.0.0.4.6161          __juniper_private1__    ESTABLISHED 162.0.0.5.62026
tcp4      0      0 *.33005

```

tcp4	0	0	162.0.0.4.9000	__juniper_private2__	LISTEN	162.0.0.4.51611
tcp4	0	0	162.0.0.4.51611	__juniper_private1__	FIN_WAIT_2	162.0.0.4.9000
tcp4	0	0	*.6151	__juniper_private1__	CLOSE_WAIT	*.*
tcp4	0	0	*.6154	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6153	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.31343	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.31341	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6152	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.32003	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.33009	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.3221	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.23	default	LISTEN	*.*
tcp4	0	0	*.22	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	default	LISTEN	*.*
tcp4	0	0	*.21	default	LISTEN	*.*
tcp4	0	0	*.79	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6234	__juniper_private1__	LISTEN	*.*
udp4	0	0	127.0.0.1.123	__juniper_private1__	LISTEN	*.*
udp4	0	0	10.255.178.11.123	default		*.*
udp4	0	0	*.123	default		*.*
udp46	0	0	*.514	default		*.*
udp4	0	0	*.514	default		*.*
udp46	0	0	*.50895	default		*.*
udp4	0	0	*.50794	default		*.*
udp4	0	0	*.31342	default		*.*
udp46	0	0	*.161	__juniper_private1__		*.*
udp4	0	0	*.161	default		*.*
				default		

udp4	0	0	*.31340	__juniper_private2__	*.*
udp4	0	0	*.31340	__juniper_private1__	*.*
udp46	0	0	*.49152	default	*.*
udp46	0	0	*.4784	default	*.*
udp46	0	0	*.3784	default	*.*
udp4	0	0	*.49152	default	*.*
udp4	0	0	*.4784	default	*.*
udp4	0	0	*.3784	default	*.*
udp4	0	0	*.6333	__juniper_private1__	*.*
ip4	0	0	*.*	default	*.*
ip4	0	0	*.*	default	*.*
ip4	0	0	*.*	default	*.*

lcc0-re0:

-----

Active Internet connections (including servers) (including routing-instances)

Proto	Recv-Q	Send-Q	Local Address	Routing Instance	(state)	Foreign Address
tcp4	0	0	*.7000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	192.168.178.3.23	default	ESTABLISHED	
tcp4	0	0	*.6234	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.9000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.33009	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.3221	default	LISTEN	*.*
tcp4	0	0	*.23	default	LISTEN	*.*
tcp4	0	0	*.22	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	default	LISTEN	*.*
tcp4	0	0	*.21	default	LISTEN	*.*
tcp4	0	0	*.79	default	LISTEN	*.*
tcp4	0	0	*.514	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.513	__juniper_private1__	LISTEN	*.*
udp46	0	0	*.514	default		*.*
udp4	0	0	*.514			*.*

```

udp46      0      0 *.59924    default          *.
udp4       0      0 *.59412    default          *.
udp46      0      0 *.161      default          *.
udp4       0      0 *.161      default          *.
udp4       0      0 *.31342    default          *.
udp4       0      0 *.6333     __juniper_private1__
                                         *.
                                         __juniper_private1__

```

```
lcc1-re0:
```

```

-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Routing Instance      (state)      Foreign Address
tcp4      0      0 *.7000          __juniper_private1__ LISTEN         *.
tcp4      0      0 *.6234          __juniper_private1__ LISTEN         *.
tcp4      0      0 *.9000          __juniper_private1__ LISTEN         *.
tcp4      0      0 *.3221          __juniper_private1__ LISTEN         *.
tcp4      0      0 *.23            default            LISTEN         *.
tcp4      0      0 *.22            default            LISTEN         *.
tcp4      0      0 *.514           default            LISTEN         *.
tcp4      0      0 *.513           default            LISTEN         *.
tcp4      0      0 *.21            default            LISTEN         *.
tcp4      0      0 *.79            default            LISTEN         *.
tcp4      0      0 *.514           __juniper_private1__ LISTEN         *.
tcp4      0      0 *.513           __juniper_private1__ LISTEN         *.
tcp4      0      0 *.33009         __juniper_private2__ LISTEN         *.
udp46     0      0 *.514           default            *.
udp4      0      0 *.514           default            *.
udp46     0      0 *.59924         default            *.
udp4      0      0 *.59412         default            *.
udp4      0      0 *.31342         default            *.
udp46     0      0 *.161           __juniper_private1__ *.
udp4      0      0 *.161           default            *.
udp4      0      0 *.6333          default            *.
                                         __juniper_private1__

```

lcc2-re0:

Active Internet connections (including servers) (including routing-instances)						
Proto	Recv-Q	Send-Q	Local Address	Routing Instance	(state)	Foreign Address
tcp4	0	0	*.7000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6234	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.9000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.33009	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.3221	default	LISTEN	*.*
tcp4	0	0	*.23	default	LISTEN	*.*
tcp4	0	0	*.22	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	default	LISTEN	*.*
tcp4	0	0	*.21	default	LISTEN	*.*
tcp4	0	0	*.79	default	LISTEN	*.*
tcp4	0	0	*.514	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.513	__juniper_private1__	LISTEN	*.*
udp46	0	0	*.514	default		*.*
udp4	0	0	*.514	default		*.*
udp4	0	0	*.31342	__juniper_private1__		*.*
udp46	0	0	*.62103	default		*.*
udp4	0	0	*.59924	default		*.*
udp46	0	0	*.161	default		*.*
udp4	0	0	*.161	default		*.*
udp4	0	0	*.6333	__juniper_private1__		*.*

lcc3-re0:

Active Internet connections (including servers) (including routing-instances)						
Proto	Recv-Q	Send-Q	Local Address	Routing Instance	(state)	Foreign Address
tcp4	0	0	*.7000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6234	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.9000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.33009	__juniper_private1__		*.*

tcp4	0	0	*.3221	__juniper_private2__	LISTEN	*,*
				default	LISTEN	*,*
tcp4	0	0	*.23	default	LISTEN	*,*
tcp4	0	0	*.22	default	LISTEN	*,*
tcp4	0	0	*.514	default	LISTEN	*,*
tcp4	0	0	*.513	default	LISTEN	*,*
tcp4	0	0	*.21	default	LISTEN	*,*
tcp4	0	0	*.79	default	LISTEN	*,*
tcp4	0	0	*.514	__juniper_private1__	LISTEN	*,*
tcp4	0	0	*.513	__juniper_private1__	LISTEN	*,*
udp46	0	0	*.514	default		*,*
udp4	0	0	*.514	default		*,*
udp46	0	0	*.62103	default		*,*
udp4	0	0	*.59924	default		*,*
udp4	0	0	*.31342	__juniper_private1__		*,*
udp46	0	0	*.161	default		*,*
udp4	0	0	*.161	default		*,*
udp4	0	0	*.6333	__juniper_private1__		*,*

#### show system connections (QFX3500 Switch)

```

user@switch> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
              (state)
tcp4          0      0 10.94.204.110.23        172.17.28.19.1308    ESTABLISHED
tcp4          0      0 128.0.0.1.6234          128.0.0.1.65142     ESTABLISHED
tcp4          0      0 128.0.0.1.65142         128.0.0.1.6234     ESTABLISHED
tcp4          0      0 128.0.0.1.33003         128.0.0.1.61441     ESTABLISHED
tcp4          0      0 128.0.0.1.61441         128.0.0.1.33003     ESTABLISHED
tcp46         0      0 *.179                   *.*                  LISTEN
tcp4          0      0 *.179                   *.*                  LISTEN
tcp4          0      0 128.0.0.16.9000         128.0.0.16.50970    ESTABLISHED
tcp4          0      0 128.0.0.16.50970        128.0.0.16.9000     ESTABLISHED
tcp4          0      0 *.38                     *.*                  LISTEN

```

			LISTEN	
tcp4	0	0 *.3491		*.*
			LISTEN	
tcp4	0	0 *.6156		*.*
			LISTEN	
tcp4	0	0 128.0.0.1.33001		128.0.0.1.59437
			ESTABLISHED	
tcp4	0	0 128.0.0.1.59437		128.0.0.1.33001
			ESTABLISHED	
tcp4	0	0 128.0.0.1.33023		128.0.0.1.63605
			ESTABLISHED	
tcp4	0	0 128.0.0.1.63605		128.0.0.1.33023
			ESTABLISHED	
tcp4	0	0 128.0.0.1.33001		128.0.0.1.63830
			ESTABLISHED	
tcp4	0	0 128.0.0.1.63830		128.0.0.1.33001
			ESTABLISHED	
tcp4	0	0 *.667		*.*
			LISTEN	
tcp4	0	0 *.6156		*.*
			LISTEN	
tcp4	0	0 128.0.0.1.7000		128.0.0.1.51580
			ESTABLISHED	
tcp4	0	0 128.0.0.1.51580		128.0.0.1.7000
			ESTABLISHED	
tcp4	0	0 128.0.0.1.6234		128.0.0.1.53646
			ESTABLISHED	
tcp4	0	0 *.33001		*.*
			LISTEN	
tcp4	0	0 *.33003		*.*
			LISTEN	
tcp4	0	0 128.0.0.1.53646		128.0.0.1.6234
			ESTABLISHED	
tcp4	0	0 128.0.0.16.9000		128.0.0.16.63454
			ESTABLISHED	
tcp4	0	0 128.0.0.16.63454		128.0.0.16.9000
			ESTABLISHED	
tcp4	0	0 *.666		*.*
			LISTEN	
tcp4	0	0 *.7000		*.*
			LISTEN	
tcp4	0	0 *.51627		*.*
			LISTEN	
tcp4	0	0 *.3492		*.*
			LISTEN	
tcp4	0	0 *.33023		*.*
			LISTEN	
tcp4	0	0 *.33013		*.*
			LISTEN	
tcp4	0	0 *.7202		*.*
			LISTEN	
tcp4	0	0 *.6151		*.*
			LISTEN	
tcp4	0	0 *.9000		*.*
			LISTEN	
tcp4	0	0 *.6161		*.*
			LISTEN	
tcp4	0	0 *.6011		*.*
			LISTEN	
tcp4	0	0 *.3221		*.*
			LISTEN	

tcp4	0	0 *.23		*. *
			LISTEN	
tcp4	0	0 *.22		*. *
			LISTEN	
tcp4	0	0 *.514		*. *
			LISTEN	
tcp4	0	0 *.513		*. *
			LISTEN	
tcp4	0	0 *.21		*. *
			LISTEN	
tcp4	0	0 *.79		*. *
			LISTEN	
tcp4	0	0 *.514		*. *
			LISTEN	
tcp4	0	0 *.513		*. *
			LISTEN	
tcp4	0	0 *.1127		*. *
			LISTEN	
tcp4	0	0 *.1129		*. *
			LISTEN	
tcp4	0	0 *.1128		*. *
			LISTEN	
tcp4	0	0 *.6234		*. *
			LISTEN	
udp46	0	0 *.514		*. *
udp4	0	0 *.514		*. *
udp4	0	0 128.0.0.1.123		*. *
udp46	0	0 *.53344		*. *
udp4	0	0 *.54261		*. *
udp46	0	0 *.161		*. *
udp4	0	0 *.161		*. *
udp4	0	0 *.31342		*. *
udp4	0	0 *.59137		*. *
udp4	0	0 *. *		*. *
udp46	0	0 *.49152		*. *
udp46	0	0 *.4784		*. *
udp46	0	0 *.3784		*. *
udp4	0	0 *.49152		*. *
udp4	0	0 *.4784		*. *
udp4	0	0 *.3784		*. *
udp4	0	0 10.255.204.110.123		*. *
udp4	0	0 *.123		*. *
udp4	0	0 *.67		*. *
udp4	0	0 *.6333		*. *
udp4	0	0 *.2293		*. *
ip4	0	0 *. *		*. *
ip4	0	0 *. *		*. *
ip4	0	0 *. *		*. *



## show system core-dumps

<b>List of Syntax</b>	<a href="#">Syntax on page 1305</a> <a href="#">Syntax (EX Series Switches) on page 1305</a> <a href="#">Syntax (TX Matrix Router) on page 1305</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1305</a> <a href="#">Syntax (QFX Series) on page 1305</a>
<b>Syntax</b>	<pre>show system core-dumps &lt;brief   detail&gt; &lt;core-filename&gt; &lt;core-file-info&gt; &lt;re0&gt; &lt;re1&gt; &lt;routing-engine&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show system core-dumps &lt;all-members&gt; &lt;brief   detail&gt; &lt;core-filename&gt; &lt;core-file-info&gt; &lt;local&gt; &lt;member member-id&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>show system core-dumps &lt;all-chassis   all-lcc   lcc number   scc&gt; &lt;brief   detail&gt; &lt;core-filename&gt; &lt;core-file-info&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>show system core-dumps &lt;all-chassis   all-lcc   lcc number   sfc number&gt; &lt;brief   detail&gt; &lt;core-filename&gt; &lt;core-file-info&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show system core-dumps &lt;brief   detail&gt; &lt;component (UUID   serial number   all)&gt; &lt;core-file-info component (UUID   serial number) core-file-name&gt; &lt;display-period (hours   minutes   seconds)&gt; &lt;display-order&gt; &lt;kernel-crashinfo component (UUID   serial number)&gt; &lt;repository (core   log)&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p><b>re0</b>, <b>re1</b>, and <b>routing-engine</b> options introduced for dual routing engines in Junos OS Release 13.1.</p>

**Description** Show core files on all routers or switches running Junos OS. You can use the **show system core-dumps** command to show a list of system core files created when the router or switch has failed. This command can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification date, and path and filename. If dual routing engines are present, you can view core-dump files for either routing engine or both routing engines together. On a QFabric system, you can view core-dump files on individual QFabric system devices as well as on the entire QFabric system.

You can use the option **core-filename** and its options **core-file-info**, **brief**, and **detail** to display more information about the specified core-dump files.

**Options** **none**—Display a list of all existing core-dump files.



**NOTE:** If dual routing engines are present, lists only the core-dump files for the active routing engine.

**all-chassis**—(TX Matrix and TX Matrix Plus routers only) (Optional) On a routing matrix based on a TX Matrix router, display system core files for the TX Matrix router switch-card chassis [SCC] and all the T640 routers [LCCs] connected to the TX Matrix router.

On a routing matrix based on a TX Matrix Plus router, display system core files for the TX Matrix Plus router (switch-fabric chassis [SFC]) and all the T1600 routers [LCCs] connected to the TX Matrix Plus router.

**<all-lcc | lcc number>**—(TX Matrix and TX Matrix Plus routers only) (Optional) On a routing matrix based on the TX Matrix router, display core dump files for all T640 routers (line-card chassis [LCCs]) or a specific T640 router [LCC] connected to the TX Matrix router.

On a routing matrix based on the TX Matrix Plus router, display logging information for all T1600 routers (line-card chassis [LCCs]) or a specific T1600 router (LCC) connected to the TX Matrix Plus router. When using the **lcc number** option, replace **number** with a value from 0 through 3.



**NOTE:** The **all-chassis** option displays system core files for the SCC or SFC and the LCCs connected to the SCC or SFC in the routing matrix while the **all-lcc** option only displays system core files for the LCCs in the routing matrix.

**all-members**—(EX4200 switches) (Optional) Display system core files on all members of the Virtual Chassis configuration.

**brief**—(Optional) View details of a binary file.

**component** (*UUID | serial number | all*)—(QFabric systems only) (Optional) Display a list of core-dump files located on individual QFabric system device or on the entire QFabric system.

**core-file-info**—(Optional) Display the stack trace of a core file.

**core-filename**—(Optional) Name of a specific core file to display.

**detail**—(Optional) View stack trace with details of the binary file.

**display-order** (*timestamp-sort | alphanumeric-sort*)—(QFabric systems only) (Optional) Display list of debug artifacts generated within the specified period—for example, within the last hour, within the last 20 minutes, or within the last 32 seconds—or according to their filename.

**display-period** (*hours | minutes | seconds*)—(QFabric systems only) (Optional) Display core-dump files generated within the specified period—for example, within the last hour, within the last 20 minutes, or within the last 32 seconds.

**kernel-crashinfo component** (*UUID | serial number*)—(QFabric systems only) (Optional) Display kernel crash information from the EEPROM on a QFabric system device.

**local**—(EX4200 switches only) (Optional) Display system core files on the local Virtual Chassis member.

**member member-id**—(EX4200 switches only) (Optional) Display system core files on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

**re0**—(Dual routing engines only) Display the core-dump files on re0.

**re1**—(Dual routing engines only) Display the core-dump files on re1.

**repository** (*core | log*)—(QFabric systems only) (Optional) Specify either the core or log repository in which to view core-dump files.

**routing-engine** (*backup | both | local | master | other*)—(Dual routing engines only) Display a list of core-dump files for either the backup, local, master, or other routing engine or both routing engines.

**scc**—(TX Matrix routers only) (Optional) Display system core files on the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display system core files on the TX Matrix Plus router (or switch-fabric chassis).

**Required Privilege Level**

view

**List of Sample Output**

[show system core-dumps on page 1309](#)

[show system core-dumps on page 1310](#)

[show system core-dumps routing-engine both on page 1310](#)

[show system core-dumps \(TX Matrix Plus Router\) on page 1310](#)

[show system core-dumps \(QFX3500 Switch\) on page 1312](#)  
[show system core-dumps \(QFabric Systems\) on page 1312](#)  
[show system core-dumps core-file-info component serial number core-file-name \(QFabric Systems\) on page 1313](#)  
[show system core-dumps component serial number display-order alphanumeric-sort repository core \(QFabric Systems\) on page 1313](#)  
[show system core-dumps display-period \(QFabric Systems\) on page 1313](#)  
[show system core-dumps kernel-crashinfo component serial number \(QFabric Systems\) on page 1315](#)  
[show system core-dumps repository core \(QFabric Systems\) on page 1317](#)  
[show system core-dumps repository log \(QFabric Systems\) on page 1317](#)

**Output Fields** Table 127 on page 1308 describes the output fields for the **show system core-dumps** command. Output fields are listed in the approximate order in which they appear.

**Table 127: show system core-dumps Output Fields**

Field Name	Field Description
<i>Permissions</i>	Read/write permissions for the file named.
<i>Links</i>	Number of links to the file.
<i>Owner</i>	Name of the file owner.
<i>Group</i>	Name of the group with file access.
<i>File size</i>	File size in bytes.
<i>Modified</i>	Last file modification date and time.
<i>Path/filename</i>	File path where the file resides and the filename.
<b>Repository scope:</b>	Repository where core-dump files and log files are stored. The core-dump files are located in the <b>core</b> repository, and the log files are located in the <b>log</b> repository. The default <b>Repository scope</b> is shared since both the <b>core</b> and <b>log</b> repositories are shared by all of the QFabric system devices.
<b>Repository head:</b>	Path to the top-level repository location.
<b>Repository name:</b>	Name of the repository: <b>core</b> or <b>log</b> .
<b>List of nodes for core repository:</b>	List of core-dump files associated with a particular QFabric system device located in the core repository.
<b>Node Group</b>	Name of the QFabric system device.
<b>Node Identifier</b>	UUID or serial number of the QFabric system device.
<b>Num</b>	Number of core-dump and log files.

Table 127: show system core-dumps Output Fields (*continued*)

Field Name	Field Description
<b>Model</b>	Model number of the QFabric system device.
<b>Usage</b>	Usage of the repository in megabytes.
<b>Total usage of core repository:</b>	Total usage of core-dump files associated with a particular QFabric system device located in the core repository. Usage is specified in megabytes and as a percentage.
<b>Total usage of log repository:</b>	Total usage of log files associated with a particular QFabric system device located in the log repository. Usage is specified in megabytes and as a percentage.
<b>List of nodes for core repository:</b>	List of core-dump files associated with a particular QFabric system device located in the core repository.
<b>List of nodes for log repository:</b>	List of log files associated with a particular QFabric system device located in the log repository.
<b>Filename</b>	Name of the core-dump file.
<b>Date</b>	Last core-dump file modification date and time.
<b>Size</b>	Size of the core-dump file.
<b>Core filename</b>	Filename of the core-dump file.
<b>Process name</b>	Name of the process that is generating a core-dump file or log file.
<b>Release</b>	Junos OS release.
<b>Build server</b>	Junos OS build server.
<b>Build date</b>	Junos OS build date.
<b>Stack trace</b>	Stack trace of the core-dump file.

## Sample Output

### show system core-dumps

This example shows the command output if core files exist.

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

### show system core-dumps

This example shows the command output if core files do not exist.

```
user@host> show system core-dumps
/var/crash/*core*: No such file or directory
/var/tmp/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
```

### show system core-dumps routing-engine both

This example shows the command output if dual routing engines are present.

```
user@host> show system core-dumps routing-engine both
re0:
-----
/var/crash/*core*: No such file or directory
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory

/var/tmp/cores:
total blocks: 496776
-rw-rw---- 1 root field 11910589 Nov 8 13:20 chassisd.core.0.201311081320
...

-rw-rw---- 1 root field 11737227 Oct 28 14:21
rpd.core-tarball.4.tgz.201310281421.3458162
total files: 10

re1:
-----
/var/crash/*core*: No such file or directory
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory

/var/tmp/cores:
total blocks: 3178420
-rw-rw---- 1 root field 19039721 Nov 8 14:29
chassisd.core.0.201311081429.3485600.gz
-rw-rw---- 1 root field 19039793 Nov 8 14:37
chassisd.core.1.201311081437.3485599.gz
..

-rw-rw---- 1 root field 11710113 Oct 17 15:26
rpd.core-tarball.1.1.tgz.201310171526.3430028
```

### show system core-dumps (TX Matrix Plus Router)

```
user@host> show system core-dumps
sfc0-re0:
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory

/var/crash/cores:
total 8

/var/tmp/cores:
total 1627592
-rw-r--r-- 1 root field 535346090 May 15 07:36
rpd.core-tarball.0.090515.0736.tgz
```

```
-rw-r--r-- 1 root field 105632057 May 15 07:37
rpd.core-tarball.1.090515.0737.tgz
-rw-r--r-- 1 root field 101981681 May 15 07:38
rpd.core-tarball.2.090515.0738.tgz
-rw-r--r-- 1 root field 85854573 May 15 07:40
rpd.core-tarball.3.090515.0740.tgz
-rw-r--r-- 1 root field 4157845 May 15 08:18
rpd.core-tarball.4.090515.0818.tgz
```

lcc0-re0:

```
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
```

```
/var/crash/cores:
total 8
```

```
/var/tmp/cores:
total 12
```

lcc1-re0:

```
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
```

```
/var/crash/cores:
total 8
```

```
/var/tmp/cores:
total 10024
```

```
-rw-r--r-- 1 root field 1875794 Apr 22 15:47
chassisd.core-tarball.0.090422.1547.tgz
-rw-r--r-- 1 root field 1894183 Apr 22 19:02
chassisd.core-tarball.0.090422.1902.tgz
-rw-r--r-- 1 root field 1290240 Apr 26 16:01 ksyncd_1558.core.0.090426.1601
```

lcc2-re0:

```
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
```

```
/var/crash/cores:
total 21124008
```

```
-rw-r--r-- 1 root wheel 1022376528 May 2 06:43
core-LCC2-EGFPC7.core.0.090502.0643
-rw-r--r-- 1 root wheel 1022376528 May 2 08:13
core-LCC2-EGFPC7.core.0.090502.0813
-rw-r--r-- 1 root wheel 1022376544 May 5 06:15
core-LCC2-EGFPC7.core.0.090505.0615
-rw-r--r-- 1 root wheel 1022376544 May 6 10:59
core-LCC2-EGFPC7.core.0.090506.1059
-rw-r--r-- 1 root wheel 1022376528 May 2 06:58
core-LCC2-EGFPC7.core.1.090502.0658
-rw-r--r-- 1 root wheel 754271232 May 5 06:33
core-LCC2-EGFPC7.core.1.090505.0633
-rw-r--r-- 1 root wheel 264897536 May 6 11:12
core-LCC2-EGFPC7.core.1.090506.1112
-rw-r--r-- 1 root wheel 1022376528 May 2 07:22
core-LCC2-EGFPC7.core.2.090502.0722
-rw-r--r-- 1 root wheel 163633152 May 5 06:52
core-LCC2-EGFPC7.core.2.090505.0652
```

```

-rw-r--r-- 1 root wheel 171312128 May 6 12:13
core-LCC2-EGFPC7.core.2.090506.1213
-rw-r--r-- 1 root wheel 1022376528 May 2 07:39
core-LCC2-EGFPC7.core.3.090502.0739
-rw-r--r-- 1 root wheel 1022376528 May 2 07:55
core-LCC2-EGFPC7.core.4.090502.0755
-rw-r--r-- 1 root wheel 427277312 May 7 04:47
core-LCC2-STFPC4.core.0.090507.0447
-rw-r--r-- 1 root wheel 419609600 May 7 04:47
core-LCC2-STFPC5.core.0.090507.0447
-rw-r--r-- 1 root wheel 432356352 May 7 04:47
core-LCC2-STFPC6.core.0.090507.0447

/var/tmp/cores:
total 2568
-rw-r--r-- 1 root field 1290240 May 14 14:26 ksyncd_1540.core.0.090514.1426
...

```

### show system core-dumps (QFX3500 Switch)

```

user@switch> show system core-dumps
/var/crash/*core*: No such file or directory
-rw-rw---- 1 root field 1545143 Jun 4 2012 /var/tmp/pafxpc.core.0.gz
-rw-rw---- 1 root field 1545146 Jun 4 2012 /var/tmp/pafxpc.core.1.gz
-rw-rw---- 1 root field 1545141 Jun 4 2012 /var/tmp/pafxpc.core.2.gz
-rw-rw---- 1 root field 1545146 Jun 4 2012 /var/tmp/pafxpc.core.3.gz
-rw-rw---- 1 root field 1545142 Jun 5 2012 /var/tmp/pafxpc.core.4.gz
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
total 5

```

### show system core-dumps (QFabric Systems)

```

user@switch> show system core-dumps
Repository scope: shared
Repository head: /pbdata/export
List of nodes for core repository: /pbdata/export/rdumps/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	OM
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	0	fx-jvre	OM
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	0	fx-jvre	OM
NW-NG-0	BBAK0394	0	qfx3500	OM
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	0	fx-jvre	OM
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	0	fx-jvre	OM
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	0	fx-jvre	OM
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	0	fx-jvre	OM
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YW3803	0	qfxc08-3008	OM
IC-WS001	WS001/YN5999	0	qfxc08-3008	OM
node-device1	BBAK0372	0	qfx3500	OM
node-device1	EE3093	0	qfx3500	OM

Total usage of core repository: 0M of 70000M (0.0%)

```

List of nodes for log repository: /pbdata/export/rlogs/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	OM
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	1	fx-jvre	OM
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	1	fx-jvre	OM



NW-NG-0	BBAK0394	1	qfx3500	OM
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	1	fx-jvre	OM
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	3	fx-jvre	OM
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	1	fx-jvre	OM
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	1	fx-jvre	OM
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YN5999	1	qfxc08-3008	OM
IC-WS001	WS001/YW3803	1	qfxc08-3008	OM
node-device1	BBAK0372	1	qfx3500	OM
node-device1	EE3093	1	qfx3500	OM

Total usage of log repository: 0M of 70000M (0.0%)

#### show system core-dumps core-file-info component serial number core-file-name (QFabric Systems)

```

user@switch> show system core-dumps core-file-info component
e8ff4b3e-7d92-11e0-be5d-00e081c1fe0e cosd.core.0.1519.05162011131846.gz
Repository scope: shared
Repository head: /pbstorage
Repository name: core
Core filename: /pbstorage/rdumps/e8ff4b3e-7d92-11e0-be5d-
00e081c1fe0e/5658.cosd.core.0.1519.05162011131846
Process name: cosd
Release: 11.3I0
Build server: /c/ssengupta/dfx_ha_v1/obj-i386-dcp/dcp/usr.sbin/cosd
Build date: 2011-05-14 01:11:44 UTC
Stack trace:
#0 0x8885d183 in select () from /usr/lib/libc.so.6
#0 0x8885d183 in select () from /usr/lib/libc.so.6
#1 0x887d4a45 in pselect () from /usr/lib/libc.so.6
#2 0x88774719 in pselect () from /usr/lib/libthr.so.2
#3 0x885de5db in __evGetNext () from /usr/lib/libisc.so.2
#4 0x885debf0 in __evMainLoop () from /usr/lib/libisc.so.2
#5 0x081125b2 in cosd_loop ()
#6 0x0812e19a in main ()

```

#### show system core-dumps component serial number display-order alphanumeric-sort repository core (QFabric Systems)

```

user@switch> show system core-dumps component BBAK8891 display-order alphanumeric-sort
repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
List of core dumps for component BBAK8891
Repository location: /pbdata/export/rdumps/BBAK8891

```

Filename	Date	Size
eswd.core.0.1361.11172011214257.gz	Nov 17 21:43:10 2011	4779553
eswd.core.1.80267.11172011214514.gz	Nov 17 21:45:19 2011	3541648
eswd.core.2.80682.11172011214535.gz	Nov 17 21:45:43 2011	2156683
vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:35 2011	375617

Number of core dumps in repository: 4

#### show system core-dumps display-period (QFabric Systems)

```

user@switch> show system core-dumps display-period 24h
show system core-dumps display-period 24h
Repository scope: shared
Repository head: /pbdata/export
List of core dumps at repository: /pbdata/export/rdumps
Delta timespec: Last 24h

```

Component: BBAK8273		
Filename	Size	Date
vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:35 2011	375794
Component: cedb7b0e-0025-11e1-9a5f-00e081c52990		
Filename	Size	Date
vccpd.core.0.1461.11182011151131.gz	Nov 18 15:11:31 2011	120951
Component: ee19c4f8-0025-11e1-ae6f-00e081c52990		
Filename	Size	Date
vccpd.core.0.1462.11182011151131.gz	Nov 18 15:11:31 2011	109420
Component: BBAK8281		
Filename	Size	Date
vccpd.core.0.1196.11182011151131.gz	Nov 18 15:11:36 2011	375373
Component: BBAK8891		
Filename	Size	Date
vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:35 2011	375617
Component: BBAK8276		
Filename	Size	Date
vccpd.core.0.1196.11182011151131.gz	Nov 18 15:11:35 2011	375350
Component: BBAK8868		
Filename	Size	Date
vccpd.core.0.1196.11182011151130.gz	Nov 18 15:11:34 2011	376211
Component: BBAK8835		
Filename	Size	Date
vccpd.core.0.1195.11182011151130.gz	Nov 18 15:11:35 2011	375700
Component: BBAK8283		
Filename	Size	Date
vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:36 2011	368298
Component: YW3781/YW3781		
Filename	Size	Date
vccpd.core.0.1220.11182011151131.gz	Nov 18 15:11:38 2011	380002
Component: 09726be2-0026-11e1-82d9-00e081c52990		
Filename	Size	Date
vccpd.core.0.1461.11182011151130.gz	Nov 18 15:11:31 2011	119965
Component: BBAK8309		
Filename	Size	Date
vccpd.core.0.1196.11182011151131.gz	Nov 18 15:11:36 2011	378930
Component: 303d476a-0026-11e1-abf4-00e081c52990		
Filename	Size	Date
vccpd.core.0.1460.11182011151131.gz	Nov 18 15:11:31 2011	118385
Component: YW3798/YW3798		
Filename	Size	Date
vccpd.core.0.1219.11182011151131.gz	Nov 18 15:11:36 2011	380455
List of log dumps at repository: /pbdata/export/rlogs		
Delta timespec: Last 24h		
Component: BBAK8273		
Filename	Size	Date

vccpd.tarball.0.1195.11182011151138.tgz	Nov 18 15:11:39 2011	20415
Component: cedb7b0e-0025-11e1-9a5f-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1461.11182011151131.tgz	Nov 18 15:11:33 2011	19651
Component: ee19c4f8-0025-11e1-aef6-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1462.11182011151133.tgz	Nov 18 15:11:36 2011	24650
Component: BBAK8281		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz	Nov 18 15:11:41 2011	19445
Component: BBAK8891		
Filename	Size	Date
vccpd.tarball.0.1195.11182011151138.tgz	Nov 18 15:11:41 2011	21916
Component: BBAK8276		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz	Nov 18 15:11:39 2011	20461
Component: BBAK8868		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz	Nov 18 15:11:41 2011	21924
Component: BBAK8835		
Filename	Size	Date
vccpd.tarball.0.1195.11182011151137.tgz	Nov 18 15:11:39 2011	19424
Component: BBAK8283		
Filename	Size	Date
vccpd.tarball.0.1195.11182011151138.tgz	Nov 18 15:11:42 2011	31186
Component: YW3781/YW3781		
Filename	Size	Date
vccpd.tarball.0.1220.11182011151141.tgz	Nov 18 15:11:45 2011	27565
Component: 09726be2-0026-11e1-82d9-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1461.11182011151130.tgz	Nov 18 15:11:34 2011	19613
Component: BBAK8309		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151138.tgz	Nov 18 15:11:46 2011	50362
Component: 303d476a-0026-11e1-abf4-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1460.11182011151133.tgz	Nov 18 15:11:33 2011	19360
Component: YW3798/YW3798		
Filename	Size	Date
vccpd.tarball.0.1219.11182011151140.tgz	Nov 18 15:11:49 2011	24473

#### show system core-dumps kernel-crashinfo component serial number (QFabric Systems)

```
user@switch> show system core-dumps kernel-crashinfo component A0001/YA0197
Node: A0001/YA0197
```

Information about previous kernel crash:

-- Kernel panic data --

Panic string: kdb\_sysctl\_panic

System uptime: 3 day 20 hr 59 min 40 sec Kernel crash time: 2011-11-15 Wed 15:25:17

Kernel build linkstamp: JUNOS 11.3I #0: 2011-11-10 20:42:27 UTC

-- Stacktrace of panicing context --

Processor 1 (crash monarch):

savectx+0x0 (c9552800,80214efc,802a7fbc,c88ad05c) ra 801b93a8 sz 0

kdm\_kcore\_save\_crashinfo+0x254 (c9552800,0,802a7fbc,c88ad05c) ra 801b9f44 sz 784

kdm\_kcore\_kern\_panic\_event\_handler+0x4b0 (c9552800,0,802a7fbc,c88ad05c) ra 8022a9b8 sz 88

panic+0x1d0 (c9552800,0,4,77fed534) ra 802540c0 sz 56

kdb\_sysctl\_panic+0x70 (c9552800,0,4,77fed534) ra 80237e58 sz 40 sysctl\_root+0x12c (c9552800,0,4,e8bc5cf8) ra 80238e50 sz 48

userland\_sysctl+0x164 (c9552800,0,4,e8bc5cf8) ra 8023956c sz 104

\_\_sysctl+0xe4 (c9552800,0,4,e8bc5cf8) ra 806d62e8 sz 160

trap+0xe1c (c9552800,0,4,e8bc5cf8) ra 80896e68 sz 128

MipsUserGenException+0x1a4 (c9552800,0,4,405cd12c) ra 0 sz 0

pid 82340, process: sysctl

Processor 0:

restoreintr+0x14 (1,81bca820,3,0) ra 806cdc3c sz 0

spinlock\_exit+0x30 (1,81bca820,3,0) ra 8025d354 sz 24

sleepq\_release+0x64 (1,81bca820,3,0) ra 8025e670 sz 24

sleepq\_timeout+0x224 (1,81bca820,3,0) ra 80240294 sz 48

softclock+0x434 (1,81bca820,3,0) ra 802067f8 sz 80

ithread\_loop+0x244 (1,81bca820,3,0) ra 80200e28 sz 64 fork\_exit+0xc0

(1,81bca820,3,0) ra 80897c28 sz 48

MipsNMIException+0x34 (1,81bca820,3,0) ra 0 sz 0

pid 82340, process: sysctl

Processor 2:

cpu\_idle+0x20 (80960000,51bbc,2031df,81bca1b8) ra 80204948 sz 24 idle\_proc+0x130

(80960000,51bbc,2031df,81bca1b8) ra 80200e28 sz 56 fork\_exit+0xc0

(80960000,51bbc,2031df,81bca1b8) ra 80897c28 sz 48

MipsNMIException+0x34 (80960000,51bbc,2031df,81bca1b8) ra 0 sz 0

pid 82340, process: sysctl

Processor 3:

cpu\_idle+0x20 (80960000,51bbc,2038df,81bca300) ra 80204948 sz 24 idle\_proc+0x130

(80960000,51bbc,2038df,81bca300) ra 80200e28 sz 56 fork\_exit+0xc0

(80960000,51bbc,2038df,81bca300) ra 80897c28 sz 48

MipsNMIException+0x34 (80960000,51bbc,2038df,81bca300) ra 0 sz 0

pid 82340, process: sysctl

Processor 4:

cpu\_idle+0x20 (80960000,51bbc,2037df,81bca448) ra 80204948 sz 24 idle\_proc+0x130

(80960000,51bbc,2037df,81bca448) ra 80200e28 sz 56 fork\_exit+0xc0

(80960000,51bbc,2037df,81bca448) ra 80897c28 sz 48

MipsNMIException+0x34 (80960000,51bbc,2037df,81bca448) ra 0 sz 0

pid 82340, process: sysctl

Processor 5:

restoreintr+0x14 (1,51bbc,203edf,81bca590) ra 806cdc3c sz 0

spinlock\_exit+0x30 (1,51bbc,203edf,81bca590) ra 80204a34 sz 24 idle\_proc+0x21c

(1,51bbc,203edf,81bca590) ra 80200e28 sz 56 fork\_exit+0xc0

(1,51bbc,203edf,81bca590) ra 80897c28 sz 48

MipsNMIException+0x34 (1,51bbc,203edf,81bca590) ra 0 sz 0

pid 82340, process: sysctl

```

Processor 6:
cpu_idle+0x20 (80960000,51bbc,205cdf,81bca6d8) ra 80204948 sz 24 idle_proc+0x130
(80960000,51bbc,205cdf,81bca6d8) ra 80200e28 sz 56 fork_exit+0xc0
(80960000,51bbc,205cdf,81bca6d8) ra 80897c28 sz 48
MipsNMIEException+0x34 (80960000,51bbc,205cdf,81bca6d8) ra 0 sz 0
pid 82340, process: sysctl

Processor 7:
lockmgr+0x5ac (c97e8484,c8dd9800,0,c8dd9800) ra 8c11c81c sz 48
sal_sem_take+0x134 (c97e8484,c8dd9800,0,c8dd9800) ra 8c351108 sz 56
_bcm_esw_linkscan_thread+0x45c (c97e8484,c8dd9800,0,c8dd9800) ra 8c11cdb4 sz 104
sal_thread_start_wrap+0x74 (c97e8484,c8dd9800,0,c8dd9800) ra 80200e28 sz 32
fork_exit+0xc0 (c97e8484,c8dd9800,0,c8dd9800) ra 80897c28 sz 48
MipsNMIEException+0x34 (c97e8484,c8dd9800,0,c8dd9800) ra 0 sz 0
pid 82340, process: sysctl
-- End of stacktrace --

```

### show system core-dumps repository core (QFabric Systems)

```

user@switch> show system core-dumps repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
List of nodes for core repository: /pbdata/export/rdumps/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	0M
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	0	fx-jvre	0M
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	0	fx-jvre	0M
NW-NG-0	BBAK0394	0	qfx3500	0M
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	0	fx-jvre	0M
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	0	fx-jvre	0M
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	0	fx-jvre	0M
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	0	fx-jvre	0M
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YW3803	0	qfxc08-3008	0M
IC-WS001	WS001/YN5999	0	qfxc08-3008	0M
node-device1	BBAK0372	0	qfx3500	0M
node-device1	EE3093	0	qfx3500	0M

Total usage of core repository: 0M of 70000M (0.0%)

### show system core-dumps repository log (QFabric Systems)

```

user@switch> show system core-dumps repository log
Repository scope: shared
Repository head: /pbdata/export
Repository name: log
List of nodes for log repository: /pbdata/export/rlogs/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	0M
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	1	fx-jvre	0M
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	1	fx-jvre	0M
NW-NG-0	BBAK0394	1	qfx3500	0M
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	1	fx-jvre	0M
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	3	fx-jvre	0M
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	1	fx-jvre	0M
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	1	fx-jvre	0M
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YN5999	1	qfxc08-3008	0M
IC-WS001	WS001/YW3803	1	qfxc08-3008	0M

node-device1	BBAK0372	1	qfx3500	0M
node-device1	EE3093	1	qfx3500	0M
Total usage of log repository:0M of 70000M (0.0%)				

## show system directory-usage

<b>List of Syntax</b>	<a href="#">Syntax on page 1319</a> <a href="#">Syntax (EX Series) on page 1319</a> <a href="#">Syntax (TX Matrix Router) on page 1319</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1319</a> <a href="#">Syntax (MX Series Router) on page 1319</a> <a href="#">Syntax (QFX Series) on page 1319</a>
<b>Syntax</b>	<pre>show system directory-usage &lt;depth <i>number</i>&gt; &lt;path&gt;</pre>
<b>Syntax (EX Series)</b>	<pre>show system directory-usage &lt;all-members&gt; &lt;depth <i>number</i>&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;path&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>show system directory-usage &lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt; &lt;depth <i>number</i>&gt; &lt;path&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>show system directory-usage &lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt; &lt;depth <i>number</i>&gt; &lt;path&gt;</pre>
<b>Syntax (MX Series Router)</b>	<pre>show system directory-usage &lt;all-members&gt; &lt;depth <i>number</i>&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;path&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show system directory-usage &lt;depth <i>number</i>&gt; &lt;path&gt; &lt;infrastructure <i>name</i>&gt; &lt;interconnect-device <i>name</i>&gt; &lt;node-group <i>name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Display directory usage information.
<b>Options</b>	<b>none</b> —Display all directory usage information.

**all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display directory usage information about all the T640 routers (in a routing matrix based on a TX Matrix router). Display directory usage information about all the T1600 or T4000 routers (in a routing matrix based on a TX Matrix Plus router) in the chassis.

**all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display directory information for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display directory information for all connected T1600 or T4000 LCCs.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Display directory information for all members of the Virtual Chassis configuration.

**depth *number***—(Optional) Depth of the directory to traverse. This option is useful when you want to limit the output shown for a large file system.

**infrastructure *name***— (QFabric systems only) (Optional) Display directory information for the fabric control Routing Engines and fabric manager Routing Engines.

**interconnect-device *name***— (QFabric systems only) (Optional) Display directory information for the Interconnect device.

**node-group *name***— (QFabric systems only) (Optional) Display directory information for the Node group.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display directory information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display directory information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display directory information for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display directory information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

***path***—(Optional) Path or root directory to traverse.



**scc**—(TX Matrix router only) (Optional) Display directory information for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display directory information for the TX Matrix Plus router. Replace *number* with 0.

**Required Privilege Level**

view

**Related Documentation**

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

[show system directory-usage scc \(TX Matrix Router\) on page 1322](#)  
[show system directory-usage sfc \(TX Matrix Plus Router\) on page 1322](#)  
[show system directory-usage \(QFX3500 Switch\) on page 1322](#)

**Output Fields**

[Table 128 on page 1321](#) describes the output fields for the **show system directory-usage** command. Output fields are listed in the approximate order in which they appear.

**Table 128: show system directory-usage Output Fields**

Field Name	Field Description
<i>bytes</i>	Number of bytes used by files in a directory.
<i>directory-name</i>	Name of the directory.

## Sample Output

### show system directory-usage scc (TX Matrix Router)

```
user@host> show system directory-usage /var/tmp scc
/var/tmp
1.0K    /var/tmp/vi.recover
2.0K    /var/tmp/instmp.tPMk8u
1.0K    /var/tmp/install
        /var/tmp/instmp.GUMpur
4.8M    /var/tmp/instmp.GUMpur/packages
6.4M    /var/tmp/troy1
297M    /var/tmp/dsw
        /var/tmp/pkg_tmp.2073
83K     /var/tmp/pkg_tmp.2073/bin
        /var/tmp/instmp.oMIDb1
89K     /var/tmp/instmp.oMIDb1/bin
        /var/tmp/instmp.byhMjR
4.6M    /var/tmp/instmp.byhMjR/packages
        /var/tmp/instmp.6fqHf3
1.7M    /var/tmp/instmp.6fqHf3/packages
        /var/tmp/instmp.mljECe
4.6M    /var/tmp/instmp.mljECe/packages
```

### show system directory-usage sfc (TX Matrix Plus Router)

```
user@switch> show system directory-usage /var/tmp sfc 0
sfc0-re0:
-----
        /var/tmp
46K     /var/tmp/gres-tp
        /var/tmp/sec-download
2.0K    /var/tmp/sec-download/sub-download
2.0K    /var/tmp/vi.recover
2.0K    /var/tmp/install
795M    /var/tmp/cores
766K    /var/tmp/pr440594
```

### show system directory-usage (QFX3500 Switch)

```
user@switch> show system directory-usage
/var/tmp
30K     /var/tmp/gres-tp
2.0K    /var/tmp/rtbdb
2.0K    /var/tmp/vi.recover
2.0K    /var/tmp/install
2.0K    /var/tmp/pics
```

## show system processes

<b>List of Syntax</b>	<a href="#">Syntax on page 1323</a> <a href="#">Syntax (EX Series Switches) on page 1323</a> <a href="#">Syntax (MX Series Routers) on page 1323</a> <a href="#">Syntax (QFX Series) on page 1323</a> <a href="#">Syntax (TX Matrix Routers) on page 1323</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1323</a>
<b>Syntax</b>	<pre>show system processes &lt;brief   detail   extensive   summary&gt; &lt;health (pid <i>process-identifer</i>   process-name <i>process-name</i>)&gt; &lt;providers&gt; &lt;resource-limits (brief   detail) <i>process-name</i>&gt; &lt;wide&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show system processes &lt;all-members&gt; &lt;brief   detail   extensive   summary&gt; &lt;health (pid <i>process-identifer</i>   process-name <i>process-name</i>)&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;providers&gt; &lt;resource-limits (brief   detail) <i>process-name</i>&gt; &lt;wide&gt;</pre>
<b>Syntax (MX Series Routers)</b>	<pre>show system processes &lt;all-members&gt; &lt;brief   detail   extensive   summary&gt; &lt;health (pid <i>process-identifer</i>   process-name <i>process-name</i>)&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;providers&gt; &lt;resource-limits (brief   detail) <i>process-name</i>&gt; &lt;wide&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show system processes &lt;brief   detail   extensive   summary &gt; &lt;health (pid <i>process-identifer</i>   process-name <i>process-name</i>)&gt; host-processes &lt;interconnect-device <i>name</i>&gt; &lt;node-group <i>name</i>&gt; &lt;providers&gt; &lt;resource-limits&gt; &lt;wide&gt;</pre>
<b>Syntax (TX Matrix Routers)</b>	<pre>show system processes &lt;brief   detail   extensive   summary&gt; &lt;all-chassis  all-lcc   lcc <i>number</i>   scc&gt; &lt;wide&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>show system processes &lt;brief   detail   extensive   summary&gt;</pre>

<all-chassis| all-lcc | lcc *number* | sfc *number*>  
<wide>

<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option <b>sfc</b> introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about software processes that are running on the router or switch and that have controlling terminals.
<b>Options</b>	<p><b>none</b>—Display standard information about system processes.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of detail.</p> <p><b>adaptive-services</b>—(Optional) Display the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.</p> <p><b>alarm-control</b>—(Optional) Display the process to configure the system alarm.</p> <p><b>all-chassis</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display standard system process information about all the T640 routers (in a routing matrix based on the TX Matrix router) or all the T1600 or T4000 routers (in a routing matrix based on the TX Matrix Plus router) in the chassis.</p> <p><b>all-lcc</b>—(TX Matrix routers and TX Matrix Plus router only) (Optional) Display standard system process information for all T640 routers (or line-card chassis) connected to the TX Matrix router. Display standard system process information for all connected T1600 or T4000 LCCs.</p> <p><b>all-members</b>—(EX4200 switches and MX Series routers only) (Optional) Display standard system process information for all members of the Virtual Chassis configuration.</p> <p><b>ancpd-service</b>—Display the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.</p> <p><b>application-identification</b>—Display the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.</p> <p><b>audit-process</b>—(Optional) Display the RADIUS accounting process.</p> <p><b>auto-configuration</b>—Display the Interface Auto-Configuration process.</p> <p><b>bootp</b>—Display the process that enables a router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent. DHCP relaying is disabled.</p>

**captive-portal-content-delivery**—Display the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

**ce-l2tp-service**—(Optional) (M10, M10i, M7i, and MX Series routers only) Display the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

**cfm**—Display Ethernet Operations, Administration, and Maintenance (OAM) connectivity fault management (CFM) process, which can be used to monitor the physical link between two switches.

**chassis-control**—(Optional) Display the chassis management process.

**class-of-service**—(Optional) Display the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

**clksyncd-service**—Display the external clock synchronization process, which uses synchronous Ethernet (SyncE).

**craft-control**—Display the process for the I/O of the craft interface.

**database-replication**—(EX Series switches and MX Series routers only) (Optional) Display the database replication process.

**datapath-trace-service**—Display the packet path tracing process.

**dhcp-service**—(EX Series switches and MX Series routers only) (Optional) Display the Dynamic Host Configuration Protocol process, which enables a DHCP server to allocate network IP addresses and deliver configuration settings to client hosts without user intervention.

**diameter-service**—(Optional) Display the diameter process.

**disk-monitoring**—(Optional) Display the disk monitoring process, which checks the health of the hard disk drive on the Routing Engine.

**dynamic-flow-capture**—(Optional) Display the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

**ecc-error-logging**—(Optional) Display the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

**ethernet-connectivity-fault-management**— Display the process that provides IEEE 802.1ag OAM connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

**ethernet-link-fault-management**—(EX Series switches and MX Series routers only) (Optional) Display the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

**event-processing**—(Optional) Display the event process (eventd).

**firewall**—(Optional) Display the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

**general-authentication-service**—(EX Series switches and MX Series routers only)  
(Optional) Display the general authentication process.

**health (pid *process-identifier* | process-name *process-name*)**—(Optional) Display process health information, either by process id (PID) or by process name.

**host-processes**—Display process information of processes running on the host system.

**iccp-service**—Display the Inter-Chassis Communication Protocol (ICCP) process.

**idp-policy**—Display the intrusion detection and prevention (IDP) protocol process.

**ilmi**—Display the Integrated Local Management Interface (ILMI) protocol process, which provides bidirectional exchange of management information between two ATM interfaces across a physical connection.

**inet-process**—Display the IP multicast family process.

**init**—Display the process that initializes the USB modem.

**interface-control**—(Optional) Display the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

**kernel-replication**—(Optional) Display the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

**l2-learning**—(Optional) Display the Layer 2 address flooding and learning process.

**l2cpd-service**—Display the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

**lACP**—(Optional) Display the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display standard system process information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display standard system process information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display standard system process information for the local Virtual Chassis member.

**local-policy-decision-function**—Display the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

**logical-system-mux**—Display the logical router multiplexer process (lrmuxd), which manages the multiple instances of the routing protocols process (rpd) on a machine running logical routers.

**mac-validation**—Display the MAC validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display standard system process information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**mib-process**—(Optional) Display the MIB II process, which provides the router's MIB II agent.

**mobile-ip**—(Optional) Display the Mobile IP process, which configures Junos OS Mobile IP features.

**moundd-service**—(EX Series switches and MX Series routers only) (Optional) Display the service for NFS mounts requests.

**mpls-traceroute**—(Optional) Display the MPLS Periodic Traceroute process.

**mspd**—(Optional) Display the Multiservice process.

**multicast-snooping**—(EX Series switches and MX Series routers only) (Optional) Display the multicast snooping process, which makes Layer 2 devices such as VLAN switches aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

**named-service**—(Optional) Display the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

**neighbor-liveness**—Display the process, which specifies the maximum length of time that the router waits for its neighbor to re-establish an LDP session.

**nfsd-service**—(Optional) Display the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

**ntp**—Display the Network Time Protocol (NTP) process, which provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network.

**packet-triggered-subscribers**—Display the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

**peer-selection-service**—(Optional) Display the Peer Selection Service process.

**periodic-packet-services**—Display the Periodic packet management process, which is responsible for processing a variety of time-sensitive periodic tasks so that other processes can more optimally direct their resources.

**pfe**—Display the Packet Forwarding Engine management process.

**pgcp-service**—(Optional) Display the pgcpd service process running on the Routing Engine.

**pgm**—Display the Pragmatic General Multicast (PGM) protocol process, which enables a reliable transport layer for multicast applications.

**pic-services-logging**—(Optional) Display the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

**ppp**—(Optional) Display the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

**ppp-service**—Display the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.

**pppoe**—(Optional) Display the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

**process-monitor**—Display the process health monitor process (pmond).

**providers**—(Optional) Display provider processes.

**redundancy-interface-process**—(Optional) Display the ASP redundancy process.

**remote-operations**—(Optional) Display the remote operations process, which provides the ping and traceroute MIBs.



**resource-cleanup**—Display the resource cleanup process.

**resource-limits (brief | detail) process-name**—(Optional) Display process resource limits.

**routing**—(Optional) Display the routing protocol process.

**sampling**—(Optional) Display the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

**sbc-configuration-process**—Display the session border controller (SBC) process of the border signaling gateway (BSG).

**scc**—(TX Matrix routers only) (Optional) Display standard system process information for the TX Matrix router (or switch-card chassis).

**sdk-service**—Display the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

**secure-neighbor-discovery**—(EX Series switches and MX Series routers only) (Optional) Display the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

**send**—(Optional) Display the Secure Neighbor Discovery Protocol (SEND) process, which provides support for protecting Neighbor Discovery Protocol (NDP) messages.

**service-deployment**—(Optional) Display the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

**sfc number**—(TX Matrix Plus routers only) (Optional) Display system process information for the TX Matrix Plus router. Replace **number** with 0.

**snmp**—Display the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

**sonet-aps**—Display the SONET Automatic Protection Switching (APS) process, which monitors any SONET interface that participates in APS.

**static-subscribers**—(Optional) Display the Static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

**tunnel-oamd**—(Optional) Display the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

**vrrp**—(EX Series switches and MX Series routers only) (Optional) Display the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

**watchdog**—Display the watchdog timer process, which enables the watchdog timer when Junos OS encounters a problem.

**wide**—(Optional) Display process information that might be wider than 80 columns.

**Additional Information** By default, when you issue the **show system processes** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation**

- [List of Junos OS Processes](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

- [show system processes on page 1332](#)
- [show system processes brief on page 1333](#)
- [show system processes detail on page 1333](#)
- [show system processes extensive on page 1333](#)
- [show system processes extensive \(EX9200 Switch\) on page 1334](#)
- [show system processes lcc wide \(TX Matrix Routing Matrix\) on page 1334](#)
- [show system processes summary on page 1335](#)
- [show system processes \(TX Matrix Plus Router\) on page 1335](#)
- [show system processes sfc \(TX Matrix Plus Router\) on page 1343](#)
- [show system processes lcc wide \(TX Matrix Plus Routing Matrix\) on page 1345](#)
- [show system processes \(QFX Series\) on page 1347](#)

**Output Fields** [Table 129 on page 1330](#) describes the output fields for the **show system processes** command. Output fields are listed in the approximate order in which they appear.

**Table 129: show system processes Output Fields**

Field Name	Field Description	Level of Output
<b>last pid</b>	Last process identifier assigned to the process.	brief extensive summary
<b>load averages</b>	Three load averages followed by the current time.	brief extensive summary
<b>processes</b>	Number of existing processes and the number of processes in each state ( <b>sleeping</b> , <b>running</b> , <b>starting</b> , <b>zombies</b> , and <b>stopped</b> ).	brief extensive summary
<b>Mem</b>	Information about physical and virtual memory allocation.	brief extensive summary
<b>Swap</b>	Information about physical and virtual memory allocation.	brief extensive summary

Table 129: show system processes Output Fields (*continued*)

Field Name	Field Description	Level of Output
PID	Process identifier.	detail extensive summary
TT	Control terminal name.	none detail
STAT	<p>Symbolic process state. The state is given by a sequence of letters. The first letter indicates the run state of the process:</p> <ul style="list-style-type: none"> <li>• <b>D</b>—In disk or other short-term, uninterruptible wait</li> <li>• <b>I</b>—Idle (sleeping longer than about 20 seconds)</li> <li>• <b>R</b>—Runnable</li> <li>• <b>S</b>—Sleeping for less than 20 seconds</li> <li>• <b>T</b>—Stopped</li> <li>• <b>Z</b>—Dead (zombie)</li> <li>• <b>+</b> —The process is in the foreground process group of its control terminal.</li> <li>• <b>&lt;</b>—The process has raised CPU scheduling priority.</li> <li>• <b>&gt;</b>—The process has specified a soft limit on memory requirements and is currently exceeding that limit; such a process is not swapped.</li> <li>• <b>A</b>—The process requested random page replacement.</li> <li>• <b>E</b>—The process is trying to exit.</li> <li>• <b>L</b>—The process has pages locked in core.</li> <li>• <b>N</b>—The process has reduced CPU scheduling priority.</li> <li>• <b>S</b>—The process requested first-in, first-out (FIFO) page replacement.</li> <li>• <b>s</b>—The process is a session leader.</li> <li>• <b>V</b>—The process is temporarily suspended.</li> <li>• <b>W</b>—The process is swapped out.</li> <li>• <b>X</b>—The process is being traced or debugged.</li> </ul>	none detail
UID	User identifier.	detail
USERNAME	Process owner.	extensive summary
PPID	Parent process identifier.	detail
CPU	<p>(D)—Short-term CPU usage.</p> <p>(E and S)—Raw (unweighted) CPU usage. The value of this field is used to sort the processes in the output.</p>	detail extensive summary
RSS	Resident set size.	detail
WCHAN	Symbolic name of the wait channel.	detail
STARTED	Local time when the process started running.	detail
PRI	Current priority of the process. A lower number indicates a higher priority.	detail extensive summary

Table 129: show system processes Output Fields (*continued*)

Field Name	Field Description	Level of Output
NI or NICE	UNIX "niceness" value. A lower number indicates a higher priority.	detail extensive summary
SIZE	Total size of the process (text, data, and stack), in kilobytes.	extensive summary
RES	Current amount of resident memory, in kilobytes.	extensive summary
STATE	Current state of the process (for example, <b>sleep</b> , <b>wait</b> , <b>run</b> , <b>idle</b> , <b>zombie</b> , or <b>stop</b> ).	extensive summary
TIME	(S)—Number of system and user CPU seconds that the process has used.  (None, D, and E)—Total amount of time that the command has been running.	detail extensive summary
WCPU	Weighted CPU usage.	extensive summary
COMMAND	Command that is currently running.	detail extensive summary
THR	Number of threads in the process	extensive

## Sample Output

### show system processes

```

user@host> show system processes
PID TT  STAT      TIME COMMAND
  0  ??  DLs      0:00.70 (swapper)
  1  ??  Is       0:00.35 /sbin/init --
  2  ??  DL       0:00.00 (pagedaemon)
  3  ??  DL       0:00.00 (vmdaemon)
  4  ??  DL       0:42.37 (update)
  5  ??  DL       0:00.00 (if_jnx)
 80  ??  Ss       0:14.66 syslogd -s
 96  ??  Is       0:00.01 portmap
128  ??  Is       0:02.70 cron
173  ??  Is       0:02.24 /usr/local/sbin/sshd (sshd1)
189  ??  S        0:03.80 /sbin/watchdog -t180
190  ??  I        0:00.03 /usr/sbin/tnetd -N
191  ??  S        2:24.76 /sbin/ifd -N
192  ??  S<       0:55.44 /usr/sbin/xntpd -N
195  ??  S        0:53.11 /usr/sbin/snmpd -N
196  ??  S        1:15.73 /usr/sbin/mib2d -N
198  ??  I        0:00.75 /usr/sbin/inetd -N
2677 ??  I        0:00.01 /usr/sbin/mgd -N
2712 ??  Ss       0:00.24 rlogind
2735 ??  R        0:00.00 /bin/ps -ax
1985 p0-  S        0:07.41 ./rpd -N
2713 p0  Is       0:00.24 -tcsh (tcsh)
2726 p0  S+       0:00.07 cli

```

## show system processes brief

```

user@host> show system processes brief
last pid: 543; load averages: 0.00, 0.00, 0.00 18:29:47
37 processes: 1 running, 36 sleeping

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

```

## show system processes detail

```

user@host> show system processes detail

```

PID	UID	PPID	CPU	PRI	NI	RSS	WCHAN	STARTED	TT	STAT	TIME	COMMAND
3151	1049	3129	2	28	0	672	-	1:13PM	p0	R+	0:00.00	ps -ax -r
1	0	0	0	10	0	376	wait	1:51PM	??	Is	0:00.29	/sbin/ini
2	0	0	0	-18	0	12	psleep	1:51PM	??	DL	0:00.00	(pagedae
3	0	0	0	28	0	12	psleep	1:51PM	??	DL	0:00.00	(vmdaemo
4	0	0	0	28	0	12	update	1:51PM	??	DL	0:07.15	(update)
5	0	0	0	2	0	12	pfesel	1:51PM	??	IL	0:02.90	(if_pfe)
27	0	1	0	10	0	17936	mfsidl	1:51PM	??	Is	0:00.46	mfs /dev/
81	0	1	0	2	0	496	select	1:52PM	??	Ss	0:31.21	syslogd -
119	1	1	0	2	0	492	select	1:52PM	??	Is	0:00.00	portmap
134	0	1	0	2	0	580	select	1:52PM	??	S	0:02.95	amd -p -a
151	0	1	0	18	0	532	pause	1:52PM	??	Is	0:00.34	cron
183	0	1	0	2	0	420	select	1:52PM	??	Ss	0:00.07	/usr/loca
206	0	1	0	18	0	72	pause	1:52PM	??	S	0:00.51	/sbin/wat
207	0	1	0	2	0	520	select	1:52PM	??	I	0:00.16	/usr/sbin
208	0	1	0	2	0	536	select	1:52PM	??	S	0:08.21	/sbin/dcd
210	0	1	255	2	-12	740	select	1:52PM	??	S<	0:05.83	/usr/sbin
211	0	1	0	2	0	376	select	1:52PM	??	S	0:00.03	/usr/sbin
215	0	1	0	2	0	548	select	1:52PM	??	I	0:00.50	/usr/sbin
219	0	1	0	3	0	540	ttyin	1:52PM	v0	Is+	0:00.02	/usr/libe
220	0	1	0	3	0	540	ttyin	1:52PM	v1	Is+	0:00.01	/usr/libe
221	0	1	0	3	0	540	ttyin	1:52PM	v2	Is+	0:00.01	/usr/libe
222	0	1	0	3	0	540	ttyin	1:52PM	v3	Is+	0:00.01	/usr/libe
735	0	1	0	2	0	468	select	2:47PM	??	S	0:19.14	/usr/sbin
736	0	1	0	2	0	212	select	2:47PM	??	S	0:14.13	/usr/sbin
1380	0	1	0	3	0	888	ttyin	7:32PM	d0	Is+	0:00.46	bash
3019	0	207	0	2	0	636	select	10:49AM	??	Ss	0:02.93	tnp.chass
3122	0	1380	0	2	0	1764	select	12:33PM	d0	S	0:00.77	./rpd -N
3128	0	215	0	2	0	580	select	12:45PM	??	Ss	0:00.12	rlogind
3129	1049	3128	0	18	0	944	pause	12:45PM	p0	Ss	0:00.14	-tcsh (tc
0	0	0	0	-18	0	0	sched	1:51PM	??	DLs	0:00.10	(swapper

## show system processes extensive

```

user@host> show system processes extensive
Mem: 241M Active, 99M Inact, 78M Wired, 325M Cache, 69M Buf, 1251M Free
Swap: 2048M Total, 2048M Free

```

PID	USERNAME	THR	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	COMMAND
11	root	1	171	52	OK	12K	RUN	807.5H	98.73%	idle
13	root	1	-20	-139	OK	12K	WAIT	36:17	0.00%	swi7: clock sio
1499	root	1	96	0	7212K	3040K	select	34:01	0.00%	license-check
1621	root	1	96	0	20968K	11216K	select	20:25	0.00%	mib2d
1465	root	2	8	-88	115M	11748K	nanslp	14:32	0.00%	chassisd
1478	root	1	96	0	6336K	3816K	select	11:28	0.00%	ppmd
20	root	1	-68	-187	OK	12K	WAIT	10:28	0.00%	irq10: em0 em1+++*
1490	root	1	96	0	11792K	4336K	select	9:44	0.00%	shm-rtssdbd

```

1618 root      1  96    0 39584K  7464K select  8:47  0.00% pfed
1622 root      1  96    0 15268K 10988K select  6:16  0.00% snmpd
1466 root      1  96    0  7408K  2896K select  5:44  0.00% alarmd
   7 root      1 -16    0    0K    12K client  5:09  0.00% ifstate notify
1480 root      1  96    0  5388K  2660K select  4:29  0.00% ksyncd
  12 root      1 -40 -159    0K    12K WAIT   4:15  0.00% swi2: netisr 0
1462 root      1  96    0 1836K   1240K select  3:57  0.00% bslockd
  55 root      1 -16    0    0K    12K -      3:44  0.00% schedcpu
1392 root      1  16    0    0K    12K bcmsem  3:37  0.00% bcmLINK.0
  47 root      1 -16    0    0K    12K psleep  3:25  0.00% vmkmemdaemon
  36 root      1  20    0    0K    12K syncer  2:46  0.00% syncer
1484 root      1  96    0  7484K  3428K select  2:38  0.00% clksyncd
1616 root      1  96    0  4848K  2848K select  2:18  0.00% irsd
1487 root      1  96    0 32800K  6992K select  2:10  0.00% smid
1623 root      1  96    0 34616K  5464K select  2:01  0.00% dcd
  15 root      1 -16    0    0K    12K -      1:59  0.00% yarrow
  49 root      1 -16    0    0K    12K .       1:51  0.00% ddostasks

```

### show system processes extensive (EX9200 Switch)

```

user@switch> show system processes extensive
last pid: 3372; load averages:  0.02,  0.02,  0.00 up 0+01:42:22   16:39:57
151 processes: 4 running, 131 sleeping, 1 zombie, 15 waiting

```

```

Mem: 935M Active, 122M Inact, 108M Wired, 838M Cache, 214M Buf, 5872M Free
Swap: 8192M Total, 8192M Free

```

PID	USERNAME	THR	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	COMMAND
10	root	1	171	52	0K	16K	RUN	96:34	92.19%	idle
3317	root	1	97	0	40412K	30944K	select	0:00	5.13%	mgd
3316	root	1	96	0	26672K	20516K	select	0:00	3.08%	cli
1626	root	2	8	-88	124M	20332K	nanslp	3:19	2.39%	chassisd
260	root	1	-8	0	0K	16K	mdwait	0:16	0.00%	md16
19	root	1	-68	-187	0K	16K	WAIT	0:12	0.00%	irq11: em0 em1
em2*										
1642	root	1	96	0	8052K	3936K	RUN	0:10	0.00%	clksyncd
11	root	1	-20	-139	0K	16K	WAIT	0:07	0.00%	swi7: clock sio
154	root	1	-8	0	0K	16K	mdwait	0:06	0.00%	md8
1784	root	1	96	0	98M	33720K	select	0:05	0.00%	authd
1646	root	1	96	0	7776K	2944K	select	0:03	0.00%	license-check
1807	root	1	96	0	41340K	9944K	select	0:02	0.00%	mib2d

[...Output truncated...]

### show system processes lcc wide (TX Matrix Routing Matrix)

```

user@host> show system processes lcc 2 wide
lcc2-re0:

```

PID	TT	STAT	TIME	COMMAND
0	??	DLs	0:00.00	(swapper)
1	??	ILs	0:00.10	/sbin/preinit -- (init)
2	??	DL	0:00.00	(pagedaemon)
3	??	DL	0:00.00	(vmdaemon)
4	??	DL	0:00.00	(bufdaemon)
5	??	DL	0:00.04	(syncer)
6	??	DL	0:00.00	(netdaemon)
7	??	IL	0:00.00	(if_pic_listen)
8	??	IL	0:00.00	(scs_housekeeping)
9	??	IL	0:00.00	(if_pfe_listen)
10	??	DL	0:00.00	(vmuncachedaemon)

```

11  ??  SL    0:00.02  (cb_poll)
172  ??  ILs    0:00.21  mfs -o noauto /dev/ad1s1b /tmp (newfs)
2909 ??  Is    0:00.00  pccardd
2932 ??  Ss    0:00.07  syslogd -r -s
3039 ??  Is    0:00.00  cron
3217 ??  I     0:00.00  /sbin/watchdog -d
3218 ??  I     0:00.02  /usr/sbin/tnetd -N
3221 ??  S     0:00.11  /usr/sbin/alarmd -N
3222 ??  S     0:00.85  /usr/sbin/craftd -N
3223 ??  S     0:00.05  /usr/sbin/mgd -N
3224 ??  I     0:00.02  /usr/sbin/inetd -N
3225 ??  I     0:00.00  /usr/sbin/tnp.sntpd -N
3226 ??  I     0:00.01  /usr/sbin/tnp.sntpc -N
3228 ??  I     0:00.01  /usr/sbin/smartd -N
3231 ??  I     0:00.01  /usr/sbin/eccd -N
3425 ??  S     0:00.09  /usr/sbin/dfwd -N
3426 ??  S     0:00.19  /sbin/dcd -N
3427 ??  I     0:00.04  /usr/sbin/pfed -N
3430 ??  S     0:00.10  /usr/sbin/ksyncd -N
3482 ??  S     1:53.63  /usr/sbin/chassisd -N
4285 ??  SL    0:00.01  (peer proxy)
4286 ??  SL    0:00.00  (peer proxy)
4303 ??  Ss    0:00.00  mgd: (mgd) (root) (mgd)
4304 ??  R     0:00.00  /bin/ps -ax -ww
3270 d0  Is+  0:00.00  /usr/libexec/getty std.9600 ttyd0

```

#### show system processes summary

```

user@host> show system processes summary
last pid: 543; load averages: 0.00, 0.00, 0.00 18:29:47
37 processes: 1 running, 36 sleeping

```

```

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

```

PID	USERNAME	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	CPU	COMMAND
527	root	2	0	176K	580K	select	0:00	0.04%	0.04%	rlogind
543	root	30	0	604K	768K	RUN	0:00	0.00%	0.00%	top

#### show system processes (TX Matrix Plus Router)

```

user@host> show system processes
sfc0-re0:

```

```

-----
PID  TT  STAT  TIME  COMMAND
0  ??  Wls    0:00.00  [swapper]
1  ??  ILs    0:00.18  /packages/mnt/jbase/sbin/init --
2  ??  DL     0:00.20  [g_event]
3  ??  DL     0:00.39  [g_up]
4  ??  DL     0:00.32  [g_down]
5  ??  DL     0:00.00  [thread taskq]
6  ??  DL     0:00.09  [kqueue taskq]
7  ??  DL     0:00.01  [pagedaemon]
8  ??  DL     0:00.00  [vmdaemon]
9  ??  DL     0:06.63  [pagezero]
10  ??  DL     0:00.00  [ktrace]
11  ??  RL    310:52.98  [idle]
12  ??  WL     0:11.03  [swi2: net]
13  ??  WL     0:27.58  [swi7: clock sio]
14  ??  WL     0:00.00  [swi6: vm]
15  ??  DL     0:03.02  [yarrow]

```

```

16 ?? WL 0:00.00 [swi9: +]
17 ?? WL 0:00.00 [swi8: +]
18 ?? WL 0:00.00 [swi5: cambio]
19 ?? WL 0:00.00 [swi9: task queue]
20 ?? WL 0:11.41 [irq16: uhci0 uhci*]
21 ?? DL 0:00.00 [usb0]
22 ?? DL 0:00.00 [usbtask]
23 ?? WL 0:39.51 [irq17: uhci1 uhci*]
24 ?? DL 0:00.00 [usb1]
25 ?? WL 0:00.00 [irq18: uhci2 uhci*]
26 ?? DL 0:00.83 [usb2]
27 ?? DL 0:00.00 [usb3]
28 ?? DL 0:00.00 [usb4]
29 ?? DL 0:00.00 [usb5]
30 ?? DL 0:00.73 [usb6]
31 ?? DL 0:00.00 [usb7]
32 ?? WL 0:00.00 [irq14: ata0]
33 ?? WL 0:00.00 [irq15: ata1]
34 ?? WL 0:00.00 [irq1: atkbd0]
35 ?? WL 0:00.00 [swi0: sio]
36 ?? WL 0:00.00 [irq11: isab0]
37 ?? WL 0:00.00 [swi3: ip6opt ipopt]
38 ?? WL 0:00.00 [swi4: ip6mismatch+]
39 ?? WL 0:00.00 [swi1: ipfwd]
40 ?? DL 0:00.02 [bufdaemon]
41 ?? DL 0:00.02 [vnlr]
42 ?? DL 0:00.39 [syncer]
43 ?? DL 0:00.05 [softdepflush]
44 ?? DL 0:00.00 [netdaemon]
45 ?? DL 0:00.02 [vmuncachedaemon]
46 ?? DL 0:00.00 [if_pic_listen]
47 ?? DL 0:00.35 [vmkmemdaemon]
48 ?? DL 0:00.00 [cb_poll]
49 ?? DL 0:00.06 [if_pfe_listen]
50 ?? DL 0:00.00 [scs_housekeeping]
51 ?? IL 0:00.00 [kern_dump_proc]
52 ?? IL 0:00.00 [nfsiod 0]
53 ?? IL 0:00.00 [nfsiod 1]
54 ?? IL 0:00.00 [nfsiod 2]
55 ?? IL 0:00.00 [nfsiod 3]
56 ?? DL 0:00.37 [schedcpu]
57 ?? DL 0:00.56 [md0]
79 ?? DL 0:02.58 [md1]
100 ?? DL 0:00.03 [md2]
118 ?? DL 0:00.01 [md3]
139 ?? DL 0:00.95 [md4]
160 ?? DL 0:00.12 [md5]
181 ?? DL 0:00.00 [md6]
217 ?? DL 0:00.02 [md7]
227 ?? DL 0:00.05 [md8]
1341 ?? SL 0:01.34 [bcmTX]
1342 ?? SL 0:01.68 [bcmXGS3AsyncTX]
1343 ?? SL 0:41.40 [bcmLINK.0]
1345 ?? SL 0:33.83 [bcmLINK.1]
1350 ?? Is 0:00.01 /usr/sbin/cron
1502 ?? S 0:00.01 /sbin/watchdog -t-1
1503 ?? S 0:00.86 /usr/libexec/bslockd -mp -N
1504 ?? S 0:00.01 /usr/sbin/tnetd -N
1507 ?? S 0:01.32 /usr/sbin/alarmd -N
1508 ?? S 0:14.54 /usr/sbin/craftd -N
1509 ?? S 0:01.19 /usr/sbin/mgd -N

```



```

1512 ?? I      0:00.05 /usr/sbin/inetd -N
1513 ?? S      0:00.10 /usr/sbin/tnp.sntpd -N
1517 ?? S      0:00.11 /usr/sbin/smartd -N
1525 ?? S      0:01.10 /usr/sbin/idpd -N
1526 ?? S      0:01.43 /usr/sbin/license-check -U -M -p 10 -i 10
1527 ?? I      0:00.01 /usr/libexec/getty Pc ttyv0
1616 ?? DL     0:00.30 [peer proxy]
1617 ?? DL     0:00.32 [peer proxy]
1618 ?? DL     0:00.34 [peer proxy]
1619 ?? DL     0:00.30 [peer proxy]
2391 ?? Is     0:00.01 telnetd
7331 ?? Ss     0:00.03 telnetd
9538 ?? DL     0:01.16 [jsr_kkcm]
9613 ?? DL     0:00.18 [peer proxy]
23781 ?? Ss     0:00.01 telnetd
23926 ?? Ss     0:00.01 mgd: (mgd) (regress)/dev/tty2 (mgd)
36867 ?? S      0:03.14 /usr/sbin/rpd -N
36874 ?? S      0:00.08 /usr/sbin/lmpd
36876 ?? S      0:00.17 /usr/sbin/lacpd -N
36877 ?? S      0:00.15 /usr/sbin/bfdd -N
36878 ?? S      0:05.05 /usr/sbin/ppmd -N
36907 ?? S      0:25.07 /usr/sbin/chassisd -N
37775 ?? S      0:00.01 /usr/sbin/bdbrepd -N
45727 ?? S      0:00.02 /usr/sbin/xntpd -j -N -g (ntpd)
45729 ?? S      0:00.38 /usr/sbin/l2ald -N
45730 ?? S<     0:00.12 /usr/sbin/apspd -N
45731 ?? SN     0:00.10 /usr/sbin/sampled -N
45732 ?? S      0:00.03 /usr/sbin/ilmid -N
45733 ?? S      0:00.09 /usr/sbin/rmopd -N
45734 ?? S      0:00.30 /usr/sbin/cosd
45735 ?? I      0:00.00 /usr/sbin/rtspd -N
45736 ?? S      0:00.06 /usr/sbin/fsad -N
45737 ?? S      0:00.05 /usr/sbin/rdd -N
45738 ?? S      0:00.10 /usr/sbin/pppd -N
45739 ?? S      0:00.05 /usr/sbin/dfcd -N
45740 ?? S      0:00.07 /usr/sbin/lfmd -N
45741 ?? S      0:00.01 /usr/sbin/mplsoamd -N
45742 ?? I      0:00.01 /usr/sbin/sendd -N
45743 ?? S      0:00.08 /usr/sbin/appidd -N
45744 ?? S      0:00.05 /usr/sbin/mspd -N
45745 ?? S      0:00.25 /usr/sbin/jdiameterd -N
45746 ?? S      0:00.10 /usr/sbin/pfed -N
45747 ?? S      0:00.19 /usr/sbin/lpdfd -N
45748 ?? S      0:00.63 /sbin/dcd -N
45750 ?? S      0:00.45 /usr/sbin/mib2d -N
45751 ?? S      0:00.15 /usr/sbin/dfwd -N
45752 ?? S      0:00.15 /usr/sbin/irsd -N
45764 ?? S      0:20.59 /usr/sbin/snmpd -N
56479 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
56480 ?? R      0:00.00 /bin/ps -ax
1142 d0- I      0:00.01 /usr/sbin/usbd -N
1160 d0- S      0:29.17 /usr/sbin/eventd -N -r -s -A
6527 d0 Is+     0:00.00 /usr/libexec/getty std.9600 ttyd0
2392 p1 Is      0:00.00 login [pam] (login)
2393 p1 I        0:00.00 -csh (csh)
2394 p1 I        0:00.00 su -
2395 p1 I+       0:00.01 -su (csh)
23782 p2 Is      0:00.00 login [pam] (login)
23881 p2 I        0:00.00 -csh (csh)
23925 p2 S+      0:00.03 cli
7332 p3 Is      0:00.00 login [pam] (login)

```

```

7333 p3 I      0:00.00 -csh (csh)
23780 p3 S+    0:00.02 telnet aj

```

```
lcc0-re0:
```

```

-----
PID TT  STAT      TIME COMMAND
  0 ??  WLS      0:00.00 [swapper]
  1 ??  ILs      0:00.16 /packages/mnt/jbase/sbin/init --
  2 ??  DL       0:00.01 [g_event]
  3 ??  DL       0:00.16 [g_up]
  4 ??  DL       0:00.11 [g_down]
  5 ??  DL       0:00.00 [thread taskq]
  6 ??  DL       0:00.00 [kqueue taskq]
  7 ??  DL       0:00.00 [pagedaemon]
  8 ??  DL       0:00.00 [vmdaemon]
  9 ??  DL       0:01.77 [pagezero]
10 ??  DL       0:00.00 [ktrace]
11 ??  RL      17:22.31 [idle]
12 ??  WL       0:00.32 [swi2: net]
13 ??  WL       0:01.21 [swi7: clock sio]
14 ??  WL       0:00.00 [swi6: vm]
15 ??  DL       0:00.10 [yarrow]
16 ??  WL       0:00.00 [swi9: +]
17 ??  WL       0:00.00 [swi8: +]
18 ??  WL       0:00.00 [swi5: cambio]
19 ??  WL       0:00.00 [swi9: task queue]
20 ??  WL       0:02.73 [irq10: bcm0 uhci1*]
21 ??  WL       0:00.02 [irq11: cb0 uhci0+*]
22 ??  DL       0:00.00 [usb0]
23 ??  DL       0:00.00 [usbtask]
24 ??  DL       0:00.00 [usb1]
25 ??  DL       0:00.05 [usb2]
26 ??  DL       0:00.00 [usb3]
27 ??  DL       0:00.00 [usb4]
28 ??  DL       0:00.00 [usb5]
29 ??  DL       0:00.04 [usb6]
30 ??  DL       0:00.00 [usb7]
31 ??  WL       0:00.00 [irq14: ata0]
32 ??  WL       0:00.00 [irq15: ata1]
33 ??  WL       0:00.00 [irq1: atkbd0]
34 ??  WL       0:00.00 [swi0: sio]
35 ??  WL       0:00.00 [swi3: ip6opt ipopt]
36 ??  WL       0:00.00 [swi4: ip6mismatch+]
37 ??  WL       0:00.00 [swi1: ipfwd]
38 ??  DL       0:00.00 [bufdaemon]
39 ??  DL       0:00.00 [vn1ru]
40 ??  DL       0:00.01 [syncer]
41 ??  DL       0:00.00 [softdepflush]
42 ??  DL       0:00.00 [netdaemon]
43 ??  DL       0:00.00 [vmuncachedaemon]
44 ??  DL       0:00.00 [if_pic_listen]
45 ??  DL       0:00.02 [vmkmemdaemon]
46 ??  DL       0:00.01 [cb_poll]
47 ??  DL       0:00.00 [if_pfe_listen]
48 ??  DL       0:00.00 [scs_housekeeping]
49 ??  IL       0:00.00 [kern_dump_proc]
50 ??  IL       0:00.00 [nfsiod 0]
51 ??  IL       0:00.00 [nfsiod 1]
52 ??  IL       0:00.00 [nfsiod 2]
53 ??  IL       0:00.00 [nfsiod 3]
54 ??  DL       0:00.01 [schedcpu]

```

```

55 ?? DL 0:00.73 [md0]
77 ?? DL 0:03.54 [md1]
98 ?? DL 0:00.37 [md2]
116 ?? DL 0:00.02 [md3]
137 ?? DL 0:00.56 [md4]
158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1078 ?? DL 0:00.00 [jsr_kkcm]
1363 ?? SL 0:00.09 [bcmTX]
1364 ?? SL 0:00.10 [bcmXGS3AsyncTX]
1365 ?? SL 0:03.08 [bcmLINK.0]
1370 ?? Is 0:00.00 /usr/sbin/cron
1522 ?? S 0:00.00 /sbin/watchdog -t-1
1523 ?? S 0:00.05 /usr/libexec/bslockd -mp -N
1524 ?? I 0:00.01 /usr/sbin/tnetd -N
1526 ?? S 0:04.98 /usr/sbin/chassisd -N
1527 ?? S 0:00.04 /usr/sbin/alarmd -N
1528 ?? I 0:00.40 /usr/sbin/craftd -N
1529 ?? S 0:00.08 /usr/sbin/mgd -N
1532 ?? I 0:00.04 /usr/sbin/inetd -N
1533 ?? I 0:00.00 /usr/sbin/tnp.snptd -N
1534 ?? I 0:00.00 /usr/sbin/tnp.snptc -N
1536 ?? S 0:00.01 /usr/sbin/smartd -N
1540 ?? I 0:00.07 /usr/sbin/jcsd -N
1541 ?? S 0:00.11 /usr/sbin/idpd -N
1542 ?? I 0:00.00 /usr/libexec/getty Pc ttyv0
2089 ?? DL 0:00.01 [peer proxy]
2090 ?? DL 0:00.01 [peer proxy]
2091 ?? DL 0:00.01 [peer proxy]
2657 ?? S 0:00.02 /usr/sbin/dfwd -N
2658 ?? S 0:00.02 /sbin/dcd -N
2659 ?? S 0:00.05 /usr/sbin/snmpd -N
2660 ?? S 0:00.01 /usr/sbin/mib2d -N
2661 ?? S 0:00.01 /usr/sbin/pfed -N
2662 ?? S 0:00.01 /usr/sbin/irsd -N
2667 ?? S 0:00.13 /usr/sbin/ksyncd -N
2690 ?? Ss 0:00.00 mgd: (mgd) (root) (mgd)
2691 ?? R 0:00.00 /bin/ps -ax
1164 d0- S 0:00.00 /usr/sbin/usbd -N
1182 d0- S 0:00.34 /usr/sbin/eventd -N -r -s -A
1543 d0 Is+ 0:00.00 /usr/libexec/getty std.9600 ttyd0

```

```
lcc1-re0:
```

```

-----
PID TT STAT TIME COMMAND
0 ?? Wls 0:00.00 [swapper]
1 ?? ILs 0:00.17 /packages/mnt/jbase/sbin/init --
2 ?? DL 0:00.01 [g_event]
3 ?? DL 0:00.16 [g_up]
4 ?? DL 0:00.11 [g_down]
5 ?? DL 0:00.00 [thread taskq]
6 ?? DL 0:00.00 [kqueue taskq]
7 ?? DL 0:00.00 [pagedaemon]
8 ?? DL 0:00.00 [vmdaemon]
9 ?? DL 0:01.77 [pagezero]
10 ?? DL 0:00.00 [ktrace]
11 ?? RL 17:22.83 [idle]
12 ?? WL 0:00.35 [swi2: net]
13 ?? WL 0:01.20 [swi7: clock sio]

```

```

14 ?? WL 0:00.00 [swi6: vm]
15 ?? DL 0:00.10 [yarrow]
16 ?? WL 0:00.00 [swi9: +]
17 ?? WL 0:00.00 [swi8: +]
18 ?? WL 0:00.00 [swi5: cambio]
19 ?? WL 0:00.00 [swi9: task queue]
20 ?? WL 0:02.87 [irq10: bcm0 uhci1*]
21 ?? WL 0:00.02 [irq11: cb0 uhci0+*]
22 ?? DL 0:00.00 [usb0]
23 ?? DL 0:00.00 [usbtask]
24 ?? DL 0:00.00 [usb1]
25 ?? DL 0:00.05 [usb2]
26 ?? DL 0:00.00 [usb3]
27 ?? DL 0:00.00 [usb4]
28 ?? DL 0:00.00 [usb5]
29 ?? DL 0:00.04 [usb6]
30 ?? DL 0:00.00 [usb7]
31 ?? WL 0:00.00 [irq14: ata0]
32 ?? WL 0:00.00 [irq15: ata1]
33 ?? WL 0:00.00 [irq1: atkbd0]
34 ?? WL 0:00.00 [swi0: sio]
35 ?? WL 0:00.00 [swi3: ip6opt ipopt]
36 ?? WL 0:00.00 [swi4: ip6mismatch+]
37 ?? WL 0:00.00 [swi1: ipfwd]
38 ?? DL 0:00.00 [bufdaemon]
39 ?? DL 0:00.00 [vnlru]
40 ?? DL 0:00.01 [syncer]
41 ?? DL 0:00.00 [softdepflush]
42 ?? DL 0:00.00 [netdaemon]
43 ?? DL 0:00.00 [vmuncachedaemon]
44 ?? DL 0:00.00 [if_pic_listen]
45 ?? DL 0:00.02 [vmkmemdaemon]
46 ?? DL 0:00.01 [cb_poll]
47 ?? DL 0:00.00 [if_pfe_listen]
48 ?? DL 0:00.00 [scs_housekeeping]
49 ?? IL 0:00.00 [kern_dump_proc]
50 ?? IL 0:00.00 [nfsiod 0]
51 ?? IL 0:00.00 [nfsiod 1]
52 ?? IL 0:00.00 [nfsiod 2]
53 ?? IL 0:00.00 [nfsiod 3]
54 ?? DL 0:00.02 [schedcpu]
55 ?? DL 0:00.75 [md0]
77 ?? DL 0:03.40 [md1]
98 ?? DL 0:00.37 [md2]
116 ?? DL 0:00.02 [md3]
137 ?? DL 0:00.56 [md4]
158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1052 ?? DL 0:00.00 [jsr_kkcm]
1337 ?? SL 0:00.09 [bcmTX]
1338 ?? SL 0:00.10 [bcmXGS3AsyncTX]
1339 ?? SL 0:03.10 [bcmLINK.0]
1344 ?? Is 0:00.00 /usr/sbin/cron
1496 ?? S 0:00.00 /sbin/watchdog -t-1
1497 ?? S 0:00.05 /usr/libexec/bslockd -mp -N
1498 ?? I 0:00.01 /usr/sbin/tnetd -N
1500 ?? S 0:04.97 /usr/sbin/chassisd -N
1501 ?? S 0:00.04 /usr/sbin/alarmd -N
1502 ?? I 0:00.40 /usr/sbin/craftd -N

```

```

1503 ?? S      0:00.08 /usr/sbin/mgd -N
1506 ?? I      0:00.04 /usr/sbin/inetd -N
1507 ?? I      0:00.00 /usr/sbin/tnp.sntpd -N
1508 ?? I      0:00.00 /usr/sbin/tnp.sntpc -N
1510 ?? S      0:00.01 /usr/sbin/smartd -N
1514 ?? I      0:00.07 /usr/sbin/jcsd -N
1515 ?? S      0:00.18 /usr/sbin/idpd -N
1516 ?? I      0:00.00 /usr/libexec/getty Pc ttyv0
2068 ?? DL     0:00.01 [peer proxy]
2069 ?? DL     0:00.01 [peer proxy]
2070 ?? DL     0:00.01 [peer proxy]
2666 ?? S      0:00.02 /sbin/dcd -N
2667 ?? S      0:00.01 /usr/sbin/irsd -N
2668 ?? S      0:00.01 /usr/sbin/pfed -N
2669 ?? S      0:00.05 /usr/sbin/snmpd -N
2670 ?? S      0:00.01 /usr/sbin/mib2d -N
2671 ?? S      0:00.02 /usr/sbin/dfwd -N
2675 ?? S      0:00.13 /usr/sbin/ksyncd -N
2699 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
2700 ?? R      0:00.00 /bin/ps -ax
1138 d0- S      0:00.00 /usr/sbin/usbd -N
1156 d0- S      0:00.37 /usr/sbin/eventd -N -r -s -A
1517 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0

```

lcc2-re0:

```

-----
PID TT  STAT    TIME COMMAND
  0 ??  Wls    0:00.00 [swapper]
  1 ??  ILs    0:00.18 /packages/mnt/jbase/sbin/init --
  2 ??  DL     0:00.01 [g_event]
  3 ??  DL     0:00.17 [g_up]
  4 ??  DL     0:00.12 [g_down]
  5 ??  DL     0:00.00 [thread taskq]
  6 ??  DL     0:00.00 [kqueue taskq]
  7 ??  DL     0:00.00 [pagedaemon]
  8 ??  DL     0:00.00 [vmdaemon]
  9 ??  DL     0:01.77 [pagezero]
 10 ??  DL     0:00.00 [ktrace]
 11 ??  RL    17:19.13 [idle]
 12 ??  WL     0:00.36 [swi2: net]
 13 ??  WL     0:01.20 [swi7: clock sio]
 14 ??  WL     0:00.00 [swi6: vm]
 15 ??  DL     0:00.13 [yarrow]
 16 ??  WL     0:00.00 [swi9: +]
 17 ??  WL     0:00.00 [swi8: +]
 18 ??  WL     0:00.00 [swi5: cambio]
 19 ??  WL     0:00.00 [swi9: task queue]
 20 ??  WL     0:03.03 [irq10: bcm0 uhci1*]
 21 ??  WL     0:00.02 [irq11: cb0 uhci0+*]
 22 ??  DL     0:00.00 [usb0]
 23 ??  DL     0:00.00 [usbtask]
 24 ??  DL     0:00.00 [usb1]
 25 ??  DL     0:00.05 [usb2]
 26 ??  DL     0:00.00 [usb3]
 27 ??  DL     0:00.00 [usb4]
 28 ??  DL     0:00.00 [usb5]
 29 ??  DL     0:00.04 [usb6]
 30 ??  DL     0:00.00 [usb7]
 31 ??  WL     0:00.00 [irq14: ata0]
 32 ??  WL     0:00.00 [irq15: ata1]
 33 ??  WL     0:00.00 [irq1: atkbd0]

```

```

34 ?? WL 0:00.00 [swi0: sio]
35 ?? WL 0:00.00 [swi3: ip6opt ipopt]
36 ?? WL 0:00.00 [swi4: ip6mismatch+]
37 ?? WL 0:00.00 [swi1: ipfwd]
38 ?? DL 0:00.00 [bufdaemon]
39 ?? DL 0:00.00 [vn1ru]
40 ?? DL 0:00.01 [syncer]
41 ?? DL 0:00.00 [softdepflush]
42 ?? DL 0:00.00 [netdaemon]
43 ?? DL 0:00.00 [vmuncachedaemon]
44 ?? DL 0:00.00 [if_pic_listen]
45 ?? DL 0:00.02 [vmkmemdaemon]
46 ?? DL 0:00.01 [cb_poll]
47 ?? DL 0:00.00 [if_pfe_listen]
48 ?? DL 0:00.00 [scs_housekeeping]
49 ?? IL 0:00.00 [kern_dump_proc]
50 ?? IL 0:00.00 [nfsiod 0]
51 ?? IL 0:00.00 [nfsiod 1]
52 ?? IL 0:00.00 [nfsiod 2]
53 ?? IL 0:00.00 [nfsiod 3]
54 ?? DL 0:00.02 [schedcpu]
55 ?? DL 0:00.75 [md0]
77 ?? DL 0:03.48 [md1]
98 ?? DL 0:00.59 [md2]
116 ?? DL 0:00.02 [md3]
137 ?? DL 0:00.56 [md4]
158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1052 ?? DL 0:00.00 [jsr_kkcm]
1337 ?? SL 0:00.09 [bcmTX]
1338 ?? SL 0:00.10 [bcmXGS3AsyncTX]
1339 ?? SL 0:03.22 [bcmLINK.0]
1344 ?? Is 0:00.00 /usr/sbin/cron
1496 ?? S 0:00.00 /sbin/watchdog -t-1
1497 ?? S 0:00.05 /usr/libexec/bslockd -mp -N
1498 ?? S 0:00.01 /usr/sbin/tnetd -N
1500 ?? R 0:05.17 /usr/sbin/chassisd -N
1501 ?? S 0:00.04 /usr/sbin/alarmd -N
1502 ?? I 0:00.39 /usr/sbin/craftd -N
1503 ?? S 0:00.08 /usr/sbin/mgd -N
1506 ?? I 0:00.05 /usr/sbin/inetd -N
1507 ?? I 0:00.00 /usr/sbin/tnp.sntpd -N
1508 ?? I 0:00.00 /usr/sbin/tnp.sntpc -N
1510 ?? S 0:00.01 /usr/sbin/smartd -N
1514 ?? I 0:00.07 /usr/sbin/jcsd -N
1515 ?? S 0:00.17 /usr/sbin/idpd -N
1516 ?? I 0:00.00 /usr/libexec/getty Pc ttyv0
2591 ?? DL 0:00.01 [peer proxy]
2592 ?? DL 0:00.01 [peer proxy]
2593 ?? DL 0:00.01 [peer proxy]
2597 ?? DL 0:00.00 [peer proxy]
3192 ?? S 0:00.01 /usr/sbin/irsd -N
3193 ?? S 0:00.05 /usr/sbin/snmpd -N
3194 ?? S 0:00.02 /sbin/dcd -N
3195 ?? S 0:00.01 /usr/sbin/pfed -N
3196 ?? S 0:00.01 /usr/sbin/mib2d -N
3197 ?? S 0:00.02 /usr/sbin/dfwd -N
3198 ?? S 0:00.13 /usr/sbin/ksyncd -N
3228 ?? Ss 0:00.00 mgd: (mgd) (root) (mgd)

```

```

3229 ?? R      0:00.00 /bin/ps -ax
1138 d0- S     0:00.00 /usr/sbin/usbd -N
1156 d0- S     0:00.42 /usr/sbin/eventd -N -r -s -A
1517 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0
...

```

### show system processes sfc (TX Matrix Plus Router)

```

user@host> show system processes sfc 0
sfc0-re0:

```

```

-----
PID  TT  STAT      TIME COMMAND
  0  ??  Wls      0:00.00 [swapper]
  1  ??  SLs      0:00.18 /packages/mnt/jbase/sbin/init --
  2  ??  DL       0:00.20 [g_event]
  3  ??  DL       0:00.39 [g_up]
  4  ??  DL       0:00.32 [g_down]
  5  ??  DL       0:00.00 [thread taskq]
  6  ??  DL       0:00.09 [kqueue taskq]
  7  ??  DL       0:00.01 [pagedaemon]
  8  ??  DL       0:00.00 [vmdaemon]
  9  ??  DL       0:06.63 [pagezero]
 10  ??  DL       0:00.00 [ktrace]
 11  ??  RL      312:09.00 [idle]
 12  ??  WL       0:11.07 [swi2: net]
 13  ??  WL       0:27.70 [swi7: clock sio]
 14  ??  WL       0:00.00 [swi6: vm]
 15  ??  DL       0:03.03 [yarrow]
 16  ??  WL       0:00.00 [swi9: +]
 17  ??  WL       0:00.00 [swi8: +]
 18  ??  WL       0:00.00 [swi5: cambio]
 19  ??  WL       0:00.00 [swi9: task queue]
 20  ??  WL       0:11.46 [irq16: uhci0 uhci*]
 21  ??  DL       0:00.00 [usb0]
 22  ??  DL       0:00.00 [usbtask]
 23  ??  WL       0:39.63 [irq17: uhci1 uhci*]
 24  ??  DL       0:00.00 [usb1]
 25  ??  WL       0:00.00 [irq18: uhci2 uhci*]
 26  ??  DL       0:00.84 [usb2]
 27  ??  DL       0:00.00 [usb3]
 28  ??  DL       0:00.00 [usb4]
 29  ??  DL       0:00.00 [usb5]
 30  ??  DL       0:00.73 [usb6]
 31  ??  DL       0:00.00 [usb7]
 32  ??  WL       0:00.00 [irq14: ata0]
 33  ??  WL       0:00.00 [irq15: ata1]
 34  ??  WL       0:00.00 [irq1: atkbd0]
 35  ??  WL       0:00.00 [swi0: sio]
 36  ??  WL       0:00.00 [irq11: isab0]
 37  ??  WL       0:00.00 [swi3: ip6opt ipopt]
 38  ??  WL       0:00.00 [swi4: ip6mismatch+]
 39  ??  WL       0:00.00 [swi1: ipfwd]
 40  ??  DL       0:00.02 [bufdaemon]
 41  ??  DL       0:00.02 [vnlr]
 42  ??  DL       0:00.39 [syncer]
 43  ??  DL       0:00.05 [softdepflush]
 44  ??  DL       0:00.00 [netdaemon]
 45  ??  DL       0:00.02 [vmuncachedaemon]
 46  ??  DL       0:00.00 [if_pic_listen]
 47  ??  DL       0:00.35 [vmkmemdaemon]
 48  ??  DL       0:00.00 [cb_poll]

```

```

49 ?? DL 0:00.06 [if_pfe_listen]
50 ?? DL 0:00.00 [scs_housekeeping]
51 ?? IL 0:00.00 [kern_dump_proc]
52 ?? IL 0:00.00 [nfsiod 0]
53 ?? IL 0:00.00 [nfsiod 1]
54 ?? IL 0:00.00 [nfsiod 2]
55 ?? IL 0:00.00 [nfsiod 3]
56 ?? DL 0:00.37 [schedcpu]
57 ?? DL 0:00.56 [md0]
79 ?? DL 0:02.58 [md1]
100 ?? DL 0:00.03 [md2]
118 ?? DL 0:00.01 [md3]
139 ?? DL 0:00.95 [md4]
160 ?? DL 0:00.12 [md5]
181 ?? DL 0:00.00 [md6]
217 ?? DL 0:00.02 [md7]
227 ?? DL 0:00.05 [md8]
1341 ?? SL 0:01.35 [bcmTX]
1342 ?? SL 0:01.69 [bcmXGS3AsyncTX]
1343 ?? SL 0:41.57 [bcmLINK.0]
1345 ?? SL 0:33.97 [bcmLINK.1]
1350 ?? Is 0:00.01 /usr/sbin/cron
1502 ?? S 0:00.01 /sbin/watchdog -t-1
1503 ?? S 0:00.86 /usr/libexec/bslockd -mp -N
1504 ?? I 0:00.01 /usr/sbin/tnetd -N
1507 ?? S 0:01.32 /usr/sbin/alarmd -N
1508 ?? S 0:14.54 /usr/sbin/craftd -N
1509 ?? S 0:01.20 /usr/sbin/mgd -N
1512 ?? S 0:00.05 /usr/sbin/inetd -N
1513 ?? S 0:00.10 /usr/sbin/tnp.sntpd -N
1517 ?? S 0:00.11 /usr/sbin/smartd -N
1525 ?? S 0:01.11 /usr/sbin/idpd -N
1526 ?? S 0:01.43 /usr/sbin/license-check -U -M -p 10 -i 10
1527 ?? I 0:00.01 /usr/libexec/getty Pc ttyv0
1616 ?? DL 0:00.30 [peer proxy]
1617 ?? DL 0:00.32 [peer proxy]
1618 ?? DL 0:00.34 [peer proxy]
1619 ?? DL 0:00.30 [peer proxy]
2391 ?? Is 0:00.01 telnetd
7331 ?? Ss 0:00.03 telnetd
9538 ?? DL 0:01.16 [jsr_kkcm]
9613 ?? DL 0:00.18 [peer proxy]
23781 ?? Ss 0:00.01 telnetd
23926 ?? Ss 0:00.03 mgd: (mgd) (regress)/dev/tty2 (mgd)
36867 ?? S 0:03.14 /usr/sbin/rpd -N
36874 ?? S 0:00.08 /usr/sbin/lmpd
36876 ?? S 0:00.17 /usr/sbin/lacpd -N
36877 ?? S 0:00.15 /usr/sbin/bfdd -N
36878 ?? S 0:05.05 /usr/sbin/ppmd -N
36907 ?? S 0:26.63 /usr/sbin/chassisd -N
37775 ?? S 0:00.01 /usr/sbin/bdbrepd -N
45727 ?? S 0:00.02 /usr/sbin/xntpd -j -N -g (ntpd)
45729 ?? S 0:00.40 /usr/sbin/l2ald -N
45730 ?? S< 0:00.13 /usr/sbin/apssd -N
45731 ?? SN 0:00.10 /usr/sbin/sampled -N
45732 ?? S 0:00.03 /usr/sbin/ilmid -N
45733 ?? S 0:00.09 /usr/sbin/rmopd -N
45734 ?? S 0:00.31 /usr/sbin/cosd
45735 ?? I 0:00.00 /usr/sbin/rtspd -N
45736 ?? S 0:00.06 /usr/sbin/fsad -N
45737 ?? S 0:00.05 /usr/sbin/rdd -N

```



```

45738 ?? S      0:00.10 /usr/sbin/pppd -N
45739 ?? S      0:00.05 /usr/sbin/dfcd -N
45740 ?? S      0:00.08 /usr/sbin/lfmd -N
45741 ?? S      0:00.01 /usr/sbin/mpiisoamd -N
45742 ?? I      0:00.01 /usr/sbin/sendd -N
45743 ?? S      0:00.08 /usr/sbin/appidd -N
45744 ?? S      0:00.05 /usr/sbin/mspd -N
45745 ?? S      0:00.27 /usr/sbin/jdiameterd -N
45746 ?? S      0:00.10 /usr/sbin/pfed -N
45747 ?? S      0:00.19 /usr/sbin/lpdfd -N
45748 ?? S      0:00.64 /sbin/dcd -N
45750 ?? S      0:00.46 /usr/sbin/mib2d -N
45751 ?? S      0:00.16 /usr/sbin/dfwd -N
45752 ?? S      0:00.15 /usr/sbin/irsd -N
45764 ?? S      0:20.60 /usr/sbin/snmpd -N
56481 ?? Ss     0:00.02 telnetd
56548 ?? Rs     0:00.19 mgd: (mgd) (regress)/dev/tty0 (mgd)
56577 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
56578 ?? R      0:00.00 /bin/ps -ax
1142 d0- S      0:00.01 /usr/sbin/usbd -N
1160 d0- S      0:29.71 /usr/sbin/eventd -N -r -s -A
6527 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0
56482 p0 Is     0:00.00 login [pam] (login)
56483 p0 S       0:00.01 -csh (csh)
56547 p0 S+     0:00.02 cli
2392 p1 Is     0:00.00 login [pam] (login)
2393 p1 I       0:00.00 -csh (csh)
2394 p1 I       0:00.00 su -
2395 p1 I+     0:00.01 -su (csh)
23782 p2 Is     0:00.00 login [pam] (login)
23881 p2 I       0:00.00 -csh (csh)
23925 p2 S+     0:00.03 cli
7332 p3 Is     0:00.00 login [pam] (login)
7333 p3 I       0:00.00 -csh (csh)
23780 p3 S+     0:00.02 telnet aj

```

### show system processes lcc wide (TX Matrix Plus Routing Matrix)

```

user@host> show system processes lcc 2 wide
lcc2-re0:

```

```

-----
PID  TT  STAT    TIME PROVIDER COMMAND
0   ??  WLS    0:00.00 (null)  [swapper]
1   ??  ILs    0:00.19          /packages/mnt/jbase/sbin/init --
2   ??  DL     0:00.02          [g_event]
3   ??  DL     0:00.19          [g_up]
4   ??  DL     0:00.13          [g_down]
5   ??  DL     0:00.00          [thread taskq]
6   ??  DL     0:00.00          [kqueue taskq]
7   ??  DL     0:00.00          [pagedaemon]
8   ??  DL     0:00.00          [vmdaemon]
9   ??  DL     0:01.77          [pagezero]
10  ??  DL     0:00.00          [ktrace]
11  ??  RL     20:33.81          [idle]
12  ??  WL     0:00.38          [swi2: net]
13  ??  WL     0:01.43          [swi7: clock sio]
14  ??  WL     0:00.00          [swi6: vm]
15  ??  DL     0:00.14          [yarrow]
16  ??  WL     0:00.00          [swi9: +]
17  ??  WL     0:00.00          [swi8: +]
18  ??  WL     0:00.00          [swi5: cambio]

```

19	??	WL	0:00.00	[swi9: task queue]
20	??	WL	0:03.18	[irq10: bcm0 uhci1*]
21	??	WL	0:00.03	[irq11: cb0 uhci0+*]
22	??	DL	0:00.00	[usb0]
23	??	DL	0:00.00	[usbtask]
24	??	DL	0:00.00	[usb1]
25	??	DL	0:00.06	[usb2]
26	??	DL	0:00.00	[usb3]
27	??	DL	0:00.00	[usb4]
28	??	DL	0:00.00	[usb5]
29	??	DL	0:00.05	[usb6]
30	??	DL	0:00.00	[usb7]
31	??	WL	0:00.00	[irq14: ata0]
32	??	WL	0:00.00	[irq15: ata1]
33	??	WL	0:00.00	[irq1: atkbd0]
34	??	WL	0:00.00	[swi0: sio]
35	??	WL	0:00.00	[swi3: ip6opt ipopt]
36	??	WL	0:00.00	[swi4: ip6mismatch+]
37	??	WL	0:00.00	[swi1: ipfwd]
38	??	DL	0:00.00	[bufdaemon]
39	??	DL	0:00.00	[vnlru]
40	??	DL	0:00.02	[syncer]
41	??	DL	0:00.01	[softdepflush]
42	??	DL	0:00.00	[netdaemon]
43	??	DL	0:00.00	[vmuncachedaemon]
44	??	DL	0:00.00	[if_pic_listen]
45	??	DL	0:00.03	[vmkmemdaemon]
46	??	DL	0:00.01	[cb_poll]
47	??	DL	0:00.00	[if_pfe_listen]
48	??	DL	0:00.00	[scs_housekeeping]
49	??	IL	0:00.00	[kern_dump_proc]
50	??	IL	0:00.00	[nfsiod 0]
51	??	IL	0:00.00	[nfsiod 1]
52	??	IL	0:00.00	[nfsiod 2]
53	??	IL	0:00.00	[nfsiod 3]
54	??	DL	0:00.02	[schedcpu]
55	??	DL	0:00.75	[md0]
77	??	DL	0:03.84	[md1]
98	??	DL	0:00.59	[md2]
116	??	DL	0:00.02	[md3]
137	??	DL	0:00.72	[md4]
158	??	DL	0:00.15	[md5]
179	??	DL	0:00.00	[md6]
215	??	DL	0:00.03	[md7]
225	??	DL	0:00.03	[md8]
1052	??	DL	0:00.00	[jsr_kkcm]
1337	??	SL	0:00.11	[bcmTX]
1338	??	SL	0:00.12	[bcmXGS3AsyncTX]
1339	??	SL	0:03.82	[bcmLINK.0]
1344	??	Is	0:00.00	/usr/sbin/cron
1496	??	I	0:00.00	/sbin/watchdog -t-1
1497	??	S	0:00.06	/usr/libexec/bslockd -mp -N
1498	??	I	0:00.01	/usr/sbin/tnetd -N
1500	??	S	0:09.93	/usr/sbin/chassisd -N
1501	??	S	0:00.05	/usr/sbin/alarmd -N
1502	??	I	0:00.39	/usr/sbin/craftd -N
1503	??	S	0:00.09	/usr/sbin/mgd -N
1506	??	I	0:00.05	/usr/sbin/inetd -N
1507	??	I	0:00.00	/usr/sbin/tnp.sntpd -N
1508	??	I	0:00.00	/usr/sbin/tnp.sntpc -N
1510	??	S	0:00.01	/usr/sbin/smartd -N

```

1514 ?? I      0:00.07      /usr/sbin/jcsd -N
1515 ?? S      0:00.17      /usr/sbin/idpd -N
1516 ?? I      0:00.00      /usr/libexec/getty Pc ttyv0
2591 ?? DL     0:00.01      [peer proxy]
2592 ?? DL     0:00.01      [peer proxy]
2593 ?? DL     0:00.01      [peer proxy]
2597 ?? DL     0:00.01      [peer proxy]
3192 ?? S      0:00.02      /usr/sbin/irsd -N
3193 ?? S      0:00.05      /usr/sbin/snmpd -N
3194 ?? S      0:00.04      /sbin/dcd -N
3195 ?? I      0:00.01      /usr/sbin/pfed -N
3196 ?? S      0:00.02      /usr/sbin/mib2d -N
3197 ?? I      0:00.03      /usr/sbin/dfwd -N
3198 ?? S      0:00.15      /usr/sbin/ksyncd -N
3559 ?? Ss     0:00.00      mgd: (mgd) (root) (mgd)
3560 ?? R      0:00.00      /bin/ps -ax -Jpww
1138 d0- S      0:00.00      /usr/sbin/usbd -N
1156 d0- S      0:00.50      /usr/sbin/eventd -N -r -s -A
1517 d0 Is+    0:00.00      /usr/libexec/getty std.9600 ttyd0

```

#### show system processes (QFX Series)

```
user@switch> show system processes
```

```

PID  TT  STAT      TIME COMMAND
  0  ??  Wls    -2341043:-31.01 [swapper]
  1  ??  SLs     0:01.34 /packages/mnt/jbase/sbin/init --
  2  ??  DL      2:48.31 [g_event]
  3  ??  DL      1:47.44 [g_up]
  4  ??  DL      1:37.82 [g_down]
  5  ??  DL      0:00.00 [kdm_tcp_poller]
  6  ??  DL      0:00.00 [thread taskq]
  7  ??  DL      0:04.86 [kqueue taskq]
  9  ??  DL      0:03.94 [pagedaemon]
 10  ??  DL      0:00.00 [ktrace]
 11  ??  RL      0:00.00 [idle: cpu31]
 12  ??  RL      0:00.00 [idle: cpu30]
 13  ??  RL      0:00.00 [idle: cpu29]
 14  ??  RL      0:00.00 [idle: cpu28]
 15  ??  RL      0:00.00 [idle: cpu27]
 16  ??  RL      0:00.00 [idle: cpu26]
 17  ??  RL      0:00.00 [idle: cpu25]
 18  ??  RL      0:00.00 [idle: cpu24]
 19  ??  RL      0:00.00 [idle: cpu23]
 20  ??  RL      0:00.00 [idle: cpu22]
 21  ??  RL      0:00.00 [idle: cpu21]
 22  ??  RL      0:00.00 [idle: cpu20]
 23  ??  RL      0:00.00 [idle: cpu19]
 24  ??  RL      0:00.00 [idle: cpu18]
 25  ??  RL      0:00.00 [idle: cpu17]
 26  ??  RL      0:00.00 [idle: cpu16]
 27  ??  RL      0:00.00 [idle: cpu15]
 28  ??  RL      0:00.00 [idle: cpu14]
 29  ??  RL      0:00.00 [idle: cpu13]
 30  ??  RL      0:00.00 [idle: cpu12]
 31  ??  RL      0:00.00 [idle: cpu11]
 32  ??  RL      0:00.00 [idle: cpu10]
 33  ??  RL      0:00.00 [idle: cpu9]
 34  ??  RL    18184:07.25 [idle: cpu8]
 35  ??  RL      0:00.00 [idle: cpu7]
 36  ??  RL    17862:11.31 [idle: cpu6]
 37  ??  RL    19343:45.16 [idle: cpu5]

```

```

38 ?? RL 5192:38.30 [idle: cpu4]
39 ?? RL 0:00.00 [idle: cpu3]
40 ?? RL 19278:02.24 [idle: cpu2]
41 ?? RL 19291:00.72 [idle: cpu1]
42 ?? RL 18910:31.21 [idle: cpu0]
43 ?? WL 19:03.74 [swi2: net]
44 ?? WL 261:43.82 [swi7: clock sio]
45 ?? WL 0:00.00 [swi6: vm]
46 ?? DL 2:18.57 [yarrow]
47 ?? WL 0:00.00 [swi9: +]
48 ?? WL 0:00.00 [swi8: +]
49 ?? WL 0:12.36 [swi5: cambio]
50 ?? WL 0:00.00 [swi9: task queue]
51 ?? WL 0:00.00 [swi0: sio]
52 ?? WL 0:32.40 [irq39: ehci0]
53 ?? DL 0:00.21 [usb0]
54 ?? DL 0:00.00 [usbtask]
55 ?? WL 0:00.00 [irq22: xlr_lbus0]
56 ?? WL 0:00.00 [irq38: xlr_lbus0]
57 ?? WL 0:00.00 [swi3: ip6opt ipopt]
58 ?? WL 0:00.00 [swi4: ip6mismatch+]
59 ?? WL 0:00.00 [swi1: ipfwd]
60 ?? DL 0:18.65 [pagezero]
61 ?? DL 0:18.59 [bufdaemon]
62 ?? DL 1:10.44 [vnlru_mem]
63 ?? DL 1:51.66 [syncer]
64 ?? DL 0:20.22 [vnlru]
65 ?? DL 0:40.48 [softdepflush]
66 ?? DL 0:00.00 [netdaemon]
67 ?? DL 20:47.67 [vmkmemdaemon]
68 ?? DL 0:00.00 [if_pfe_listen]
69 ?? SL 0:02.80 [kdm_checkkcore]
70 ?? SL 0:03.34 [kdm_savekcore]
71 ?? SL 0:04.31 [kdm_livekcore]
72 ?? SL 0:06.14 [kdm_logger]
73 ?? SL 0:04.31 [kdm_kdb]
74 ?? SL 0:00.02 [devrt_kernel_thread]
75 ?? DL 0:21.54 [vmuncachedaemon]
76 ?? DL 0:00.00 [if_pic_listen0]
77 ?? SL 0:00.00 [nfsiod 0]
78 ?? SL 0:00.00 [nfsiod 1]
79 ?? SL 0:00.00 [nfsiod 2]
80 ?? SL 0:00.00 [nfsiod 3]
81 ?? WL 5:59.98 [irq13: +]
82 ?? RL 105:06.81 [pkt_sender: cpu0]
83 ?? DL 0:03.62 [md0]
95 ?? DL 0:37.04 [md1]
115 ?? DL 0:06.01 [md2]
135 ?? DL 0:00.75 [md3]
155 ?? DL 0:21.17 [md4]
175 ?? DL 0:01.90 [md5]
195 ?? DL 0:06.26 [md6]
231 ?? DL 0:00.01 [md7]
755 ?? Ss 0:04.17 /usr/sbin/cron
847 ?? S 0:00.10 /usr/sbin/tinetd -N
849 ?? S 0:06.82 /usr/sbin/mgd -N
850 ?? S 0:00.32 /usr/sbin/inetd -N
852 ?? S 1:05.34 /usr/sbin/dhcpd -N
853 ?? S 0:00.18 /usr/sbin/inetd -p /var/run/inetd_4.pid -N -JU __juni
855 ?? L 1181:02.21 /usr/sbin/dc-pfe -N (pafxpc)
857 ?? S 17:55.86 /usr/sbin/vccpd -N

```

```

896 ?? S      93:43.45 /usr/sbin/chassism -N
953 ?? S      0:02.89 /sbin/watchdog -t-1
954 ?? S      3:34.00 /sbin/dcd -N
955 ?? S     10:30.13 /usr/sbin/chassisd -N
956 ?? DL     0:00.21 [peer proxy]
957 ?? S      4:07.43 /usr/sbin/alarmd -N
958 ?? S      0:31.69 /usr/sbin/craftd -N
959 ?? S      0:55.16 /usr/sbin/mib2d -N
960 ?? S      3:40.64 /usr/sbin/rpd -N
961 ?? S      0:00.03 /usr/sbin/tnp.sntpd -N
962 ?? S      0:51.94 /usr/sbin/pfed -N
963 ?? S      0:47.31 /usr/sbin/rmopd -N
964 ?? S      0:33.65 /usr/sbin/cosd
965 ?? S      1:48.41 /usr/sbin/ppmd -N
966 ?? S      0:07.18 /usr/sbin/dfwd -N
967 ?? S      1:02.56 /usr/sbin/bfdd -N
968 ?? S      0:00.63 /usr/sbin/rdd -N
969 ?? S      0:40.61 /usr/sbin/dfcd -N
971 ?? S      0:07.81 /usr/sbin/bdbrepd -N
972 ?? S      0:00.28 /usr/sbin/sendd -N
973 ?? S      1:37.69 /usr/sbin/xntpd -j -N -g -JU __juniper_private4__ (nt
974 ?? S      5:56.28 /usr/sbin/snmpd -N -JU __juniper_private4__
975 ?? S     16:46.82 /usr/sbin/jdiameterd -N
976 ?? S      2:34.13 /usr/sbin/eswd -N
977 ?? S      1:03.05 /usr/sbin/sflowd -N
978 ?? S      0:22.30 /usr/sbin/fcd -N
979 ?? S      1:07.01 /usr/sbin/vccpdf -N
982 ?? S      0:25.25 /usr/sbin/mcsnoopd -N
983 ?? S      3:45.68 /usr/sbin/rpdf -N
1043 ?? S      0:37.87 /usr/sbin/lacpd -N
1048 ?? DL     0:01.29 [peer proxy]
1111 ?? WL     0:00.00 [swi2: FMNITHRD+]
1112 ?? DL     0:00.03 [peer proxy]
12816 ?? S     15:35.32 /usr/sbin/sfid -N
30893 ?? Ss    0:00.65 sshd: tlewis@tty0 (sshd)
30897 ?? Ss    0:00.15 mgd: (mgd) (tlewis)/dev/tty0 (mgd)
30905 ?? Ss    0:00.64 sshd: tlewis@tty1 (sshd)
30909 ?? Ss    0:00.15 mgd: (mgd) (tlewis)/dev/tty1 (mgd)
30910 ?? Ss    0:01.26 sshd: tcheng@tty2 (sshd)
30914 ?? Ss    0:00.80 mgd: (mgd) (tcheng)/dev/tty2 (mgd)
30937 ?? R      0:00.03 /bin/ps -ax
661  d0- S      0:21.24 /usr/sbin/eventd -N -r -s -A
860  d0 Ss+    0:00.07 /usr/libexec/getty std.9600 ttyd0
30896 p0 Ss+    0:00.55 -cli (cli)
30908 p1 Ss+    0:00.50 -cli (cli)
30913 p2 Ss+    0:00.85 -cli (cli)

```



## PART 8

# System Services

- [Overview on page 1353](#)
- [Configuration on page 1393](#)
- [Administration on page 1691](#)





## CHAPTER 24

# Overview

- [Software Overview on page 1353](#)
- [DHCP Local Server on page 1355](#)
- [DHCP Relay Agent on page 1378](#)
- [Public Key Cryptography Overview on page 1388](#)
- [Self-Signed Certificates Overview on page 1390](#)

## Software Overview

---

- [Understanding Software Infrastructure and Processes on page 1353](#)

### Understanding Software Infrastructure and Processes

Each switch runs the Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches on its general-purpose processors. Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the chassis.

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

With the J-Web interface and the command-line interface (CLI) to the Junos OS, you configure switching features and routing protocols and set the properties of network interfaces on your switch. After activating a software configuration, use either the J-Web or CLI user interface to monitor the switch, manage operations, and diagnose protocol and network connectivity problems.

- [Routing Engine and Packet Forwarding Engine on page 1353](#)
- [Junos OS Processes on page 1354](#)

### Routing Engine and Packet Forwarding Engine

---

A switch has two primary software processing components:

- **Packet Forwarding Engine**—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.

- Routing Engine—Provides three main functions:
  - Creates the packet forwarding switch fabric for the switch, providing route lookup, filtering, and switching on incoming data packets, then directing outbound packets to the appropriate interface for transmission to the network
  - Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
  - Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

### Junos OS Processes

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

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

[Table 6 on page 30](#) describes the primary Junos OS processes.

**Table 130: Junos OS Processes**

Process	Name	Description
Chassis process	chassisd	<p>Detects hardware on the system that is used to configure network interfaces.</p> <p>Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered.</p> <p>Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.</p>
Ethernet switching process	eswd	<p>Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree protocol and access port security. The process is also responsible for managing Ethernet switching interfaces, VLANs, and VLAN interfaces.</p> <p>Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.</p>
Forwarding process	pfem	<p>Defines how routing protocols operate on the switch. The overall performance of the switch is largely determined by the effectiveness of the forwarding process.</p>
Interface process	dcd	<p>Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.</p>

Table 130: Junos OS Processes (*continued*)

Process	Name	Description
Management process	mgd	<p>Provides communication between the other processes and an interface to the configuration database.</p> <p>Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured.</p> <p>Interacts with the other processes when commands are issued through one of the user interfaces on the switch.</p> <p>If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.</p>
Routing protocol process	rpd	Defines how routing protocols such as RIP, OSPF, and BGP operate on the device, including selecting routes and maintaining forwarding tables.

**Related Documentation**

- [For more information about processes, see \*Junos OS Network Operations Guide\*](#)
- [For more information about basic system parameters, supported protocols, and software processes, see \*Junos OS System Basics Configuration Guide\*](#)

## DHCP Local Server

- [Extended DHCP Local Server Overview on page 1356](#)
- [DHCPv6 Local Server Overview on page 1361](#)
- [DHCP Local Server Handling of Client Information Request Messages on page 1363](#)
- [Configuring Group-Specific DHCP Local Server Options on page 1364](#)
- [Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients on page 1364](#)
- [DHCP Snooping Support on page 1367](#)
- [DHCP Auto Logout Overview on page 1368](#)
- [Address-Assignment Pools Overview on page 1370](#)
- [Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA\\_NA Option on page 1371](#)
- [Multiple Address Assignment for DHCPv6 Clients on page 1371](#)
- [Centrally Configured Opaque DHCP Options on page 1373](#)
- [Port Number Requirements for DHCP Firewall Filters on page 1377](#)

## Extended DHCP Local Server Overview

Junos OS includes an extended DHCP local server that enhances traditional DHCP server operation by providing additional address assignment and client configuration functionality and flexibility in a subscriber-aware environment. The extended DHCP local server enables service providers to take advantage of external address-assignment pools and integrated RADIUS-based configuration capabilities in addition to the continued support of traditional local address pools. The address-assignment pools are considered external because they are external to the DHCP local server. The pools are managed independently of the DHCP local server, and can be shared by different client applications, such as DHCP or PPPoE access. [Table 131 on page 1357](#) provides a comparison of the extended DHCP local server and a traditional DHCP local server.

The extended DHCP local server provides an IP address and other configuration information in response to a client request. The server supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide DHCP client authentication. You can configure the dynamic profile and authentication support on a global basis or for a specific group of interfaces.

**Table 131: Comparing the Extended DHCP Local Server to the Traditional DHCP Local Server**

Feature	Extended DHCP Local Server	Traditional DHCP Local Server
Local address pools	X	X
External, centrally-managed address pools	X	—
Local configuration	X	X
External configuration using information from address-assignment pools or RADIUS servers	X	—
Dynamic-profile attachment	X	—
RADIUS-based subscriber authentication, and configuration using RADIUS attributes and Juniper Networks VSAs	X	—
IPv6 client support	X	—
Default minimum client configuration	X	X

You can also configure the extended DHCP local server to support IPv6 clients. Both DHCP local server and DHCPv6 local server support the specific address request feature, which enables you to assign a particular address to a client.



**NOTE:** If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

This overview covers:

- [Interaction Among the DHCP Client, Extended DHCP Local Server, and Address-Assignment Pools on page 1358](#)
- [Providing DHCP Client Configuration Information on page 1358](#)
- [Minimal Configuration for Clients on page 1360](#)
- [DHCP Local Server and Address-Assignment Pools on page 1360](#)

### [Interaction Among the DHCP Client, Extended DHCP Local Server, and Address-Assignment Pools](#)

---

The pattern of interaction between the DHCP local server, the DHCP client, and address-assignment pools is the same regardless of whether you are using a router or a switch. However, there are some differences in the details of usage.

- On routers—In a typical carrier edge network configuration, the DHCP client is on the subscriber's computer or customer premises equipment (CPE), and the DHCP local server is configured on the router.
- On switches—In a typical network configuration, the DHCP client is on an access device, such as a personal computer, and the DHCP local server is configured on the switch.

The following steps provide a high-level description of the interaction among the DHCP local server, DHCP client, and address-assignment pools:

1. The DHCP client sends a discover packet to one or more DHCP local servers in the network to obtain configuration parameters and an IP address for the subscriber (or DHCP client).
2. Each DHCP local server that receives the discover packet then searches its address-assignment pool for the client address and configuration options. Each local server creates an entry in its internal client table to keep track of the client state, then sends a DHCP offer packet to the client.
3. On receipt of the offer packet, the DHCP client selects the DHCP local server from which to obtain configuration information and sends a request packet indicating the DHCP local server selected to grant the address and configuration information.
4. The selected DHCP local server sends an acknowledgement packet to the client that contains the client address lease and configuration parameters. The server also installs the host route and ARP entry, and then monitors the lease state.

### [Providing DHCP Client Configuration Information](#)

---

When the extended DHCP application receives a response from an external authentication server, the response might include information in addition to the IP address and subnet

mask. The extended DHCP application uses the information from the authentication grant for the response the DHCP application sends to the DHCP client. The DHCP application can either send the information in its original form or the application might merge the information with local configuration specifications. For example, if the authentication grant includes an address pool name and a local configuration specifies DHCP attributes for that pool (such as, DNS server address), the extended DHCP application merges the authentication results and the attributes in the reply that the server sends to the client.

A local configuration is optional — a client can be fully configured by the external authentication service. However, if the external authentication service does not provide client configuration, you might need to configure the local address-assignment pool to provide the configuration information, such as DNS server, for the client. When a local configuration specifies options, the extended DHCP application adds the local configuration options to the offer PDU the server sends to the client. If the two sets of options overlap, the options in the authentication response from the external service take precedence.

When you use RADIUS to provide the authentication, the additional information might be in the form of RADIUS attributes and Juniper Networks VSAs. [Table 132 on page 1359](#) lists the information that RADIUS might include in the authentication grant. See *RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework* for a complete list of RADIUS attributes and Juniper Networks VSAs that the extended DHCP applications supports for subscriber access management or DHCP management.

**Table 132: Information in Authentication Grant**

Attribute Number	Attribute Name	Description
RADIUS attribute 8	Framed-IP-Address	Client IP address
RADIUS attribute 9	Framed-IP-Netmask	Subnet mask for client IP address (DHCP option 1)
Juniper Networks VSA 26-4	Primary-DNS	Primary domain server (DHCP option 6)
Juniper Networks VSA 26-5	Secondary-DNS	Secondary domain server (DHCP option 6)
Juniper Networks VSA 26-6	Primary-WINS	Primary WINS server (DHCP option 44)
Juniper Networks VSA 26-7	Secondary-WINS	Secondary WINS server (DHCP option 44)
RADIUS attribute 27	Session-Timeout	Lease time
RADIUS attribute 88	Framed-Pool	Address assignment pool name
Juniper Networks VSA 26-109	DHCP-Guided-Relay-Server	DHCP relay server

### Minimal Configuration for Clients

---

The extended DHCP local server provides a minimal configuration to the DHCP client if the client does not have DHCP option 55 configured. The server provides the subnet mask of the address-assignment pool that is selected for the client. In addition to the subnet mask, the server provides the following values to the client if the information is configured in the selected address-assignment pool:

- **router**—A router located on the client's subnet. This statement is the equivalent of DHCP option 3.
- **domain name**—The name of the domain in which the client searches for a DHCP server host. This is the default domain name that is appended to hostnames that are not fully qualified. This is equivalent to DHCP option 15.
- **domain name server**—A Domain Name System (DNS) name server that is available to the client to resolve hostname-to-client mappings. This is equivalent to DHCP option 6.

### DHCP Local Server and Address-Assignment Pools

---

In the traditional DHCP server operation, the client address pool and client configuration information reside on the DHCP server. With the extended DHCP local server, the client address and configuration information reside in external address-assignment pools (external to the DHCP local server). The external address-assignment pools are managed by the **authd** process, independently of the DHCP local server, and can be shared by different client applications.

The extended DHCP local server also supports advanced pool matching and the use of named address ranges. You can also configure the local server to use DHCP option 82 information in the client PDU to determine which named address range to use for a particular client. The client configuration information, which is configured in the address-assignment pool, includes user-defined options, such as boot server, grace period, and lease time.

Configuring the DHCP environment that includes the extended DHCP local server requires two independent configuration operations, which you can complete in any order. In one operation, you configure the extended DHCP local server on the router and specify how the DHCP local server determines which address-assignment pool to use. In the other operation, you configure the address-assignment pools used by the DHCP local server. The address-assignment pools contain the IP addresses, named address ranges, and configuration information for DHCP clients.



**NOTE:** The extended DHCP local server and the address-assignment pools used by the server must be configured in the same logical system and routing instance.

---

#### Related Documentation

- [Address-Assignment Pools Overview on page 1370](#)



- [Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 1442](#)
- [Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview](#)
- [Using External AAA Authentication Services with DHCP on page 1420](#)
- [Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA\\_NA Option on page 1371](#)
- [Graceful Routing Engine Switchover for DHCP](#)
- [High Availability Using Unified ISSU in the PPP Access Network](#)
- [Tracing Extended DHCP Operations](#)
- [Verifying and Managing DHCP Local Server Configuration on page 1696](#)
- [Example: Minimum Extended DHCP Local Server Configuration on page 1393](#)
- [Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 1394](#)
- [Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine](#)

## DHCPv6 Local Server Overview

The DHCPv6 local server enhances the extended DHCP local server by providing support for IPv6. When a DHCPv6 client logs in, the DHCPv6 local server can optionally use the AAA service framework to interact with the RADIUS server. The RADIUS server, which is configured independently of DHCP, authenticates the client and supplies the IPv6 prefix and client configuration parameters.

You can configure DHCPv6 local server to communicate the following attributes to the AAA service framework and RADIUS at login time:

- Client username
- Client password



**NOTE:** The client username, which uniquely identifies a subscriber or a DHCP client, must be present in the configuration in order for DHCPv6 local server to use RADIUS authentication.

Based on the attributes that the DHCPv6 local server provides, RADIUS returns the information listed in [Table 133 on page 1361](#) to configure the client:

**Table 133: RADIUS Attributes and VSAs for DHCPv6 Local Server**

Attribute Number	Attribute Name	Description
27	Session-Timeout	Lease time, in seconds. If not supplied, the lease does not expire

**Table 133: RADIUS Attributes and VSAs for DHCPv6 Local Server (*continued*)**

Attribute Number	Attribute Name	Description
123	Delegated-IPv6-Prefix	Prefix that is delegated to the client
26-143	Max-Clients-Per-Interface	Maximum number of clients allowed per interface

The DHCPv6 local server is compatible with the extended DHCP local server and the extended DHCP relay agent, and can be enabled on the same interface as either the extended DHCP local server or DHCP relay agent.

The DHCPv6 local server provides many of the same features as the extended DHCP local server, including:

- Configuration for a specific interface or for a group of interfaces
- Site-specific usernames and passwords
- Numbered Ethernet interfaces
- Statically configured CoS and filters
- AAA directed login
- Use of the IA\_NA option to assign a specific address to a client

To configure the extended DHCPv6 local server on the router (or switch), you include the **dhcpv6** statement at the **[edit system services dhcp-local-server]** hierarchy level. See the *[edit system services dhcp-local-server] Hierarchy Level* for the complete DHCP local server syntax, including the DHCPv6 syntax.

You can also include the **dhcpv6** statement at the following hierarchy levels:

- **[edit logical-systems *logical-system-name* system services dhcp-local-server]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server]**
- **[edit routing-instances *routing-instance-name* system services dhcp-local-server]**

#### Related Documentation

- [Extended DHCP Local Server Overview on page 1356](#)
- [Using External AAA Authentication Services with DHCP on page 1420](#)
- [Grouping Interfaces with Common DHCP Configurations on page 1421](#)
- [Configuring Group-Specific DHCP Local Server Options on page 1364](#)
- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)
- [Configuring Passwords for Usernames on page 1439](#)
- [Creating Unique Usernames for DHCP Clients on page 1439](#)

- [Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA\\_NA Option on page 1371](#)
- [Verifying and Managing DHCPv6 Local Server Configuration on page 1696](#)
- *Example: Extended DHCPv6 Local Server Configuration*

## DHCP Local Server Handling of Client Information Request Messages

DHCP clients that already have externally provided addresses may solicit further configuration information from a DHCP server by sending a DHCP information request that indicates what information is desired. By default, DHCP local server and DHCPv6 local server ignore any DHCP information requests that they receive. You can override this default behavior to enable processing of these messages.

If you enable processing of information requests, DHCP local server responds to the client with a DHCP acknowledgment message that includes the requested information—if it is available. DHCPv6 local server responds in the same manner but uses a DHCP reply message. No subscriber management or DHCP-management is applied as a result of the DHCP information request message.

By default, DHCP relay and DHCP relay proxy automatically forward DHCP information request messages without modification if the messages are received on an interface configured for a DHCP server group. DHCP relay and relay proxy drop information request messages received on any other interfaces. You cannot disable this default DHCP relay and relay proxy behavior.

The information requested by these clients has typically been configured with the **dhcp-attributes** statement for an address pool defined by the **address-assignment pool pool-name** statement at the **[edit access]** hierarchy level.

When you enable processing of DHCP information requests, you can optionally specify the name of the pool from which the local server retrieves the requested configuration information for the client. If you do not specify a local pool, then the local server requests that AAA selects and returns only the name of the relevant pool.



**NOTE:** PPP interfaces are not supported on EX Series switches.

When DHCPv6 is configured over PPP interfaces, the PPP RADIUS authentication data can be used to select the pool from which the response information is taken. Additionally other RADIUS attributes can also be inserted into the DHCPv6 reply message. If an overlap exists between RADIUS attributes and local pool attributes, the RADIUS values are used instead of the local configuration data. If no RADIUS information is received from the underlying PPP interface, then the behavior is the same as described previously for non-PPP interfaces.

### Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)
- [Enabling Processing of Client Information Requests on page 1427](#)

## Configuring Group-Specific DHCP Local Server Options

You can include the following statements at the **[edit system services dhcp-local-server group group-name]** hierarchy level to set group-specific DHCP local server configuration options. Statements configured at the **[edit system services dhcp-local-server group group-name]** hierarchy level apply only to the named group of interfaces, and override any global DHCP local server settings configured with the same statements at the **[edit system services dhcp-local-server]** hierarchy level.

DHCPv6 local server supports the same set of statements with the exception of the **dynamic-profile** statement.

- **authentication**—Configure the parameters the router sends to the external AAA server.
- **dynamic-profile**—Specify the dynamic profile that is attached to a group of interfaces.
- **interface**—Specify one or more interfaces, or a range of interfaces, that are within the specified group.
- **overrides**—Override the default configuration settings for the extended DHCP local server. For information, see “[Overriding Default DHCP Local Server Configuration Settings](#)” on page 1423.

### Related Documentation

- [Grouping Interfaces with Common DHCP Configurations on page 1421](#)

## Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients

Dynamic reconfiguration of clients enables the extended DHCP local server to initiate a client update without waiting for the client to initiate a request.

### Default Client/Server Interaction

---

Typically the DHCP client initiates all of the basic DHCP client/server interactions. The DHCP server sends information to a client only in response to a request from that client. This behavior does not enable a client to be quickly updated with its network address and configuration in the event of server changes:



**NOTE:** Technically, the DHCP client/server interactions are the same on routers and switches. However, the primary usage of this technology on the routers is for subscriber management. The switches are not used for subscriber management. Therefore, this topic provides two sample scenarios. The actions are the same, but the implementation details are different.

- On routers—Suppose a service provider restructures its addressing scheme or changes the server IP addresses that it provided to clients. Without dynamic reconfiguration, the service provider typically clears the DHCP server binding table, but cannot inform the DHCP clients that their bindings have been cleared. Consequently, the DHCP client operates as though its IP address is still valid, but it is now unable to communicate over the access network, resulting in an outage. The DHCP local server needs to wait

for the client to send a message to renew its lease or rebind to the server. In response, the server sends a NAK message to the client to force it to begin the DHCP connection process again. Alternatively, the provider can wait for customers to make a service call about the network failures and then instruct them to power cycle their customer premises equipment to reinitiate the connection. Neither of these actions is timely or convenient for customers.

- On switches—Suppose you restructure the addressing scheme or change the server IP addresses that the DHCP server provides to clients. Without dynamic reconfiguration, the network typically clears the DHCP server binding table, but cannot inform the DHCP clients that their bindings have been cleared. Consequently, the DHCP client operates as though its IP address is still valid, but it is now unable to communicate over the access network, resulting in an outage. The DHCP local server needs to wait for the client to send a message to renew its lease or rebind to the server. In response, the server sends a NAK message to the client to force it to begin the DHCP connection process again. Alternatively, you can wait for users to notify you of the network failures and then instruct them to power cycle their equipment to reinitiate the connection. Neither of these actions is timely or convenient for users.

#### Dynamic Client/Server Interaction for DHCPv4

Dynamic reconfiguration for DHCPv4 is available through a partial implementation of RFC 3203, *DHCP Reconfigure Extension* for DHCPv4. It enables the DHCPv4 local server to send a message to the client to force reconfiguration.

The server sends a `forcerenew` message to a DHCPv4 client, initiating a message exchange. In response, DHCPv4 clients that support the `forcerenew` message then send a lease renewal message to the server. The server rejects the lease renewal request and sends a NAK to the client, causing the client to reinitiate the DHCP connection. A successful reconnection results in the reconfiguration of the DHCP client. Only the exchange of `forcerenew`, `renew`, and NAK messages is supported from RFC 3202. DHCP relay and DHCP relay proxy do not participate in the client reconfiguration or react to `forcerenew` messages other than to forward them to the client.

When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfiguring state and the local server sends a `forcerenew` message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber services or DHCP-managed services, such as forwarding and statistics, continue to work. Client statistics are not maintained in the interval between a successful reconfiguration and the subsequent client binding. When the server responds to the client renewal request with a NAK, the client entry is removed from the binding table and final statistics are reported. New statistics are collected when the client sends a discover message to establish a new session.

#### Dynamic Client/Server Interaction for DHCPv6

Dynamic reconfiguration for DHCPv6 is available through a partial implementation of RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*. It enables the DHCPv6 local server to send a message to the client to force reconfiguration.

DHCPv6 servers send reconfigure messages to DHCPv6 clients, initiating a message exchange. In response, DHCPv6 clients that support the reconfigure message transition to the renewing state and send a renew message to the server. The server returns a reply message with a lifetime of zero (0). The client transitions to the init state and sends a solicit message. The server sends an advertise message to indicate that it is available for service. The client sends a request for configuration parameters, which the server then includes in its reply. DHCP relay and DHCP relay proxy do not participate in the client reconfiguration or react to reconfigure messages other than to forward them to the client.

When a DHCPv6 server is triggered to initiate reconfiguration on a bound DHCPv6 client, the client transitions to the reconfigure state. All subscriber services, such as forwarding and statistics, continue to work. The server then sends the reconfigure message to the client. If the DHCPv6 client is already in the reconfigure state, the DHCPv6 server ignores the reconfiguration trigger. For clients in any state other than bound or reconfigure, the server clears the binding state of the client, as if the **clear dhcpv6 server binding** command had been issued.

### Manually Forcing the Local Server to Initiate the Reconfiguration Process

You can force the local server to initiate the reconfiguration process for clients by issuing the **request dhcp server reconfigure** command for DHCPv4 clients, and the **request dhcpv6 server reconfigure** command for DHCPv6 clients. Command options determine whether reconfiguration is then attempted for all clients or specified clients.

### Action Taken for Events That Occur During a Reconfiguration

Events that take place while a reconfiguration is in process take precedence over the reconfiguration. [Table 134 on page 1366](#) lists the actions taken in response to several different events.

**Table 134: Action Taken for Events That Occur During a Reconfiguration**

Event	Action
Server receives a discover (DHCPv4) or solicit (DHCPv6) message from the client.	Server drops packet and deletes client.
Server receives a request, renew, rebind, or init-reboot message from the client.	DHCPv4—Server sends NAK message and deletes client.  DHCPv6—Server drops packet and deletes client. Server replies to renew message with lease time of zero (0).
Server receives a release or decline message from the client.	Server deletes client.
The client lease times out.	Server deletes client.
The <b>clear dhcp server binding</b> command is issued.	Server deletes client.
The <b>request dhcp server reconfigure</b> (DHCPv4) or <b>request dhcpv6 server reconfigure</b> (DHCPv6) command is issued.	Command is ignored.

**Table 134: Action Taken for Events That Occur During a Reconfiguration (*continued*)**

Event	Action
GRES or DHCP restart occurs.	Reconfiguration process is halted.

**Related Documentation**

- [Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429](#)

## DHCP Snooping Support

DHCP snooping provides DHCP security on the router or switch by filtering incoming messages. When DHCP snooping is enabled, the router differentiates between trusted and untrusted interfaces, and forwards messages from trusted sources while rejecting the untrusted messages.

In Junos OS, DHCP snooping is enabled in a routing instance when you configure either the **dhcp-relay** statement at the **[edit forwarding-options]** hierarchy level, or the **dhcp-local-server** statement at the **[edit system services]** hierarchy level in that routing instance. The router discards snooped packets by default. To enable normal processing of snooped packets, you must explicitly configure the **allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level.

You can configure DHCP snooping support for the following:

- DHCPv4 relay agent—Override the router's (or switch's) default snooping configuration and specify that DHCP snooping is enabled or disabled globally, for a named group of interfaces, or for a specific interface within a named group.

In a separate procedure, you can set a global configuration to specify whether the DHCPv4 relay agent forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces. The router also uses the global DHCP relay agent snooping configuration to determine whether to forward or drop snooped BOOTREPLY packets.

- DHCPv6 relay agent—As you can with snooping support for the DHCPv4 relay agent, you can override the default DHCPv6 relay agent snooping configuration on the router to explicitly enable or disable snooping support globally, for a named group of interfaces, or for a specific interface with a named group of interfaces.

In multi-relay topologies where more than one DHCPv6 relay agent is between the DHCPv6 client and the DHCPv6 server, snooping enables intervening DHCPv6 relay agents between the client and the server to correctly receive and process the unicast traffic from the client and forward it to the server. The DHCPv6 relay agent snoops incoming unicast DHCPv6 packets by setting up a filter with UDP port 547 (the DHCPv6 UDP server port) on a per-forwarding table basis. The DHCPv6 relay agent then processes the packets intercepted by the filter and forwards the packets to the DHCPv6 server.

Unlike the DHCPv4 relay agent, the DHCPv6 relay agent does not support global configuration of forwarding support for DHCPv6 snooped packets.

- DHCP local server—Configure whether DHCP local server forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces.

**Related Documentation**

- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 1437](#)
- [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 1455](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 1460](#)
- [Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 1403](#)
- [Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent](#)

## DHCP Auto Logout Overview

This topic provides an introduction to the DHCP auto logout feature and includes the following sections:

- [Auto Logout Overview on page 1368](#)
- [How DHCP Identifies and Releases Clients on page 1368](#)
- [Option 60 and Option 82 Requirements on page 1369](#)

### Auto Logout Overview

---

Auto logout is supported for DHCP local server and DHCP relay agent. It improves the efficiency of DHCP IP address assignment by allowing IP addresses to be immediately released and returned to the address pool when DHCP clients are no longer using the addresses. DHCP can then assign the addresses to other clients. Without auto logout, an IP address is blocked for the entire lease period, and DHCP must wait until the address lease time expires before reusing the address.

Auto logout is particularly useful when DHCP uses long lease times for IP address assignments and to help avoid allocating duplicate IP addresses for a single client.

For example, you might have an environment that includes set-top boxes (STB) that are often upgraded or replaced. Each time a STB is changed, the new STB repeats the DHCP discover process to obtain client configuration information and an IP address. DHCP views the new STB as a completely new client and assigns a new IP address—the previous IP address assigned to the client (the old STB) remains blocked and unavailable until the lease expires. If auto logout is configured in this situation, DHCP recognizes that the new STB is actually the same client and then immediately releases the original IP address. DHCP relay agent acts as a proxy client for auto logout and sends a DHCP release message to the DHCP server.

### How DHCP Identifies and Releases Clients

---

The auto logout feature requires that DHCP explicitly identify clients. By default, DHCP local server and DHCP relay agent identify clients based on MAC address or Client Identifier, and subnet. However, in some cases this type of identification might not be



sufficient. For example, in the previous STB example, each STB has a different MAC address, so DHCP incorrectly assumes that an upgraded or replacement STB is a new client.

In order to explicitly identify clients, auto logout uses a secondary identification method when the primary identification method is unsuccessful—the primary method is considered unsuccessful if the MAC address or Client Identifier does not match that of an existing client. Subscriber management supports two secondary identification methods that you can configure.

- Incoming interface method—DHCP views a new client connection on the interface as if it comes from the same client. DHCP deletes the existing client binding before creating a binding for the newly connected device. This method allows only one client device to connect on the interface.



**NOTE:** The incoming interface method differs from the `overrides interface-client-limit 1` statement, which retains the existing binding and rejects the newly connected client.

- Option 60 and option 82 method—DHCP considers two clients as different if they have the same option 60 and option 82 information, but different subnets.

DHCP local server and DHCP relay agent perform the following operations when auto logout is enabled and the secondary identification method identifies a duplicate client (that is, the Discover packet is from an existing client).

- DHCP local server immediately releases the existing address.
- DHCP relay agent immediately releases the existing client and then sends a DHCP release packet to the DHCP server. Sending the release packet ensures that DHCP relay and the DHCP server are synchronized.

If the DHCP relay receives a Discover message from an existing client, the DHCP relay forwards the Discover message to the DHCP server. The DHCP relay preserves the binding if the client's existing IP address is returned by the DHCP server. This behavior is not applicable if the proxy-mode override or client-discover-match functionality are enabled.



**NOTE:** If the DHCP relay agent is in snoop mode, DHCP relay releases the client but does not send a release packet to the DHCP server if the discover packet is for a passive client (a client added as a result of snooped packets) or if the discover packet is a snooped packet.

### Option 60 and Option 82 Requirements

DHCP local server requires that the received discover packet include both DHCP option 60 and option 82. If either option is missing, DHCP local server cannot perform the secondary identification method and auto logout is not used.

DHCP relay agent requires that the received discover packet contain DHCP option 60. DHCP relay determines the option 82 value based on the guidelines provided in [“How DHCP Relay Agent Uses Option 82 for Auto Logout” on page 1454](#).

**Related  
Documentation**

- [Automatically Logging Out DHCP Clients on page 1426](#)
- [How DHCP Relay Agent Uses Option 82 for Auto Logout on page 1454](#)
- *Allowing Only One DHCP Client Per Interface*
- *Clearing DHCP Bindings for Subscriber Access*

## Address-Assignment Pools Overview

The address-assignment pool feature supports subscriber management and DHCP management functionality by enabling you to create centralized IPv4 and IPv6 address pools independently of the client applications that use the pools. The **authd** process manages the pools and the address allocation, whether the addresses come from local pools or from a RADIUS server. For example, multiple client applications, such as DHCP, can use the same address-assignment pool to provide addresses for their particular clients. Client applications can acquire addresses for either authenticated or unauthenticated clients.

Address-assignment pools support both dynamic and static address assignment. In dynamic address assignment, a client is automatically assigned an address from the address-assignment pool. In static address assignment, which is supported for IPv4 pools only, you reserve an address that is then always used by a particular client. Addresses that are reserved for static assignment are removed from the dynamic address pool and cannot be assigned to other clients.

You can configure named address ranges within an address-assignment pool. A named range is a subset of the overall address range. A client application can use named ranges to manage address assignment based on client-specific criteria. For example, for IPv4 address-assignment pools, you might create a named range that is based on a specific DHCP option 82 value. Then, when a DHCP client request matches the specified option 82 value, an address from the specified range is assigned to the client.

You can link address-assignment pools together to provide backup pools for address assignment. When the primary pool is fully allocated, the router or switch automatically switches to the linked, or secondary, pool and begins allocating addresses from that pool.

You can also explicitly identify that an address-assignment pool is used for ND/RA.

**Related  
Documentation**

- *Configuring Address-Assignment Pools*
- *Address-Assignment Pools Licensing Requirements*
- *Example: Configuring an Address-Assignment Pool*

## Assign a Specific IP Address to a Client Using DHCP Option 50 and DHCPv6 IA\_NA Option

Subscriber management or DHCP management enables you to specify that DHCP local server assign a particular address to a client. For example, if a client is disconnected, you might use this capability to assign the same address that the client was using prior to being disconnected. If the requested address is available, DHCP assigns it to the client. If the address is unavailable, the DHCP local server offers another address, based on the address allocation process.

Both DHCP local server and DHCPv6 local server support the specific address request feature. DHCP local server uses DHCP option 50 in DHCP DISCOVER messages to request a particular address, while DHCPv6 local server uses the IA\_NA option (Identity Association for Non-Temporary Addresses) in DHCPv6 SOLICIT messages.



**NOTE:** Subscriber management (DHCP management) supports only one address for each of the DHCPv6 IA\_NA or IA\_PD address types. If the DHCPv6 client requests more than one address for a given type, the DHCPv6 local server uses only the first address and ignores the other addresses.

### Related Documentation

- [Extended DHCP Local Server Overview on page 1356](#)
- [DHCPv6 Local Server Overview on page 1361](#)

## Multiple Address Assignment for DHCPv6 Clients

Subscriber management (on the routers) or DHCP management (on the switches) enables you to assign multiple addresses to a single DHCPv6 client. Multiple address support is enabled by default, and is activated when the DHCPv6 local server receives a DHCPv6 Solicit message from a subscriber (or DHCP client) that contains multiple addresses.

For example, if you are implementing this feature on the routers, you might use the multiple address assignment feature in a networking environment in which a customer premises equipment (CPE) device requires a host address and a delegated prefix. In such an environment, you can configure subscriber management to assign both a DHCPv6 IA\_NA (Identity Association for Non-Temporary Addresses) and an IA\_PD (Identity Association for Prefix Delegation) address to the client (the CPE device).

- [Multiple Address Assignment Using Local Address Pools or RADIUS on page 1371](#)
- [Junos OS Predefined Variable for Multiple DHCPv6 Address Assignment on page 1372](#)

### Multiple Address Assignment Using Local Address Pools or RADIUS

You can use either local address pools or RADIUS when assigning multiple addresses to a DHCP client. When at least one address is successfully allocated, the router or switch creates a subscriber (or DHCP client) entry and binds the entry to the assigned address. If both addresses are successfully allocated, the router (or switch) creates a single subscriber (or DHCP client) entry and binds both addresses to that entry.

You can also configure a delegated address pool, which explicitly specifies the address pool that subscriber management (or DHCP management) uses to assign IPv6 prefixes for subscribers (or DHCP clients).

### Junos OS Predefined Variable for Multiple DHCPv6 Address Assignment

---



**NOTE:** EX Series switches do not support demux.

---

(On the routers only) Subscriber management provides a predefined variable that you can use to dynamically configure DHCPv6 multiple address assignment. You apply the Junos OS predefined variable, **\$junos-subscriber-ipv6-multi-address**, as a demux source address in a dynamic profile. When the dynamic profile is attached to a subscriber, the variable is expanded to include both the host and prefix addresses. You use this variable instead of the **\$junos-subscriber-ipv6-address** variable, which supports a single IPv6 address.

You include the **\$junos-subscriber-ipv6-multi-address** variable at the **[edit dynamic-profile profile-name interfaces interface-name unit logical-unit-number family inet6 demux-source]** hierarchy level.

#### **Related Documentation**

- [Specifying the Delegated Address Pool for IPv6 Prefix Assignment on page 1428](#)
- [Junos OS Predefined Variables](#)

## Centrally Configured Opaque DHCP Options

Subscriber management (on the routers) or DHCP management (on the switches) enables you to centrally configure DHCP options on a RADIUS server and then distribute the options on a per-subscriber or per DHCP-client basis. This method results in RADIUS-sourced DHCP options—the DHCP options originate at the RADIUS server and are sent to the subscriber (or DHCP client). This differs from the traditional client-sourced method (also called DHCP-sourced) of configuring DHCP options, in which the options originate at the client and are sent to the RADIUS server. The subscriber management (DHCP management) RADIUS-sourced DHCP options are also considered to be *opaque*, because DHCP local server performs minimal processing and error checking for the DHCP options string before passing the options to the subscriber (DHCP client).

Subscriber management (or DHCP management) uses Juniper Networks VSA 26-55 (DHCP-Options) to distribute the RADIUS-sourced DHCP options. The RADIUS server includes VSA 26-55 in the Access-Accept message that the server returns during subscriber authentication or DHCP client authentication. The RADIUS server sends the Access-Accept message to the RADIUS client, and then on to DHCP local server for return to the DHCP subscriber. The RADIUS server can include multiple instances of VSA 26-55 in a single Access-Accept message. The RADIUS client concatenates the multiple instances and uses the result as a single instance.

There is no CLI configuration required to enable subscriber management (DHCP management) to use the centrally configured DHCP options—the procedure is triggered by the presence of VSA 26-55 in the RADIUS Access-Accept message.

When building the offer packet for the DHCP client, DHCP local server uses the following sequence:

1. Processes any RADIUS-configured parameters that are passed as separate RADIUS attributes; for example, RADIUS attribute 27 (Session Timeout).
2. Processes any client-sourced parameters; for example, RADIUS attributes 53 (DHCP Message Type) and 54 (Server Identifier).
3. Appends (without performing any processing) the opaque DHCP options string contained in the VSA 26-55 received from the RADIUS server.

In addition to supporting central configuration of DHCP options directly on the RADIUS server (RADIUS-sourced options), subscriber management (DHCP management) also supports the traditional client-sourced options configuration, in which the router's (switch's) DHCP component sends the options to the RADIUS server. The client-sourced DHCP options method is supported for both DHCP local server and DHCP relay agent; however, the RADIUS-sourced central configuration method is supported on DHCP local server only. Both the RADIUS-sourced and client-sourced methods support DHCPv4 and DHCPv6 subscribers (clients).



**NOTE:** You can use the RADIUS-sourced and client-sourced methods simultaneously on DHCP local server. However, you must ensure that the central configuration method does not include options that override client-sourced DHCP options, because this can create unpredictable results.


---

- [Data Flow for RADIUS-Sourced DHCP Options on page 1375](#)
- [Multiple VSA 26-55 Instances Configuration on page 1376](#)
- [DHCP Options That Cannot Be Centrally Configured on page 1376](#)

### Data Flow for RADIUS-Sourced DHCP Options

Figure 5 on page 1375 shows the procedure subscriber management (DHCP management) uses when configuring DHCP options for subscribers (DHCP clients).

**Figure 5: DHCP Options Data Flow**



The following general sequence describes the data flow when subscriber management (DHCP management) uses RADIUS-sourced DHCP options and VSA 26-55 to configure a DHCP subscriber (client):

1. The subscriber (DHCP client) sends a DHCP discover message (or DHCPv6 solicit message) to the DHCP local server. The message includes client-sourced DHCP options.
2. The DHCP local server initiates authentication with the Junos OS RADIUS client.
3. The RADIUS client sends an Access-Request message on behalf of the subscriber (DHCP client) to the external RADIUS server. The message includes the subscriber's (DHCP client's) client-sourced DHCP options.
4. The external RADIUS server responds by sending an Access-Accept message to the RADIUS client. The Access-Accept message includes the RADIUS-sourced opaque DHCP options in VSA 26-55.
5. The RADIUS client sends the DHCP options string to DHCP local server. If there are multiple VSA 26-55 instances, the RADIUS client first assembles them into a single options string.
6. DHCP local server processes all options into the DHCP offer (or DHCPv6 reply) message, except for the RADIUS-sourced VSA 26-55 DHCP options. After processing all other options, DHCP local server then appends the unmodified VSA 26-55 DHCP options to the message and sends the message to the subscriber (DHCP client).

7. The subscriber (DHCP client) is configured with the DHCP options.
8. The following operations occur after the subscriber (DHCP client) receives the DHCP options:
  - Accounting—The RADIUS client sends Acct-Start and Interim-Accounting requests to the RADIUS server, including the RADIUS-sourced DHCP options in VSA 26-55. By default, the DHCP options are included in accounting requests.
  - Renewal—When the subscriber (DHCP client) renews, the cached DHCP options value is returned in the DHCP renew (or DHCPv6 ACK) message. The originally assigned DHCP options cannot be modified during a renew cycle.
  - Logout—When the subscriber (DHCP client) logs out, the RADIUS client sends an Acct-Stop message to the RADIUS server, including the RADIUS-sourced VSA 26-55.

### Multiple VSA 26-55 Instances Configuration

---

VSA 26-55 supports a maximum size of 247 bytes. If your RADIUS-sourced DHCP options field is greater than 247 bytes, you must break the field up and manually configure multiple instances of VSA 26-55 for the RADIUS server to return. When using multiple instances for an options field, you must place the instances in the packet in the order in which the fragments are to be reassembled by the RADIUS client. The fragments can be of any size of 247 bytes or less.



**BEST PRACTICE:** For ease of configuration and management of your DHCP options, you might want to have one DHCP option per VSA 26-55 instance, regardless of the size of the option field.

When the RADIUS client returns a reassembled opaque options field in an accounting request to the RADIUS server, the client uses 247-byte fragments. If you had originally created instances of fewer than 247 bytes, the returned fragments might not be the same as you originally configured on the RADIUS server.



**NOTE:** If you are configuring Steel-Belted Radius (SBR) to support multiple VSA 26-55 instances, ensure that you specify VSA 26-55 with the RO flags in the Subscriber Management RADIUS dictionary file. The R value indicates a multivalued reply attribute and the O value indicates an ordered attribute.

### DHCP Options That Cannot Be Centrally Configured

---

Table 135 on page 1377 shows the DHCP options that you must not centrally configure on the RADIUS server.



Table 135: Unsupported Opaque DHCP Options

DHCP Option	Option Name	Comments
Option 0	Pad Option	Not supported.
Option 51	IP Address Lease Time	Value is provided by RADIUS attribute 27 (Session-Timeout).
Option 52	Option Overload	Not supported.
Option 53	DHCP Message Type	Value is provided by DHCP local server.
Option 54	Server Identifier	Value is provided by DHCP local server.
Option 55	Parameter Request List	Value is provided by DHCP local server.
Option 255	End	Value is provided by DHCP local server.
–	DHCP magic cookie	Not supported.

**Related Documentation** • *Monitoring DHCP Options Configured on RADIUS Servers*

## Port Number Requirements for DHCP Firewall Filters

When you configure a firewall filter to perform some action on DHCP packets at the Routing Engine, such as protecting the Routing Engine by allowing only proper DHCP packets, you must specify both port 67 (bootps) and port 68 (bootpc) for both the source and destination. The firewall filter acts at both the line cards and the Routing Engine.

This requirement applies to both DHCP local server and DHCP relay, but it applies only when DHCP is provided by the `jdhcpd` process. MX Series routers, M120 routers, and M320 routers use `jdhcpd`. For DHCP relay, that means the configuration is required only at the `[edit forwarding-options dhcp-relay]` hierarchy level and not at the `[edit forwarding-options helpers bootp]` hierarchy level.

DHCP packets received on the line cards are encapsulated by `jdhcpd` with a new UDP header where their source and destination addresses are set to port 68 before being forwarded to the Routing Engine.

For DHCP relay and DHCP proxy, packets sent to the DHCP server from the router have both the source and destination UDP ports set to 67. The DHCP server responds using the same ports. However, when the line card receives these DHCP response packets, it changes both port numbers from 67 to 68 before passing the packets to the Routing Engine. Consequently the filter needs to accept port 67 for packets relayed from the client to the server, and port 68 for packets relayed from the server to the client.

Failure to include both port 67 and port 68 as described here results in most DHCP packets not being accepted.

For information about firewall filters, see *Firewall Filters Overview*.

**Related  
Documentation**

- *Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine*
- [Extended DHCP Local Server Overview on page 1356](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)
- *Understanding Dynamic Firewall Filters*

## DHCP Relay Agent

---

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [DHCP Relay Proxy Overview on page 1382](#)
- [DHCPv6 Relay Agent Overview on page 1384](#)
- [Configuring Group-Specific DHCP Relay Options on page 1384](#)
- [DHCP Snooping Support on page 1385](#)
- [DHCP Auto Logout Overview on page 1387](#)

## Extended DHCP Relay Agent Overview

You can configure extended DHCP relay options on the router or on the switch and enable the router (or switch) to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or DHCP client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.



**NOTE:** The PTX Series Packet Transport Routers do not support authentication for DHCP relay agents.

On the routers, you can use DHCP relay in carrier edge applications such as video/IPTV to obtain configuration parameters, including an IP address, for your subscribers. For information about how to use the DHCP relay agent in a video/IPTV application, see *Broadband Subscriber Management Edge Router Overview*.

On the switches, you can use DHCP relay to obtain configuration parameters including an IP address for DHCP clients.



**NOTE:** The extended DHCP relay agent options configured with the `dhcp-relay` statement are incompatible with the DHCP/BOOTP relay agent options configured with the `bootp` statement. As a result, you cannot enable both the extended DHCP relay agent and the DHCP/BOOTP relay agent on the router at the same time.

For information about the DHCP/BOOTP relay agent, see *Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents*.

You can also configure the extended DHCP relay agent to support IPv6 clients. See “[DHCPv6 Relay Agent Overview](#)” on [page 1384](#) for information about the DHCPv6 relay agent feature.

To configure the extended DHCP relay agent on the router (or switch), include the `dhcp-relay` statement at the `[edit forwarding-options]` hierarchy level. See the `[edit forwarding-options dhcp-relay]` *Hierarchy Level* for the complete DHCP relay agent syntax.

You can also include the `dhcp-relay` statement at the following hierarchy levels:

- `[edit logical-systems logical-system-name forwarding-options]`
- `[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options]`
- `[edit routing-instances routing-instance-name forwarding-options]`

This overview covers:

- [Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP Servers on page 1380](#)
- [DHCP Liveness Detection on page 1381](#)

### **Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP Servers**

The pattern of interaction among the DHCP Relay agent, DHCP client, and DHCP servers is the same regardless of whether the software installation is on a router or a switch. However, there are some difference in the details of usage.

On routers—In a typical carrier edge network configuration, the DHCP client is on the subscriber's computer, and the DHCP relay agent is configured on the router between the DHCP client and one or more DHCP servers.

On switches—In a typical network configuration, the DHCP client is on an access device such as a personal computer and the DHCP relay agent is configured on the switch between the DHCP client and one or more DHCP servers.

The following steps describe, at a high level, how the DHCP client, DHCP relay agent, and DHCP server interact in a configuration that includes two DHCP servers.

1. The DHCP client sends a discover packet to find a DHCP server in the network from which to obtain configuration parameters for the subscriber (or DHCP client), including an IP address.
2. The DHCP relay agent receives the discover packet and forwards copies to each of the two DHCP servers. The DHCP relay agent then creates an entry in its internal client table to keep track of the client's state.
3. In response to receiving the discover packet, each DHCP server sends an offer packet to the client. The DHCP relay agent receives the offer packets and forwards them to the DHCP client.
4. On receipt of the offer packets, the DHCP client selects the DHCP server from which to obtain configuration information. Typically, the client selects the server that offers the longest lease time on the IP address.
5. The DHCP client sends a request packet that specifies the DHCP server from which to obtain configuration information.
6. The DHCP relay agent receives the request packet and forwards copies to each of the two DHCP servers.
7. The DHCP server requested by the client sends an acknowledgement (ACK) packet that contains the client's configuration parameters.
8. The DHCP relay agent receives the ACK packet and forwards it to the client.
9. The DHCP client receives the ACK packet and stores the configuration information.
10. If configured to do so, the DHCP relay agent installs a host route and Address Resolution Protocol (ARP) entry for this client.
11. After establishing the initial lease on the IP address, the DHCP client and the DHCP server use unicast transmission to negotiate lease renewal or release. The DHCP relay

agent “snoops” on all of the packets unicast between the client and the server that pass through the router (or switch) to determine when the lease for this client has expired or been released. This process is referred to as *lease shadowing* or *passive snooping*.

### DHCP Liveness Detection

Liveness detection for DHCP subscriber or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients are expected to respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.



**NOTE:** DHCP liveness detection either globally or per DHCP group.

#### Related Documentation

- [DHCPv6 Relay Agent Overview on page 1384](#)
- [Access and Access-Internal Routes for Subscriber Management](#)
- [Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview](#)
- [Using External AAA Authentication Services with DHCP on page 1420](#)
- [DHCP Relay Proxy Overview on page 1382](#)
- [Graceful Routing Engine Switchover for DHCP](#)
- [High Availability Using Unified ISSU in the PPP Access Network](#)
- [Verifying and Managing DHCP Relay Configuration on page 1697](#)
- [Tracing Extended DHCP Operations](#)
- [Example: Minimum DHCP Relay Agent Configuration on page 1398](#)
- [Example: DHCP Relay Agent Configuration with Multiple Clients and Servers](#)
- [Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 1399](#)
- [Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing](#)
- [Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine](#)

## DHCP Relay Proxy Overview

DHCP relay proxy mode is an enhancement to extended DHCP relay. DHCP relay proxy supports all DHCP relay features while providing additional features and benefits.

Normally, extended DHCP relay operates as a helper application for DHCP operations. Except for the ability to add DHCP relay agent options and the gateway address (giaddr) to DHCP packets, DHCP relay is transparent to DHCP clients and DHCP servers, and simply forwards messages between DHCP clients and servers.

When you configure DHCP relay to operate in proxy mode, the relay is no longer transparent. In proxy mode, DHCP relay conceals DHCP server details from DHCP clients, which interact with a DHCP relay in proxy mode as though it is the DHCP server. For DHCP servers there is no change, because proxy mode has no effect on how the DHCP server interacts with the DHCP relay.

DHCP relay proxy provides the following benefits:

- DHCP server isolation and DoS protection—DHCP clients are unable to detect the DHCP servers, learn DHCP server addresses, or determine the number of servers that are providing DHCP support. Server isolation also provides denial-of-service (DoS) protection for the DHCP servers.
- Multiple lease offer selection—DHCP relay proxy receives lease offers from multiple DHCP servers and selects a single offer to send to the DHCP client, thereby reducing traffic in the network. Currently, the DHCP relay proxy selects the first offer received.
- Support for both numbered and unnumbered Ethernet interfaces—For DHCP clients connected through Ethernet interfaces, when the DHCP client obtains an address, the DHCP relay proxy adds an access internal host route specifying that interface as the outbound interface. The route is automatically removed when the lease time expires or when the client releases the address.
- Logical system support—DHCP relay proxy can be configured in a logical system, whereas a non-proxy mode DHCP relay cannot.



**NOTE:** You cannot configure both DHCP relay proxy and extended DHCP local server on the same interface.

---

## Interaction Among DHCP Relay Proxy, DHCP Client, and DHCP Servers

---

The DHCP relay agent is configured on the router (or switch), which operates between the DHCP client and one or more DHCP servers.

The following steps provide a high-level description of how DHCP relay proxy interacts with DHCP clients and DHCP servers.

1. The DHCP client sends a discover packet to locate a DHCP server in the network from which to obtain configuration parameters for the subscriber.
2. The DHCP relay proxy receives the discover packet from the DHCP client and forwards copies of the packet to each supporting DHCP server. The DHCP relay proxy then creates a client table entry to keep track of the client state.
3. In response to the discover packet, each DHCP server sends an offer packet to the client, which the DHCP relay proxy receives. The DHCP relay proxy does the following:
  - a. Selects the first offer received as the offer to sent to the client
  - b. Replaces the DHCP server address with the address of the DHCP relay proxy
  - c. Forwards the offer to the DHCP client.
4. The DHCP client receives the offer from the DHCP relay proxy.
5. The DHCP client sends a request packet that indicates the DHCP server from which to obtain configuration information—the request packet specifies the address of the DHCP relay proxy.
6. The DHCP relay proxy receives the request packet and forwards copies, which include the address of selected server, to all supporting DHCP servers.
7. The DHCP server requested by the client sends an acknowledgement (ACK) packet that contains the client configuration parameters.
8. The DHCP relay proxy receives the ACK packet, replaces the DHCP server address with its own address, and forwards the packet to the client.
9. The DHCP client receives the ACK packet and stores the configuration information.
10. If configured to do so, the DHCP relay proxy installs a host route and Address Resolution Protocol (ARP) entry for the DHCP client.
11. After the initial DHCP lease is established, the DHCP relay proxy receives all lease renewals and lease releases from the DHCP client and forwards them to the DHCP server.

**Related  
Documentation**

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [Enabling DHCP Relay Proxy Mode on page 1470](#)
- [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 1473](#)

## DHCPv6 Relay Agent Overview

The DHCPv6 relay agent enhances the extended DHCP relay agent by providing support in an IPv6 network. The DHCPv6 relay agent passes messages between the DHCPv6 client and the DHCPv6 server, similar to the way DHCP relay agent supports an IPv4 network.

When a DHCPv6 client logs in, the DHCPv6 relay agent uses the AAA service framework to interact with the RADIUS server to provide authentication and accounting. The RADIUS server, which is configured independently of DHCP, authenticates the client and supplies the IPv6 prefix and client configuration parameters, such as session timeout and the maximum number of clients allowed per interface.



**NOTE:** The PTX Series Packet Transport Routers do not support authentication for DHCPv6 relay agents.

The DHCPv6 relay agent is compatible with the extended DHCP local server and the extended DHCP relay agent, and can be enabled on the same interface as either the extended DHCP local server or DHCP relay agent.

To configure the DHCPv6 relay agent on the router (or switch), you include the **dhcpcv6** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level.

You can also include the **dhcpcv6** statement at the following hierarchy levels:

- **[edit logical-systems *logical-system-name* forwarding-options dhcp-relay]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay]**
- **[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay]**

### Related Documentation

- [Using External AAA Authentication Services with DHCP on page 1420](#)
- [Grouping Interfaces with Common DHCP Configurations on page 1421](#)
- [Configuring Group-Specific DHCP Relay Options on page 1384](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)
- [Configuring Passwords for Usernames on page 1439](#)
- [Creating Unique Usernames for DHCP Clients on page 1439](#)
- [Verifying and Managing DHCPv6 Local Server Configuration on page 1696](#)
- [Example: Extended DHCPv6 Local Server Configuration](#)

## Configuring Group-Specific DHCP Relay Options

You can include the following statements at the **[edit forwarding-options dhcp-relay group *group-name*]** hierarchy level to set group-specific DHCP relay agent configuration



options. Group-specific statements apply only to the named group of interfaces, and override any global DHCP relay agent settings for the same statement.

Include the statements at the **[edit forwarding-options dhcp-relay dhcpv6 group group-name]** hierarchy level to configure group-specific options for DHCPv6 relay agent.

- **active-server-group**—Configure an active server group to apply a common DHCP relay agent configuration to a named group of DHCP server addresses. For information, see [“Configuring Active Server Groups” on page 1469](#).
- **authentication**—Configure the parameters the router (or switch) sends to the external AAA server.
- **dynamic-profile**—Specify the dynamic profile that is attached to a group of interfaces.
- **interface**—Specify one or more interfaces, or a range of interfaces, that are within the specified group.
- **liveness-detection**—Configure bidirectional failure detection timers and authentication criteria for static routes. For more information, see [“DHCP Liveness Detection Overview” on page 1472](#).
- **overrides**—Override the default configuration settings for the extended DHCP relay agent. For information, see [“Overriding the Default DHCP Relay Configuration Settings” on page 1447](#).
- **relay-agent-interface-id**—(DHCPv6 only) Insert the DHCPv6 Relay Agent Interface-ID option (option 18) in DHCPv6 packets destined for the DHCPv6 server.
- **relay-agent-remote-id**—(DHCPv6 only) Insert the DHCPv6 Relay Agent Remote-ID option (option 37) in DHCPv6 packets destined for the DHCPv6 server.
- **relay-option**—Configure selective processing, which uses DHCP options in client packets to identify and filter client traffic, and to specify the action DHCP relay agent takes with the traffic. For more information, see [Using DHCP Option Information to Selectively Process DHCP Client Traffic](#).
- **relay-option-82**—(DHCPv4 only) Enable or disable the insertion of option 82 information in packets destined for a DHCP server. For information, see [“Using DHCP Relay Agent Option 82 Information” on page 1464](#).
- **service-profile**—Specify the default subscriber service, (or default profile) which is activated when the subscriber (or DHCP client) logs in and no other service is activated by a RADIUS server or a provisioning server. For more information, see [Default Subscriber Service Overview](#).

#### Related Documentation

- [Grouping Interfaces with Common DHCP Configurations on page 1421](#)

## DHCP Snooping Support

DHCP snooping provides DHCP security on the router or switch by filtering incoming messages. When DHCP snooping is enabled, the router differentiates between trusted and untrusted interfaces, and forwards messages from trusted sources while rejecting the untrusted messages.

In Junos OS, DHCP snooping is enabled in a routing instance when you configure either the **dhcp-relay** statement at the **[edit forwarding-options]** hierarchy level, or the **dhcp-local-server** statement at the **[edit system services]** hierarchy level in that routing instance. The router discards snooped packets by default. To enable normal processing of snooped packets, you must explicitly configure the **allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level.

You can configure DHCP snooping support for the following:

- DHCPv4 relay agent—Override the router's (or switch's) default snooping configuration and specify that DHCP snooping is enabled or disabled globally, for a named group of interfaces, or for a specific interface within a named group.

In a separate procedure, you can set a global configuration to specify whether the DHCPv4 relay agent forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces. The router also uses the global DHCP relay agent snooping configuration to determine whether to forward or drop snooped BOOTREPLY packets.

- DHCPv6 relay agent—As you can with snooping support for the DHCPv4 relay agent, you can override the default DHCPv6 relay agent snooping configuration on the router to explicitly enable or disable snooping support globally, for a named group of interfaces, or for a specific interface with a named group of interfaces.

In multi-relay topologies where more than one DHCPv6 relay agent is between the DHCPv6 client and the DHCPv6 server, snooping enables intervening DHCPv6 relay agents between the client and the server to correctly receive and process the unicast traffic from the client and forward it to the server. The DHCPv6 relay agent snoops incoming unicast DHCPv6 packets by setting up a filter with UDP port 547 (the DHCPv6 UDP server port) on a per-forwarding table basis. The DHCPv6 relay agent then processes the packets intercepted by the filter and forwards the packets to the DHCPv6 server.

Unlike the DHCPv4 relay agent, the DHCPv6 relay agent does not support global configuration of forwarding support for DHCPv6 snooped packets.

- DHCP local server—Configure whether DHCP local server forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces.

#### Related Documentation

- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 1437](#)
- [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 1455](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 1460](#)
- [Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 1403](#)
- [Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent](#)

## DHCP Auto Logout Overview

This topic provides an introduction to the DHCP auto logout feature and includes the following sections:

- [Auto Logout Overview on page 1387](#)
- [How DHCP Identifies and Releases Clients on page 1387](#)
- [Option 60 and Option 82 Requirements on page 1388](#)

### Auto Logout Overview

Auto logout is supported for DHCP local server and DHCP relay agent. It improves the efficiency of DHCP IP address assignment by allowing IP addresses to be immediately released and returned to the address pool when DHCP clients are no longer using the addresses. DHCP can then assign the addresses to other clients. Without auto logout, an IP address is blocked for the entire lease period, and DHCP must wait until the address lease time expires before reusing the address.

Auto logout is particularly useful when DHCP uses long lease times for IP address assignments and to help avoid allocating duplicate IP addresses for a single client.

For example, you might have an environment that includes set-top boxes (STB) that are often upgraded or replaced. Each time a STB is changed, the new STB repeats the DHCP discover process to obtain client configuration information and an IP address. DHCP views the new STB as a completely new client and assigns a new IP address—the previous IP address assigned to the client (the old STB) remains blocked and unavailable until the lease expires. If auto logout is configured in this situation, DHCP recognizes that the new STB is actually the same client and then immediately releases the original IP address. DHCP relay agent acts as a proxy client for auto logout and sends a DHCP release message to the DHCP server.

### How DHCP Identifies and Releases Clients

The auto logout feature requires that DHCP explicitly identify clients. By default, DHCP local server and DHCP relay agent identify clients based on MAC address or Client Identifier, and subnet. However, in some cases this type of identification might not be sufficient. For example, in the previous STB example, each STB has a different MAC address, so DHCP incorrectly assumes that an upgraded or replacement STB is a new client.

In order to explicitly identify clients, auto logout uses a secondary identification method when the primary identification method is unsuccessful—the primary method is considered unsuccessful if the MAC address or Client Identifier does not match that of an existing client. Subscriber management supports two secondary identification methods that you can configure.

- Incoming interface method—DHCP views a new client connection on the interface as if it comes from the same client. DHCP deletes the existing client binding before creating a binding for the newly connected device. This method allows only one client device to connect on the interface.



**NOTE:** The incoming interface method differs from the overrides `interface-client-limit 1` statement, which retains the existing binding and rejects the newly connected client.

- Option 60 and option 82 method—DHCP considers two clients as different if they have the same option 60 and option 82 information, but different subnets.

DHCP local server and DHCP relay agent perform the following operations when auto logout is enabled and the secondary identification method identifies a duplicate client (that is, the Discover packet is from an existing client).

- DHCP local server immediately releases the existing address.
- DHCP relay agent immediately releases the existing client and then sends a DHCP release packet to the DHCP server. Sending the release packet ensures that DHCP relay and the DHCP server are synchronized.

If the DHCP relay receives a Discover message from an existing client, the DHCP relay forwards the Discover message to the DHCP server. The DHCP relay preserves the binding if the client's existing IP address is returned by the DHCP server. This behavior is not applicable if the proxy-mode override or client-discover-match functionality are enabled.



**NOTE:** If the DHCP relay agent is in snoop mode, DHCP relay releases the client but does not send a release packet to the DHCP server if the discover packet is for a passive client (a client added as a result of snooped packets) or if the discover packet is a snooped packet.

---

### Option 60 and Option 82 Requirements

DHCP local server requires that the received discover packet include both DHCP option 60 and option 82. If either option is missing, DHCP local server cannot perform the secondary identification method and auto logout is not used.

DHCP relay agent requires that the received discover packet contain DHCP option 60. DHCP relay determines the option 82 value based on the guidelines provided in [“How DHCP Relay Agent Uses Option 82 for Auto Logout”](#) on page 1454.

#### Related Documentation

- [Automatically Logging Out DHCP Clients on page 1426](#)
- [How DHCP Relay Agent Uses Option 82 for Auto Logout on page 1454](#)
- [Allowing Only One DHCP Client Per Interface](#)
- [Clearing DHCP Bindings for Subscriber Access](#)

---

## Public Key Cryptography Overview

- [Understanding Public Key Cryptography on Switches on page 1389](#)

## Understanding Public Key Cryptography on Switches

Cryptography describes the techniques related to the following aspects of information security:

- Privacy or confidentiality
- Integrity of data
- Authentication
- Nonrepudiation or nonrepudiation of origin—Nonrepudiation of origin means that signers cannot claim that they did not sign a message while claiming that their private key remains secret. In some nonrepudiation schemes used in digital signatures, a timestamp is attached to the digital signature, so that even if the private key is exposed, the signature remains valid. Public and private keys are described in the following text.

In practice, cryptographic methods protect the data transferred from one system to another over public networks by encrypting the data using an encryption key. Public key cryptography (PKC), which is used on Juniper Networks EX Series Ethernet Switches, uses a pair of encryption keys: a public key and a private key. The public and private keys are created simultaneously using the same encryption algorithm. The private key is held by a user secretly and the public key is published. Data encrypted with a public key can be decrypted only with the corresponding private key and vice versa. When you generate a public/private key pair, the switch automatically saves the key pair in a file in the certificate store, from which it is subsequently used in certificate request commands. The generated key pair is saved as ***certificate-id.priv***.



**NOTE:** The default RSA and DSA key size is 1024 bits. If you are using the Simple Certificate Enrollment Protocol (SCEP), Juniper Networks Junos operating system (Junos OS) supports RSA only.

This topic describes:

- [Public Key Infrastructure \(PKI\) and Digital Certificates on page 1389](#)

### Public Key Infrastructure (PKI) and Digital Certificates

Public key infrastructure (PKI) allows the distribution and use of the public keys in public key cryptography with security and integrity. PKI manages the public keys by using digital certificates. A digital certificate provides an electronic means of verifying the identity of an individual, an organization, or a directory service that can store digital certificates.

A PKI typically consists of a Registration Authority (RA) that verifies the identities of entities, authorizes their certificate requests, and generates unique asymmetric key pairs (unless the users' certificate requests already contain public keys); and a Certificate Authority (CA) that issues corresponding digital certificates for the requesting entities. Optionally, you can use a Certificate Repository that stores and distributes certificates and a certificate revocation list (CRL) identifying the certificates that are no longer valid.

Each entity possessing the authentic public key of a CA can verify the certificates issued by that CA.

Digital signatures exploit the public key cryptographic system as follows:

1. A sender digitally signs data by applying a cryptographic operation, involving its private key, on a digest of the data.
2. The resulting signature is attached to the data and sent to the receiver.
3. The receiver obtains the digital certificate of the sender, which provides the sender's public key and confirmation of the link between its identity and the public key. The sender's certificate is often attached to the signed data.
4. The receiver either trusts this certificate or attempts to verify it. The receiver verifies the signature on the data by using the public key contained in the certificate. This verification ensures the authenticity and integrity of the received data.

As an alternative to using a PKI, an entity can distribute its public key directly to all potential signature verifiers, so long as the key's integrity is protected. The switch does it by using a self-signed certificate as a container for the public key and the corresponding entity's identity.

**Related  
Documentation**

- [Understanding Self-Signed Certificates on EX Series Switches on page 1390](#)

---

## Self-Signed Certificates Overview

---

- [Understanding Self-Signed Certificates on EX Series Switches on page 1390](#)

### Understanding Self-Signed Certificates on EX Series Switches

When you initialize a Juniper Networks EX Series Ethernet Switch with the factory default configuration, the switch generates a self-signed certificate, allowing secure access to the switch through the Secure Sockets Layer (SSL) protocol. Hypertext Transfer Protocol over Secure Sockets Layer (HTTPS) and XML Network Management over Secure Sockets Layer (XNM-SSL) are the two services that can make use of the self-signed certificates.



**NOTE:** Self-signed certificates do not provide additional security as do those generated by Certificate Authorities (CAs). This is because a client cannot verify that the server he or she has connected to is the one advertised in the certificate.

The switches provide two methods for generating a self-signed certificate:

- Automatic generation

In this case, the creator of the certificate is the switch. An automatically generated (also called “system-generated”) self-signed certificate is configured on the switch by default.

After the switch is initialized, it checks for the presence of an automatically generated self-signed certificate. If it does not find one, the switch generates one and saves it in the file system.

A self-signed certificate that is automatically generated by the switch is similar to an SSH host key. It is stored in the file system, not as part of the configuration. It persists when the switch is rebooted, and it is preserved when a **request system snapshot** command is issued.

The switch uses the following distinguished name for the automatically generated certificate:

**"CN=<device serial number>, CN=system generated, CN=self-signed"**

If you delete the system-generated self-signed certificate on the switch, the switch generates a self-signed certificate automatically.

- **Manual generation**

In this case, you create the self-signed certificate for the switch. At any time, you can use the CLI to generate a self-signed certificate. Manually generated self-signed certificates are stored in the file system, not as part of the configuration.

Self-signed certificates are valid for five years from the time they are generated. When the validity of an automatically generated self-signed certificate expires, you can delete it from the switch so that the switch generates a new self-signed certificate.

System-generated self-signed certificates and manually generated self-signed certificates can coexist on the switch.

**Related  
Documentation**

- [Understanding Public Key Cryptography on Switches on page 1389](#)
- [Manually Generating Self-Signed Certificates on Switches \(CLI Procedure\) on page 1417](#)





## CHAPTER 25

# Configuration

- [DHCP Local Server Examples on page 1393](#)
- [DHCP Relay Agent Examples on page 1398](#)
- [Configuration Tasks on page 1404](#)
- [Configuration Tasks for DHCP Local Server on page 1419](#)
- [Configuration Tasks for DHCP Relay Agent on page 1443](#)
- [DHCP Local Server Configuration Statements on page 1475](#)
- [DHCP Relay Agent Configuration Statements on page 1548](#)
- [Other Configuration Statements on page 1637](#)

### DHCP Local Server Examples

---

- [Example: Minimum Extended DHCP Local Server Configuration on page 1393](#)
- [Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 1394](#)
- [Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395](#)

#### Example: Minimum Extended DHCP Local Server Configuration

This example shows the minimum configuration you need to use for the extended DHCP local server on the router or switch:

```
[edit system services]
dhcp-local-server {
  group group_one {
    interface fe-0/0/2.0;
  }
}
```



**NOTE:** The interface type in this topic is just an example. The fe- interface type is not supported by EX Series switches.

This example creates the server group named **group\_one**, and specifies that the DHCP local server is enabled on interface **fe-0/0/2.0** within the group. The DHCP local server uses the default pool match configuration of **ip-address-first**.



**NOTE:** If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the **clear dhcp server binding** command before you delete the DHCP server configuration.

**Related  
Documentation**

- [Extended DHCP Local Server Overview on page 1356](#)

### Example: Extended DHCP Local Server Configuration with Optional Pool Matching

This example shows an extended DHCP local server configuration that includes optional IPv4 address-assignment pool matching and interface groups. For pool matching, this configuration specifies that the DHCP local server first check the response from an external authentication authority (for example, RADIUS) and use the Framed-IPv6-Pool attribute to determine the address-assignment pool to use for the client address. If no external authority match is found, the DHCP local server then uses **ip-address-first** matching together with the option 82 information to match the named address range for client IPv4 address assignment. The option 82 matching must also be included in the address-assignment pool configuration.

```
[edit system services]
dhcp-local-server {
  group group_one {
    interface fe-0/0/2.0;
    interface fe-0/0/2.1;
  }
  group group_two {
    interface fe-0/0/3.0;
    interface fe-0/0/3.1;
  }
  pool-match-order {
    external-authority
    ip-address-first;
    option-82;
  }
}
```



**NOTE:** The interface type in this topic is just an example. The **fe-** interface type is not supported by EX Series switches.

**Related  
Documentation**

- [Extended DHCP Local Server Overview on page 1356](#)
- [Address-Assignment Pools Overview on page 1370](#)

## Example: Configuring Group Liveness Detection for DHCP Local Server Clients

This example shows how to configure group liveness detection for DHCP local server subscribers or DHCP clients using Bidirectional Forwarding Detection (BFD) as the liveness detection method.

- [Requirements on page 1395](#)
- [Overview on page 1395](#)
- [Configuration on page 1395](#)

### Requirements

- Juniper Networks MX Series routers
- Juniper Networks EX Series switches
- Configure DHCP local server. See "[Extended DHCP Local Server Overview](#)" on page 1356.

### Overview

In this example, you configure group liveness detection for DHCP local server subscribers (clients) by completing the following operations:

1. Enable liveness detection for DHCP local server subscriber (or DHCP client) groups.
2. Specify BFD as the liveness detection method for all dynamically created DHCP local server subscribers (clients).
3. Configure BFD-specific statements to define how the protocol behaves.
4. Configure the action the router (switch) takes when a liveness detection failure occurs.



**NOTE:** This example explains how to configure liveness detection for a DHCPv4 network. Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the [liveness-detection](#) statement, and any subsequent configuration statements, at the `[edit system services dhcp-local-server dhcpv6]` or `[edit system services dhcp-local-server dhcpv6 group group-name]` hierarchy level.

### Configuration

#### Step-by-Step Procedure

To configure group liveness detection for DHCP local server:

1. Specify that you want to configure liveness detection.  

```
[edit system services dhcp-local-server ]
user@host# edit liveness-detection
```
2. Specify that you want to configure liveness detection for a specific DHCP local server group.  

```
[edit system services dhcp-local-server liveness-detection]
user@host# edit group local_group_1
```

3. Specify that you want to configure the liveness detection method.  

```
[edit system services dhcp-local-server group local_group_1 liveness-detection]  
user@host# edit method
```
4. Specify BFD as the liveness detection method that you want DHCP to use.  

```
[edit system services dhcp-local-server group local_group_1 liveness-detection  
method]  
user@host# edit bfd
```
5. Configure the detection time threshold (in milliseconds) at which a trap is produced.  

```
[edit system services dhcp-local-server group local_group_1 liveness-detection  
method bfd]  
user@host# set detection-time threshold 30000
```
6. Configure the time (in milliseconds) for which BFD holds a session up notification.  

```
[edit system services dhcp-local-server group local_group_1 liveness-detection  
method bfd]  
user@host# set holddown-interval 50
```
7. Configure the BFD minimum transmit and receive interval (in milliseconds).



**NOTE:** You do not need to configure the BFD minimum transmit and receive interval if you configure the minimum-interval for the BFD transmit-interval statement and the minimum-receive-interval.

```
[edit system services dhcp-local-servergroup local_group_1 liveness-detection method  
bfd]  
user@host# set minimum-interval 45000
```

8. Configure the minimum receive interval (in milliseconds).



**NOTE:** You do not need to configure the BFD minimum receive interval if you configure the BFD minimum transmit and receive interval.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection  
method bfd]  
user@host# set minimum-receive-interval 60000
```

9. Configure a multiplier value for the detection time.  

```
[edit system services dhcp-local-server group local_group_1 liveness-detection  
method bfd]  
user@host# set multiplier 100
```
10. Disable the ability for BFD interval timers to change or adapt to network situations.  

```
[edit system services dhcp-local-server group local_group_1 liveness-detection  
method bfd]  
user@host# set no-adaptation
```
11. Configure the BFD session mode.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
```

```
user@host# set session-mode automatic
```

12. Configure the threshold and minimum interval for the BFD transmit interval.



**NOTE:** You do not need to configure the transmit interval values if you have already configured the minimum transmit and receive interval for BFD.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
```

```
user@host# set transmit-interval threshold 60000 minimum-interval 45000
```

13. Configure the BFD protocol version you want to detect.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
```

```
user@host# set version automatic
```

14. Configure the action the router (switch) takes when a liveness detection failure occurs. In this example, the failure action is to clear the client session only when a liveness detection failure occurs and the local interface is detected as being up.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection]
user@host# edit failure-action action
```

**Results** From configuration mode, confirm your configuration by entering the **show system** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
regress@montag# show system
services {
  dhcp-local-server {
    group local_group_1 {
      liveness-detection {
        failure-action clear-binding-if-interface-up;
        method {
          bfd {
            version automatic;
            minimum-interval 45000;
            minimum-receive-interval 60000;
            multiplier 100;
            no-adaptation;
            transmit-interval {
              minimum-interval 45000;
              threshold 60000;
            }
          }
          detection-time {
            threshold 30000;
          }
        }
        session-mode automatic;
        holddown-interval 50;
      }
    }
  }
}
```

```
    }  
  }  
}  
}
```

If you are done configuring the device, enter **commit** from configuration mode.

**Related  
Documentation**

- [Extended DHCP Local Server Overview on page 1356](#)
- [DHCP Liveness Detection Overview on page 1472](#)
- [Configuring Detection of DHCP Local Server Client Connectivity on page 1436](#)

---

## DHCP Relay Agent Examples

- [Example: Minimum DHCP Relay Agent Configuration on page 1398](#)
- [Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 1399](#)
- [Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 1403](#)

### Example: Minimum DHCP Relay Agent Configuration

This example shows the minimum configuration you need to use the extended DHCP relay agent on the router or switch:

```
[edit forwarding-options]  
dhcp-relay {  
  server-group {  
    test 10.0.2.1;  
  }  
  active-server-group test;  
  group all {  
    interface fe-0/0/2.0;  
  }  
}
```



**NOTE:** The interface type in this topic is just an example. The **fe-** interface type is not supported by EX Series switches.

---

This example creates a server group and an active server group named **test** with IP address 10.0.2.1. The DHCP relay agent configuration is applied to a group named **all**. Within this group, the DHCP relay agent is enabled on interface fe-0/0/2.0.

**Related  
Documentation**

- [Extended DHCP Relay Agent Overview on page 1379](#)

## Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings

This example shows how to configure DHCP relay agent to use DHCP option strings to selectively identify, filter, and process client traffic.

- [Requirements on page 1399](#)
- [Overview on page 1399](#)
- [Configuration on page 1400](#)
- [Verification on page 1401](#)

### Requirements

---

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Routers or EX Series Switches

Before you configure DHCP relay agent selective processing support, be sure you:

- Configure DHCP relay agent.

See [“Extended DHCP Relay Agent Overview” on page 1379](#).

- (Optional) Configure a named DHCP local server group if you want to forward client traffic to a server group.

See [“Grouping Interfaces with Common DHCP Configurations” on page 1421](#).

### Overview

---

In this example, you configure DHCP relay agent to use DHCP option strings in client packets to selectively identify, filter, and process client traffic. To configure selective processing, you perform the following procedures:

1. Identify the client traffic—Specify the DHCP option that DHCP relay agent uses to identify the client traffic you want to process. The option you specify matches the option in the client traffic.
2. Configure a default action—Specify the default processing action, which DHCP relay uses for identified client traffic that does not satisfy any configured match criteria.
3. Create match filters and associate an action with each filter—Specify match criteria that filter the client traffic. The criteria can be an exact match or a partial match with the option string in the client traffic. Associate a processing action with each match criterion.

## Configuration

---

To configure DHCP relay agent selective processing based on DHCP option information, perform these tasks:

- [Configuring DHCP Relay Agent To Selectively Process Client Traffic Based on DHCP Option Strings on page 1400](#)
- [Results on page 1401](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the command into the CLI at the **[edit]** hierarchy level.

```
set forwarding-options dhcp-relay relay-option option-number 60
set forwarding-options dhcp-relay relay-option equals ascii video-gold forward-only
set forwarding-options dhcp-relay relay-option equals ascii video-bronze local-server-group
  servergroup-15
set forwarding-options dhcp-relay relay-option starts-with hexadecimal ffff
  local-server-group servergroup-east
set forwarding-options dhcp-relay relay-option default-action drop
```

### *Configuring DHCP Relay Agent To Selectively Process Client Traffic Based on DHCP Option Strings*

### Step-by-Step Procedure

To configure DHCP relay selective processing:

1. Specify that you want to configure DHCP relay agent support.  

```
[edit forwarding-options]
user@host# edit dhcp-relay
```
2. Specify the DHCP option that DHCP relay agent uses to identify incoming client traffic.  

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option option-number 60
```
3. Configure a default action, which DHCP relay agent uses when the incoming client traffic does not satisfy any configured match criteria.  

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option default-action drop
```
4. Configure an exact match condition and associated action that DHCP relay uses to process the identified client traffic.  

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option equals ascii video-gold forward-only
```
5. Configure a second exact match condition and associated action that DHCP relay uses to process client traffic.  

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option equals ascii video-bronze local-server-group
  servergroup-15
```



6. Configure a partial match criteria and associated action that DHCP relay uses to process client traffic.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option starts-with hexadecimal ffff local-server-group
servergroup-east
```

### Results

From configuration mode, confirm the results of your configuration by issuing the **show** statement at the **[edit forwarding-options]** hierarchy level. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit forwarding-options]
user@host# show
dhcp-relay {
  relay-option {
    option-number 60;
    equals {
      ascii video-gold {
        forward-only;
      }
    }
    equals {
      ascii video-bronze {
        local-server-group servergroup-15;
      }
    }
    default-action {
      drop;
    }
    starts-with {
      hexadecimal ffff {
        local-server-group servergroup-east;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

To verify the status of DHCP relay agent selective traffic processing, perform this task:

- [Verifying the Status of DHCP Relay Agent Selective Traffic Processing on page 1401](#)

#### *Verifying the Status of DHCP Relay Agent Selective Traffic Processing*

**Purpose** Verify the DHCP relay agent selective traffic processing status.

**Action** Display statistics for DHCP relay agent.

```
user@host> show dhcp relay statistics
Packets dropped:
  Total                30
  Bad hardware address  1
  Bad opcode           1
  Bad options          3
  Invalid server address 5
  No available addresses 1
  No interface match    2
  No routing instance match 9
  No valid local address 4
  Packet too short      2
  Read error            1
  Send error            1
  Option 60             1
  Option 82             2

Messages received:
  BOOTREQUEST          116
  DHCPDECLINE           0
  DHCPDISCOVER          11
  DHCPINFORM            0
  DHCPRELEASE           0
  DHCPREQUEST          105

Messages sent:
  BOOTREPLY             0
  DHCPOFFER             2
  DHCPACK               1
  DHCPNAK                0
  DHCPFORCERENEW        0

Packets forwarded:
  Total                4
  BOOTREQUEST           2
  BOOTREPLY             2
```

**Meaning** The **Packets forwarded** field in the **show dhcp relay statistics** command output displays the number of client packets that have been forwarded as a result of the selective traffic processing configuration. In this example, the output indicates the total number of packets that DHCP relay agent has forwarded, as well as a breakdown for the number of **BOOTREQUEST** and **BOOTREPLY** packets forwarded.

- Related Documentation**
- [Extended DHCP Relay Agent Overview on page 1379](#)
  - *DHCP Options and Selective Traffic Processing Overview*
  - *Using DHCP Option Information to Selectively Process DHCP Client Traffic*
  - *Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings*
  - *Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing*

## Example: Configuring DHCP Snooping Support for DHCP Relay Agent

This example shows how to configure DHCP snooping support for DHCP relay agent.

- [Requirements on page 1403](#)
- [Overview on page 1403](#)
- [Configuration on page 1403](#)

### Requirements

- Configure DHCP relay agent. See “[Extended DHCP Relay Agent Overview](#)” on page 1379.

### Overview

In this example, you configure DHCP snooping support for DHCP relay agent by completing the following operations:

- Override the default DHCP snooping configuration and enable DHCP snooping support for the interfaces in group **frankfurt**.
- Configure DHCP relay agent to forward snooped packets to only configured interfaces.



**NOTE:** By default, DHCP snooping is disabled globally.

### Configuration

#### Step-by-Step Procedure

To configure DHCP relay support for DHCP snooping:

1. Specify that you want to configure DHCP relay agent.  

```
[edit]
user@host# edit forwarding-options dhcp-relay
```
2. Specify the named group of interfaces on which DHCP snooping is supported.  

```
[edit forwarding-options dhcp-relay]
user@host# edit group frankfurt
```
3. Specify the interfaces that you want to include in the group. DHCP relay agent considers these as the configured interfaces when determining whether to forward or drop traffic.  

```
[edit forwarding-options dhcp-relay group frankfurt]
user@host# set interface fe-1/0/1.3 upto fe-1/0/1.9
```
4. Specify that you want to override the default configuration for the group.  

```
[edit forwarding-options dhcp-relay group frankfurt]
user@host# edit overrides
```
5. Enable DHCP snooping support for the group.  

```
[edit forwarding-options dhcp-relay group frankfurt overrides]
user@host# set allow-snooped-clients
```

6. Return to the **[edit forwarding-options dhcp-relay]** hierarchy level to configure the forwarding action and specify that DHCP relay agent forward snooped packets on only configured interfaces:

```
[edit forwarding-options dhcp-relay group frankfurt overrides]
user@host# up 2
```

7. Enable DHCP snooped packet forwarding for DHCP relay agent.

```
[edit forwarding-options dhcp-relay]
user@host# edit forward-snooped-clients
```

8. Specify that snooped packets are forwarded on only configured interfaces (the interfaces in group **frankfurt**).

```
[edit forwarding-options dhcp-relay forward-snooped-clients]
user@host# set configured-interfaces
```

**Results** From configuration mode, confirm your configuration by entering the **show forwarding-options** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it. The following output also shows a range of configured interfaces in group **frankfurt**.

```
[edit]
regress@montag# show forwarding-options
dhcp-relay {
  forward-snooped-clients configured-interfaces;
  group frankfurt {
    overrides {
      allow-snooped-clients;
    }
    interface fe-1/0/1.3 {
      upto fe-1/0/1.9;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

- Related Documentation**
- [DHCP Snooping Support on page 1367](#)
  - [Configuring DHCP Snooping for DHCP Relay Agent on page 1455](#)

---

## Configuration Tasks

- [Configuring DHCP Services \(J-Web Procedure\) on page 1405](#)
- [Configuring a DHCP SIP Server \(CLI Procedure\) on page 1412](#)
- [Configuring a DHCP Client \(CLI Procedure\) on page 1413](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 1414](#)
- [Configuring an Extended DHCP Relay Server on EX Series Switches \(CLI Procedure\) on page 1416](#)

- [Enabling HTTPS and XNM-SSL Services on Switches Using Self-Signed Certificates \(CLI Procedure\) on page 1417](#)
- [Manually Generating Self-Signed Certificates on Switches \(CLI Procedure\) on page 1417](#)
- [Deleting Self-Signed Certificates \(CLI Procedure\) on page 1419](#)

## Configuring DHCP Services (J-Web Procedure)

- [Configuring DHCP Services \(J-Web Procedure\) on EX Series Switches on page 1405](#)
- [Configuring DHCP Services on EX4300 Switches \(J-Web Procedure\) on page 1408](#)

### Configuring DHCP Services (J-Web Procedure) on EX Series Switches



**NOTE:** This topic applies only to the J-Web Application package.

Use the J-Web DHCP Configuration pages to configure DHCP pools for subnets and static bindings for DHCP clients on an ACX Series Universal Access Gateway router or an EX Series Ethernet Switch. If DHCP pools or static bindings are already configured, use the [Configure Global DHCP Parameters Configuration](#) page to add settings for these pools and static bindings. Settings that have been previously configured for DHCP pools or static bindings are not overridden when you use the [Configure Global DHCP Parameters Configuration](#) page.

To configure the DHCP server:

1. Select **Configure > Services > DHCP**
2. Access a DHCP Configuration page:
  - To configure a DHCP pool for a subnet, click **Add** in the DHCP Pools box.
  - To configure a static binding for a DHCP client, click **Add** in the DHCP Static Binding box.
  - To globally configure settings for existing DHCP pools and static bindings, click **Configure Global DHCP Parameters**.
3. Enter information into the DHCP Service Configuration pages as described in [Table 136 on page 1406](#)
4. To apply the configuration, click **Apply**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

Table 136: DHCP Service Configuration Pages Summary

Field	Function	Your Action
<b>DHCP Pool Information</b>		
DHCP Subnet (required)	Specifies the subnet on which DHCP is configured.	Type an IP address prefix.
Address Range (Low) (required)	Specifies the lowest address in the IP address pool range.	Type an IP address that is part of the subnet specified in DHCP Subnet field .
Address Range (High) (required)	Specifies the highest address in the IP address pool range.	Type an IP address that is part of the subnet specified in DHCP Subnet. This address must be greater than the address specified in the Address Range (Low) field.
Exclude Addresses	Specifies addresses to exclude from the IP address pool.	<ul style="list-style-type: none"> <li>To add an excluded address, type the address next to the <b>Add</b> button, and click <b>Add</b>.</li> <li>To delete an excluded address, select the address in the Exclude Addresses box, and click <b>Delete</b>.</li> </ul>
<b>Lease Time</b>		
Maximum Lease Time (Seconds)	Specifies the maximum length of time a client can hold a lease. (Dynamic BOOTP lease lengths can exceed this maximum time.)	Type a number from 60 through 4,294,967,295 (seconds). You can also type <b>infinite</b> to specify a lease that never expires.
Default Lease Time (Seconds)	Specifies the length of time a client can hold a lease for clients that do not request a specific lease length.	Type a number from 60 through 2,147,483,647 (seconds). You can also type <b>infinite</b> to specify a lease that never expires.
<b>Server Information</b>		
Server Identifier	Specifies the IP address of the DHCP server reported to a client.	Type the IP address of the server. If you do not specify a server identifier, the primary address of the interface on which the DHCP exchange occurs is used.
Domain Name	Specifies the domain name that clients must use to resolve hostnames.	Type the name of the domain.
Domain Search	Specifies the order—from top to bottom—in which clients must append domain names when resolving hostnames using DNS.	<ul style="list-style-type: none"> <li>To add a domain name, type the name next to the <b>Add</b> button, and click <b>Add</b>.</li> <li>To delete a domain name, select the name in the Domain Search box, and click <b>Delete</b>.</li> </ul>
DNS Name Servers	Defines a list of DNS servers that the client can use, in the specified order—from top to bottom.	<ul style="list-style-type: none"> <li>To add a DNS server, type an IP address next to the <b>Add</b> button, and click <b>Add</b>.</li> <li>To remove a DNS server, select the IP address in the DNS Name Servers box, and click <b>Delete</b>.</li> </ul>

Table 136: DHCP Service Configuration Pages Summary (*continued*)

Field	Function	Your Action
Gateway Routers	Defines a list of relay agents on the subnet, in the specified order—from top to bottom.	<ul style="list-style-type: none"> <li>To add a relay agent, type an IP address next to the <b>Add</b> button, and click <b>Add</b>.</li> <li>To remove a relay agent, select the IP address in the Gateway Routers box, and click <b>Delete</b>.</li> </ul>
WINS Servers	Defines a list of NetBIOS name servers, in the specified order—from top to bottom.	<ul style="list-style-type: none"> <li>To add a NetBIOS name server, type an IP address next to the <b>Add</b> button, and click <b>Add</b>.</li> <li>To remove a NetBIOS name server, select the IP address in the WINS Servers box, and click <b>Delete</b>.</li> </ul>
<b>Boot Options</b>		
Boot File	Specifies the path and filename of the initial boot file to be used by the client.	Type a path and filename.
Boot Server	Specifies the Trivial File Transfer Protocol (TFTP) server that the client uses to obtain the client configuration file.	Type the IP address or hostname of the TFTP server.
<b>DHCP Static Binding Information</b>		
DHCP MAC Address (required)	Specifies the MAC address of the client to be permanently assigned a static IP address.	Type the hexadecimal MAC address of the client.
Fixed IP Addresses (required)	Defines a list of IP addresses permanently assigned to the client. A static binding must have at least one fixed address assigned to it, but multiple addresses are also allowed.	<ul style="list-style-type: none"> <li>To add an IP address, type it next to the <b>Add</b> button, and click <b>Add</b>.</li> <li>To remove an IP address, select it in the Fixed IP Addresses box, and click <b>Delete</b>.</li> </ul>
Host Name	Specifies the name of the client used in DHCP messages exchanged between the server and the client. The name must be unique to the client within the subnet on which the client resides.	Type a client hostname.
Client Identifier	Specifies the name of the client used by the DHCP server to index its database of address bindings. The name must be unique to the client within the subnet on which the client resides.	Type a client identifier in string form.
Hexadecimal Client Identifier	Specifies the name of the client, in hexadecimal form, used by the DHCP server to index its database of address bindings. The name must be unique to the client within the subnet on which the client resides.	Type a client identifier in hexadecimal form.

### Configuring DHCP Services on EX4300 Switches (J-Web Procedure)

On EX4300 switches, use the DHCP Configuration page to create DHCP pools and set the DHCP parameters for them and to configure DHCP settings for existing DHCP pools and static bindings.

To configure the DHCP services on EX4300 switches:

1. Select **Configure > Services > DHCP**
2. Access a DHCP Configuration page:
  - To configure a DHCP pool for a subnet, click **Add** in the DHCP Pools box.
  - To configure DHCP groups, click **Add** in the DHCP Groups box.
  - To globally configure settings for existing DHCP pools and static bindings, click **Configure Global DHCP Parameters**.
3. Enter information into the DHCP Service Configuration pages as described in [Table 137 on page 1408](#)
4. To apply the configuration, click **OK**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

**Table 137: DHCP Service Configuration Pages Summary for EX4300 Switches**

Field	Function	Your Action
<b>DHCP Groups</b>		
Group Name	Specifies the name of the group.	Enter the name of the group.
Interfaces	Family inet interface is listed, only if it is already configured with family inet.	Select the interface for the specific group.
<b>DHCP Pool Information</b>		
Pool Name	Specifies the name of an address-assignment pool.	Type the pool name.
Link Pool	Specifies the pool name to which it is linked.	Select the option from the list.
<b>Network Address</b>		
IP Address	Specifies the IP address pool range.	Type an IP address that is part of the subnet specified in the DHCP Subnet field.



Table 137: DHCP Service Configuration Pages Summary for EX4300 Switches (*continued*)

Field	Function	Your Action
Subnet mask	Specifies the subnet specified in DHCP Subnet.	Type a subnet mask that is specified in the DHCP Subnet field.
<b>DHCP Pool Attributes</b>		
Pool Name	Displays the name of an address-assignment pool.	The pool name is displayed.
Server Identifier	Specifies the IP address of the DHCP server reported to a client.	Type the IP address of the server. If you do not specify a server identifier, the primary address of the interface on which the DHCP exchange occurs is used.
TFTP Server	Specifies the Trivial File Transfer Protocol (TFTP) server that the client uses to obtain the client configuration file.	Enter the IP address of the TFTP server.
Maximum Lease Time (Seconds)	Specifies the maximum length of time a client can hold a lease. (Dynamic BOOTP lease lengths can exceed this maximum time.)	Type a number.
Boot File	Specifies the path and filename of the initial boot file to be used by the client.	Type a path and filename.
Boot Server	Specifies the TFTP server that provides the initial boot file to the client.	Type the IP address or hostname of the TFTP server.
Grace Period	Specifies the grace period for which a client can hold a lease.	Type the grace period in seconds.
DNS Name Servers	Defines a list of DNS servers the client can use.	<ul style="list-style-type: none"> <li>To add a DNS server, click <b>Add</b>. Type an IP address in the <b>Add IP Address</b> pop-up window.</li> <li>Click <b>OK</b>.</li> <li>To remove a DNS server, select the IP address in the DNS Name Servers box, and click <b>Remove</b>.</li> </ul>
WINS Servers	Defines a list of NetBIOS name servers.	<ul style="list-style-type: none"> <li>To add a NetBIOS name server, click <b>Add</b>. Type an IP address in the <b>Add IP Address</b> pop-up window.</li> <li>Click <b>OK</b>.</li> <li>To remove a NetBIOS name server, select the IP address in the WINS Servers box, and click <b>Remove</b>.</li> </ul>
Domain Name	Specifies the domain name that clients must use to resolve hostnames.	Type the name of the domain.
NetBIOS Node Type	Specifies the NetBOIS node that provides the initial node file to the client.	Select the type from the list.

Table 137: DHCP Service Configuration Pages Summary for EX4300 Switches (*continued*)

Field	Function	Your Action
Gateway Routers	Defines a list of relay agents on the subnet, in the specified order—from top to bottom.	<ul style="list-style-type: none"> <li>To add a relay agent, click <b>Add</b>. Type an IP address in the <b>Add IP Address</b> pop-up window.</li> <li>Click <b>OK</b>.</li> <li>To remove a relay agent, select the IP address in the Gateway Routers box, and click <b>Remove</b>.</li> </ul>
Option	Specifies the DHCP options.	<ul style="list-style-type: none"> <li>To add a DHCP option, click <b>Add</b>. The Add DHCP Option pop-up window is displayed. Enter the following: <ul style="list-style-type: none"> <li>Enter the DHCP <b>Code</b> in the Code box.</li> <li>Select the DHCP type from the <b>Type</b> list.</li> <li>Select the DHCP subtype from the <b>Sub Type</b> list.</li> <li>Enter the DHCP value in the <b>Value</b> box.</li> </ul> </li> <li>Click <b>OK</b>.</li> <li>To remove a DHCP option, select the option in the Option box, and click <b>Remove</b>.</li> </ul>
Option-82		
Circuit Identifier	Identifies the circuit (interface or VLAN or both) on the switch on which the request was received.	Type the circuit identifier.
Ranges	Specifies the circuit identifier range.	Type the range for the circuit identifier.
Remote Identifier	By default, the remote ID is the MAC address of the switch	Type the remote identifier.
Ranges	Specifies the remote identifier range.	Type the range for the remote identifier.
Address Range		
Range Name	Specifies the name of the range.	Click <b>Add</b> . The Add Address Range pop-up window is displayed: <ul style="list-style-type: none"> <li>Type the range name in the <b>Range Name</b> box.</li> </ul>
Address Range		
Address Range (Low)	Specifies the lowest address in the IP address pool range.	Type an IP address that is part of the subnet specified in DHCP Subnet
Address Range (High)	Specifies the highest address in the IP address pool range.	Type an IP address that is part of the subnet specified in DHCP Subnet. This address must be greater than the address specified in Address Range (Low).

Table 137: DHCP Service Configuration Pages Summary for EX4300 Switches (*continued*)

Field	Function	Your Action
Static Bindings		
Host Name	Specifies the name of the client used in DHCP messages exchanged between the server and the client. The name must be unique to the client within the subnet on which the client resides.	Type a client hostname.
MAC Address	Specifies the MAC address of the client to be permanently assigned a static IP address.	Type the hexadecimal MAC address of the client.
Fixed IP Address	Specifies the IP address of the client.	Type the IP address.
Global Settings		
General		
Duplicate clients on interface	Specifies the DHCP local server to include the client subinterface when distinguishing between duplicate DHCP clients (clients with the same MAC address or client ID) in the same subnet.	To enable this option, select the check box.
Pool Match Order	Specifies the order in which the DHCP local server uses information in the DHCP client PDU to determine how to obtain an address for the client.	Select the pool match order.
Authentication		
Password	Specifies the password that is sent to the external AAA authentication server for subscriber authentication.	Type the password.
Username-include		
Circuit Type	Specifies the circuit type that is linked with the username.	To enable this option, select the check box.
Interface Name	Name of the interface.	To enable this option, select the check box.
Mac Address	Specifies the MAC address of the client PDU that is linked with the username during the subscriber authentication process.	To enable this option, select the check box.
Logical System Name	Specifies that the logical system name that is linked with the username during the subscriber authentication process.	To enable this option, select the check box.

Table 137: DHCP Service Configuration Pages Summary for EX4300 Switches (*continued*)

Field	Function	Your Action
Option-60	Specifies the payload of Option 60 (Vendor Class Identifier) from the client PDU be linked with the username during the subscriber authentication process.	To enable this option, select the check box.
Routing Instance Name	Specifies the routing instance name that is linked with the username during the subscriber authentication process.	To enable this option, select the check box.
Option-82		
Circuit Identifier	Specifies the name of the client used by the DHCP server to index its database of address bindings. The name must be unique to the client within the subnet on which the client resides.	To enable this option, select the check box.
Remote Identifier	Specifies the remote ID option in the client.	To enable this option, select the check box.
Domain Name	Specifies the domain name that clients must use to resolve hostnames.	Type the domain name.
User Prefix	Specifies the prefix to the username as defined by the user.	Type the prefix.
Delimiter	Specifies a character that separates components that make up the username.	Type the delimiter.

- Related Documentation**
- [Understanding DHCP Services for Switches](#)
  - [Monitoring DHCP Services on page 1691](#)

## Configuring a DHCP SIP Server (CLI Procedure)

You can use the **sip-server** statement on the EX Series switch to configure option 120 on a DHCP server. The DHCP server sends configured option values—Session Initiation Protocol (SIP) server addresses or names—to DHCP clients when they request them. Previously, you were only allowed to specify a SIP server by address using **[edit system services dhcp option 120]**. You specify either an IPv4 address or a fully qualified domain name to be used by SIP clients to locate a SIP server. You cannot specify both an address and name in the same statement.

To configure a SIP server using the **address** option:

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

For example, to configure one address:

```
[edit system services dhcp]
user@switch set sip-server 172.168.0.11
```

To configure a SIP server using the *name* option:

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

For example, to configure a name:

```
[edit system services dhcp]
user@switch set sip-server abc.example.com
```

- Related Documentation**
- [Configuring a DHCP Client \(CLI Procedure\) on page 1413](#)
  - [Understanding DHCP Services for Switches](#)

## Configuring a DHCP Client (CLI Procedure)

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

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

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

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

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

The options that you can configure are listed in [Table 138 on page 1413](#). Replace the variable *configuration-statement* with one or more of the statements listed in this table. If you do not explicitly configure these options, the switch uses default values for them.

**Table 138: DHCP Client Settings**

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

Table 138: DHCP Client Settings (*continued*)

Configuration Statement	Description
<code>retransmission-interval</code>	Time between transmission attempts.
<code>server-address</code>	IP address of the server that the client queries for an IP address.
<code>update-server</code>	TCP/IP settings learned from an external DHCP server to the DHCP server running on the switch are propagated.
<code>vendor-option</code>	Vendor class ID (CPU's manufacturer ID string) for the DHCP client.

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

## Configuring a DHCP Server on Switches (CLI Procedure)



**NOTE:** This topic applies to Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Configuring a DHCP Server on Switches \(CLI Procedure\)](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

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

A DHCP configuration consists of two components—an optional reconfiguration of default settings on DHCP clients and the configuration of a DHCP server. This topic covers configuration of the DHCP server. For information about reconfiguring a DHCP client, see [“Configuring a DHCP Client \(CLI Procedure\)” on page 1413](#).



**NOTE:** Junos OS for EX Series switches with support for ELS does not support the DHCP legacy server version.

This topic includes the following tasks:

1. [Configuring an Extended DHCP Server on a Switch on page 1415](#)

### Configuring an Extended DHCP Server on a Switch

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

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

1. Create an address pool for DHCP IP addresses:

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

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

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

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

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

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

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

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

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

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

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

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

#### Related Documentation

- [Configuring a DHCP Client \(CLI Procedure\) on page 1413](#)
- [Configuring a DHCP SIP Server \(CLI Procedure\) on page 1412](#)
- [Understanding DHCP Services for Switches](#)

## Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)

You can configure an EX Series switch to act as an extended DHCP relay agent. This means that a locally attached host can issue a DHCP request as a broadcast message and the switch configured for DHCP relay relays the message to a specified DHCP server. Configure a switch to be a DHCP relay agent if you have locally attached hosts and a remote DHCP server.

Before you begin:

- Ensure that the switch can connect to the DHCP server.

To configure a switch to act as an extended DHCP relay agent server:

1. Create at least one DHCP server group, which is a group of 1 through 5 DHCP server IP addresses:

```
[edit forwarding-options dhcp-relay]
user@switch# set server-group server-group-name ip-address
```

2. Set the global active DHCP server group. The DHCP relay server relays DHCP client requests to the DHCP servers defined in the active server group:

```
[edit forwarding-options dhcp-relay]
user@switch# set active-server-group server-group-name
```

3. Create a DHCP relay group that includes at least one interface. DHCP relay runs on the interfaces defined in DHCP groups:

```
[edit forwarding-options dhcp-relay]
user@switch# set group group-name interface interface-name
```

4. (Optional) Configure overrides of default DHCP relay behaviors, at the global level. See the override options in the [overrides](#) statement.

```
[edit forwarding-options dhcp-relay]
user@switch# set overrides
```

5. (Optional) Configure DHCP relay to use the DHCP vendor class identifier option (option 60) in DHCP client packets, at the global level:

```
[edit forwarding-options dhcp-relay]
user@switch# set relay-option option-number 60
```

6. (Optional) Configure settings for a DHCP relay group that override the settings at the global level, using these statements:

```
[edit forwarding-options dhcp-relay group group-name]
user@switch# set active-server-group server-group-name
user@switch# set overrides
user@switch# set relay-option option-number 60
```

7. (Optional) Configure settings for a DHCP relay group interface that override the settings at the global and **group** levels, using these statements:

```
[edit forwarding-options dhcp-relay group group-name interface interface-name]
user@switch# exclude
user@switch# set overrides
```



```
user@switch# set trace
user@switch# set upto upto-interface-name
```

#### Related Documentation

- [Configuring a DHCP Server on Switches \(CLI Procedure\)](#)
- [Configuring a DHCP Client \(CLI Procedure\) on page 1413](#)
- [Understanding the Extended DHCP Relay Agent for EX Series Switches](#)

## Enabling HTTPS and XNM-SSL Services on Switches Using Self-Signed Certificates (CLI Procedure)

You can use the system-generated self-signed certificate or a manually generated self-signed certificate to enable Web management HTTPS and XNM-SSL services.

- To enable HTTPS services using the automatically generated self-signed certificate:

```
[edit]
user@switch# set system services web-management https system-generated-certificate
```

- To enable HTTPS services using a manually generated self-signed certificate:

```
[edit]
user@switch# set system services web-management https pki-local-certificate
certificate-id-name
```



**NOTE:** The value of the *certificate-id-name* must match the name you specified when you generated the self-signed certificate manually.

- To enable XNM-SSL services using a manually generated self-signed certificate:

```
[edit]
user@switch# set system services xnm-ssl local-certificate certificate-id-name
```



**NOTE:** The value of the *certificate-id-name* must match the name you specified when you generated the self-signed certificate manually.

#### Related Documentation

- [Manually Generating Self-Signed Certificates on Switches \(CLI Procedure\) on page 1417](#)
- [Understanding Self-Signed Certificates on EX Series Switches on page 1390](#)

## Manually Generating Self-Signed Certificates on Switches (CLI Procedure)

EX Series switches allow you to generate custom self-signed certificates and store them in the file system. The certificate you generate manually can coexist with the automatically generated self-signed certificate on the switch. To enable secure access to the switch

over SSL, you can use either the system-generated self-signed certificate or a certificate you have generated manually.

To generate self-signed certificates manually, you must complete the following tasks:

- [Generating a Public-Private Key Pair on Switches on page 1418](#)
- [Generating Self-Signed Certificates on Switches on page 1418](#)

### Generating a Public-Private Key Pair on Switches

---

A digital certificate has an associated cryptographic key pair that is used to sign the certificate digitally. The cryptographic key pair comprises a public key and a private key. When you generate a self-signed certificate, you must provide a public-private key pair that can be used to sign the self-signed certificate. Therefore, you must generate a public-private key pair before you can generate a self-signed certificate.

To generate a public-private key pair:

```
user@switch> request security pki generate-key-pair certificate-id certificate-id-name
```



**NOTE:** Optionally, you can specify the encryption algorithm and the size of the encryption key. If you do not specify the encryption algorithm and encryption key size, default values are used. The default encryption algorithm is RSA, and the default encryption key size is 1024 bits.

After the public-private key pair is generated, the switch displays the following:

```
generated key pair certificate-id-name, key size 1024 bits
```

### Generating Self-Signed Certificates on Switches

---

To generate the self-signed certificate manually, include the certificate ID name, the subject of the distinguished name (DN), the domain name, the IP address of the switch, and the e-mail address of the certificate holder:

```
user@switch> request security pki local-certificate generate-self-signed certificate-id  
certificate-id-name domain-name domain-name email email-address ip-address switch-ip-address  
subject subject-of-distinguished-name
```

The certificate you have generated is stored in the switch's file system. The certificate ID you have specified while generating the certificate is a unique identifier that you can use to enable the HTTPS or XNM-SSL services.

To verify that the certificate was generated and loaded properly, enter the **show security pki local-certificate** operational command.

#### Related Documentation

- [Enabling HTTPS and XNM-SSL Services on Switches Using Self-Signed Certificates \(CLI Procedure\) on page 1417](#)
- [Understanding Self-Signed Certificates on EX Series Switches on page 1390](#)

## Deleting Self-Signed Certificates (CLI Procedure)

You can delete a self-signed certificate that is automatically or manually generated from the EX Series switch. When you delete the automatically generated self-signed certificate, the switch generates a new self-signed certificate and stores it in the file system.

- To delete the automatically generated certificate and its associated key pair from the switch:

```
user@switch> clear security pki local-certificate system-generated
```

- To delete a manually generated certificate and its associated key pair from the switch:

```
user@switch> clear security pki local-certificate certificate-id certificate-id-name
```

- To delete all manually generated certificates and their associated key pairs from the switch:

```
user@switch> clear security pki local-certificate all
```

### Related Documentation

- [Manually Generating Self-Signed Certificates on Switches \(CLI Procedure\) on page 1417](#)
- [Understanding Self-Signed Certificates on EX Series Switches on page 1390](#)

## Configuration Tasks for DHCP Local Server

- [Using External AAA Authentication Services with DHCP on page 1420](#)
- [Grouping Interfaces with Common DHCP Configurations on page 1421](#)
- [Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces on page 1422](#)
- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)
- [Specifying the Maximum Number of DHCP Clients Per Interface on page 1424](#)
- [Automatically Logging Out DHCP Clients on page 1426](#)
- [Enabling Processing of Client Information Requests on page 1427](#)
- [Specifying the Delegated Address Pool for IPv6 Prefix Assignment on page 1428](#)
- [Enabling DHCPv6 Rapid Commit Support on page 1428](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)
- [Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429](#)
- [Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 1432](#)
- [Configuring Deletion of the Client When Dynamic Reconfiguration Fails on page 1433](#)
- [Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 1433](#)
- [Configuring a Token for DHCP Local Server Authentication on page 1434](#)
- [Preventing Binding of Clients That Do Not Support Reconfigure Messages on page 1434](#)
- [Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings on page 1435](#)
- [Configuring Detection of DHCP Local Server Client Connectivity on page 1436](#)

- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 1437](#)
- [Configuring Passwords for Usernames on page 1439](#)
- [Creating Unique Usernames for DHCP Clients on page 1439](#)
- [Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 1442](#)

## Using External AAA Authentication Services with DHCP

The extended DHCP local server, including DHCPv6 local server, and the extended DHCP relay agent, including DHCPv6 relay agent, support the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients. When the extended DHCP local server or relay agent receives a discover PDU from a client, the extended DHCP application contacts the AAA server to authenticate the DHCP client. The extended DHCP application can obtain client addresses and DHCP configuration options from the external AAA authentication server.



**NOTE:** This section uses the term *extended DHCP application* to refer to both the extended DHCP local server and the extended DHCP relay agent.

The external authentication feature also supports AAA directed logout. If the external AAA service supports a user logout directive, the extended DHCP application honors the logout and responds as though it were requested by a CLI management command. All of the client state information and allocated resources are deleted at logout. The extended DHCP application supports directed logout using the list of configured authentication servers you specify with the **authentication-server** statement at the **[edit access profile profile-name]** hierarchy level.

You can configure either global authentication support or group-specific support.

You must configure the **username-include** statement to enable the use of authentication. The **password** statement is not required and does not cause DHCP to use authentication if the **username-include** statement is not included.

To configure DHCP local server and DHCP relay agent authentication support:

1. Specify that you want to configure authentication options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit authentication
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit authentication
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit authentication
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit authentication
```

2. (Optional) Configure a password that authenticates the username to the external authentication service.

See “Configuring Passwords for Usernames” on page 1439.

3. (Optional) Configure optional features to create a unique username.

See “Creating Unique Usernames for DHCP Clients” on page 1439.

#### Related Documentation

- [Extended DHCP Local Server Overview on page 1356](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)
- [DHCPv6 Local Server Overview on page 1361](#)
- [DHCPv6 Relay Agent Overview on page 1384](#)

## Grouping Interfaces with Common DHCP Configurations

You use the group feature to group a set of interfaces and then apply a common DHCP configuration to the named interface group. The extended DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent all support interface groups.

The following steps create a DHCP local server group; the steps are similar for the DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent.

To configure a DHCP local server interface group:

1. Specify that you want to configure DHCP local server.

```
[edit system services]
user@host# edit dhcp-local-server
```

2. Create the group and assign a name.

```
[edit system services dhcp-local-server]
user@host# edit group boston
```

3. Specify the names of one or more interfaces on which the extended DHCP application is enabled. You can repeat the `interface interface-name` statement to specify multiple interfaces within the group, but you cannot use the same interface in more than one group.

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.1
user@host# set interface fe-1/0/1.2
```

4. (Optional) You can use the `upto` option to specify a range of interfaces for a group.

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.3 upto fe-1/0/1.9
```

5. (Optional) You can use the `exclude` option to exclude a specific interface or a specified range of interfaces from the group. For example:

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.1 upto fe-1/0/1.102
user@host# set interface fe-1/0/1.6 exclude
user@host# set interface fe-1/0/1.70 upto fe-1/0/1.80 exclude
```

**Related  
Documentation**

- [Extended DHCP Local Server Overview on page 1356](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)
- [DHCPv6 Local Server Overview on page 1361](#)
- [DHCPv6 Relay Agent Overview on page 1384](#)
- [Configuring Group-Specific DHCP Local Server Options on page 1364](#)
- [Configuring Group-Specific DHCP Relay Options on page 1384](#)
- [Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces on page 1422](#)

## Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces

This topic describes guidelines to consider when configuring interface ranges for named interface groups for DHCP local server and DHCP relay. The guidelines refer to the following configuration statement:

```
user@host# set interface interface-name upto upto-interface-name
```

- The start subunit, **interface *interface-name***, serves as the key for the stanza. The remaining configuration settings are considered attributes.
- If the subunit is not included, an implicit **.0** subunit is enforced. The implicit subunit is applied to all interfaces when autoconfiguration is enabled. For example, **interface ge-2/2/2** is treated as **interface ge-2/2/2.0**.
- Ranged entries contain the **upto** option, and the configuration applies to all interfaces within the specified range. The start of a ranged entry must be less than the end of the range. Discrete entries apply to a single interface, except in the case of autoconfiguration, in which a **0** (zero) subunit acts as a wildcard.
- Interface stanzas defined within the same router or switch context are dependent and can constrain each other—both DHCP local server and DHCP relay are considered. Interface stanzas defined across different router (switch) contexts are independent and do not constrain one another.
- Each interface stanza, whether discrete or ranged, has a unique start subunit across a given router context. For example, the following configuration is not allowed within the same group because **ge-1/0/0.10** is the start subunit for both.

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.10
```

- Two groups cannot share interface space. For example, the following configuration is not allowed because the three stanzas share the same space and interfere with one another—interface **ge-1/0/0.26** is common to all three.

```
dhcp-relay group diamond interface ge-1/0/0.10 upto ge-1/0/0.30
dhcp-local-server group ruby interface ge-1/0/0.26
```

```
dhcp-relay group sapphire interface ge-1/0/0.25 upto ge-1/0/0.35
```

- Two ranges cannot overlap, either within a group or across groups. Overlapping occurs when two interface ranges share common subunit space but neither range is a proper subset of the other. The following ranges overlap:

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.20 upto ge-1/0/0.40
```

- A range can contain multiple nested ranges. A nested range is a proper subset of another range. When ranges are nested, the smallest matching range applies.

In the following example, the three ranges nest properly:

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.12 upto ge-1/0/0.15 exclude
interface ge-1/0/0.25 upto ge-1/0/0.29 exclude
```

- Discrete interfaces take precedence over ranges. In the following example, interface `ge-1/0/0.20` takes precedence and enforces an interface client limit of 5.

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.15 upto ge-1/0/0.25 exclude
interface ge-1/0/0.20 overrides interface-client-limit 5
```

#### Related Documentation

- [Grouping Interfaces with Common DHCP Configurations on page 1421](#)

## Overriding Default DHCP Local Server Configuration Settings

Subscriber management enables you to override certain default DHCP and DHCPv6 local server configuration settings. You can override settings at the global level, for a named group of interfaces, or for a specific interface within a named group.

- To override global default DHCP local server configuration options, include the **overrides** statement and its subordinate statements at the **[edit system services dhcp-local-server]** or **[edit system services dhcp-local-server dhcpv6]** hierarchy level.
- To override DHCP local server configuration options for a named group of interfaces, include the statements at the **[edit system services dhcp-local-server group group-name]** or **[edit system services dhcp-local-server dhcpv6 group]** hierarchy level.
- To override DHCP local server configuration options for a specific interface within a named group of interfaces, include the statements at the **[edit system services dhcp-local-server group group-name interface]** or **[edit system services dhcp-local-server dhcpv6 group group-name interface]** hierarchy level.

To override default DHCP local server configuration settings:

- Specify that you want to configure override options.

Global override:

```
[edit system services dhcp-local-server]
user@host# edit overrides
```

Group level override:

```
[edit system services dhcp-local-server]
user@host# edit group boston overrides
```

Per-interface override:

```
[edit system services dhcp-local-server]
user@host# edit group boston overrides interface fe-1/0/1.1
```

2. (Optional) Override the maximum number of DHCP clients allowed per interface.  
See [“Specifying the Maximum Number of DHCP Clients Per Interface” on page 1424](#).
3. (Optional) Configure DHCP client auto logout.  
See [“Automatically Logging Out DHCP Clients” on page 1426](#).
4. (Optional) Enable processing of information requests from clients.  
See [“Enabling Processing of Client Information Requests” on page 1427](#).
5. (Optional, DHCPv6 only) Specify a delegated pool name to use for DHCPv6 multiple address assignment.  
See [“Specifying the Delegated Address Pool for IPv6 Prefix Assignment” on page 1428](#).
6. (Optional, DHCPv6 only) Enable DHCPv6 rapid commit support.  
See [“Enabling DHCPv6 Rapid Commit Support” on page 1428](#).
7. (Optional, DHCPv6 only) Specify that DHCPv6 local server return DNS server addresses as IA\_NA or IA\_PD suboptions rather than as a global DHCPv6 option..  
*See Overriding How the DNS Server Address Is Returned in a DHCPv6 Multiple Address Environment.*
8. (Optional) Delete DHCP override settings.  
See [“Deleting DHCP Local Server and DHCP Relay Override Settings” on page 1429](#).

**Related  
Documentation**

- [Configuring Group-Specific DHCP Local Server Options on page 1364](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)

## Specifying the Maximum Number of DHCP Clients Per Interface

By default, there is no limit to the number of DHCP local server or DHCP relay clients allowed on an interface. However, you can override the default setting and specify the maximum number of clients allowed per interface, in the range 1 through 500,000. When the number of clients on the interface reaches the specified limit, no additional DHCP Discover PDUs or DHCPv6 Solicit PDUs are accepted. When the number of clients subsequently drops below the limit, new clients are again accepted.





**NOTE:** The maximum number of DHCP (and DHCPv6) local server clients or DHCP (and DHCPv6) relay clients can also be specified by Juniper Networks VSA 26-143 during client login. The VSA-specified value always takes precedence if the `interface-client-limit` statement specifies a different number.

If the VSA-specified value differs with each client login, DHCP uses the largest limit set by the VSA until there are no clients on the interface.

To configure the maximum number of DHCP clients allowed per interface:

1. Specify that you want to configure override options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit overrides
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit overrides
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit overrides
```

2. Configure the maximum number of clients allowed per interface. (DHCP local server, DHCPv6 local server, DHCP relay agent and DHCPv6 relay agent all support the `interface-client-limit` statement.)

```
[edit system services dhcp-local-server overrides]
user@host# set interface-client-limit number
```



**NOTE:** For DHCP local server and DHCP relay agent, you can use either the `interface-client-limit` statement or the `client-discover-match incoming-interface` statement to set a limit of one client per interface. The `interface-client-limit` statement with a value of 1 retains the existing client and rejects any new client connections. The `client-discover-match incoming-interface` statement deletes the existing client and allows a new client to connect.

#### Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)
- [Allowing Only One DHCP Client Per Interface](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)
- [Extended DHCP Local Server Overview on page 1356](#)

- [Extended DHCP Relay Agent Overview on page 1379](#)

## Automatically Logging Out DHCP Clients

You can configure the extended DHCP local server and extended DHCP relay to automatically log out DHCP clients. Auto logout immediately releases an existing client when DHCP receives a discover packet from a client whose identity matches an existing client. DHCP then releases the existing client IP address without waiting for the normal lease expiration.



**NOTE:** When the existing client is released, the new client undergoes the normal authentication process. The new client might not receive the same IP address as the original client.

To configure DHCP client auto logout:

1. Specify that you want to configure override options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit overrides
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Enable auto logout and specify the secondary identification method you want to use when the primary identification method is unsuccessful.

- For example, to configure DHCP local server to use the incoming interface method:

```
[edit system services dhcp-local-server overrides]
user@host# set client-discover-match incoming-interface
```

- For example, to configure DHCP relay agent to use the option 60 and option 82 method:

```
[edit forwarding-options dhcp-relay overrides]
user@host# set client-discover-match option60-and-option82
```



**NOTE:** If you change the auto logout configuration, existing clients continue to use the auto logout setting that was configured when they logged in. New clients use the new setting.

### Related Documentation

- [DHCP Auto Logout Overview on page 1368](#)
- [How DHCP Relay Agent Uses Option 82 for Auto Logout on page 1454](#)
- [Allowing Only One DHCP Client Per Interface](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)

- [Extended DHCP Local Server Overview on page 1356](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)

## Enabling Processing of Client Information Requests

By default, DHCP local server and DHCPv6 local server do not respond to information request messages from the client. You can enable DHCP local server and DHCPv6 local server to process these messages and respond to them with an acknowledgment (ack or reply message, respectively) and the requested information.

DHCP relay agent automatically forwards the information request messages without modification to the configured server group by means of the interfaces configured for the respective server group. The messages are dropped if they are received on an unconfigured interface. DHCP relay proxy also supports forwarding these messages. You cannot disable forwarding of the information request messages.

Configure one or more local address pools if you want to use a local pool rather than one provided by AAA. See *Configuring an Address-Assignment Pool Name and Addresses*. For processing information request messages, the address configuration is not necessary. For DHCP local server, you must specify the IPv4 family; for DHCPv6 local server, you must specify the IPv6 family.

See *Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address* for details about how to configure the information sought by clients that send information request messages.

To enable processing of DHCP client information request messages:

1. Specify that you want to configure override options.
  - For DHCP local server:
 

```
[edit system services dhcp-local-server overrides]
user@host# set process-inform
```
  - For DHCPv6 local server:
 

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set process-inform
```
2. (Optional) Specify a pool name from which DHCP information is returned to the client.
  - For DHCP local server:
 

```
[edit system services dhcp-local-server overrides process-inform]
user@host# set pool pool-name
```
  - For DHCPv6 local server:
 

```
[edit system services dhcp-local-server dhcpv6 overrides process-inform]
user@host# set pool pool-name
```

### Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)

- [Extended DHCP Local Server Overview on page 1356](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)

## Specifying the Delegated Address Pool for IPv6 Prefix Assignment

You can explicitly specify a delegated address pool:

- On routers—Subscriber management uses the pool to assign IPv6 prefixes for subscribers. You can specify the delegated address pool globally, for a specific group of interfaces, or for a particular interface.
- On switches—DHCP management uses the pool to assign IPv6 prefixes for DHCP clients. You can specify the delegated address pool globally, for a specific group of interfaces, or for a particular interface.



**NOTE:** You can also use by Juniper Networks VSA 26-161 to specify the delegated address pool. The VSA-specified value always takes precedence over the **delegated-address** statement.

To configure the delegated address pool for DHCPv6 local server:

1. Specify that you want to configure override options.

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit overrides
```

2. Configure the delegated address pool.

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set delegated-pool paris-cable-12
```

### Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)
- [Extended DHCP Local Server Overview on page 1356](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)

## Enabling DHCPv6 Rapid Commit Support

You can configure the extended DHCPv6 local server to support the DHCPv6 Rapid Commit option (DHCPv6 option 14). When rapid commit is enabled on the extended DHCPv6 local server, the server recognizes the Rapid Commit option in Solicit messages sent from the DHCPv6 client. (DHCPv6 clients are configured separately to include the DHCPv6 Rapid Commit option in the Solicit messages.) The server and client then use a two-message exchange (Solicit and Reply) to configure clients, rather than the default four-method exchange (Solicit, Advertise, Request, and Reply). The two-message exchange provides faster client configuration, and is beneficial in environments in which networks are under a heavy load.

You can configure the DHCPv6 local server to support the Rapid Commit option globally, for a specific group, or for a specific interface. By default, rapid commit support is disabled on the DHCPv6 local server.

To configure the DHCPv6 local server to support the DHCPv6 Rapid Commit option:

1. Specify that you want to configure the **overrides** options.

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit overrides
```

2. Enable rapid commit support.

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set rapid-commit
```

#### Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)
- [Extended DHCP Local Server Overview on page 1356](#)

## Deleting DHCP Local Server and DHCP Relay Override Settings

You can delete override settings for DHCP local server and DHCP relay globally, for a named group, or for a specific interface within a named group. You can delete a specific override setting or all overrides.

- To delete a specific DHCP override setting at a particular hierarchy level, include the **overrides** statement with the appropriate subordinate statements. For example, to delete the DHCP local server override **interface-client-limit** setting for a group named **marin20**:

```
[edit system services dhcp-local-server]
user@host# delete group marin20 overrides interface-client-limit
```

- To delete all DHCP override settings at a hierarchy level, include the **overrides** statement without any subordinate statements. For example, to delete all DHCP relay overrides for interface **fxp0.0**, which is in group **marin20**:

```
[edit forwarding-options dhcp-relay]
user@host# delete group marin20 interface fxp0.0 overrides
```

#### Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)
- [Extended DHCP Local Server Overview on page 1356](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)

## Configuring Dynamic Client Reconfiguration of Extended Local Server Clients

The DHCP local server can initiate reconfiguration of its clients to avoid extended outages because of server configuration changes. You can enable dynamic reconfiguration for all DHCP clients or only the DHCP clients serviced by a specified group of interfaces, and you can modify the behavior accordingly.

You can modify the behavior of the reconfiguration process by including the appropriate statements at the **[edit system services dhcp-local-server reconfigure]** hierarchy level for all DHCPv4 clients, and at the **[edit system services dhcp-local-server dhcpv6 reconfigure]** hierarchy level for all DHCPv6 clients. To override this global configuration for only the DHCP clients serviced by a specified group of interfaces, you can include the statements with different values at the **[edit system services dhcp-local-server group *group-name* reconfigure]** hierarchy level for DHCPv4 clients, and at the **[edit system services dhcp-local-server dhcpv6 group *group-name* reconfigure]** hierarchy level for DHCPv6 clients.

To configure dynamic reconfiguration of DHCP clients:

1. Enable dynamic reconfiguration with default values for all clients.

For DHCPv4:

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

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6]
user@host# set reconfigure
```

2. (Optional) Enable dynamic reconfiguration for only the DHCP clients serviced by a group of interfaces.

For DHCPv4:

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

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 group group-name]
user@host# set reconfigure
```

3. (Optional) Configure an authentication token. The DHCP local server then includes this token inside the authentication option when it sends forcerenew or reconfigure messages. If the service provider has previously configured the DHCP client with this token, then the client can compare that token against the newly received token, and reject the message if the tokens do not match. This functionality corresponds to RFC 3118, *Authentication for DHCP Messages*, section 4.

- a. For all clients:

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set token token-value
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set token token-value
```

- b. For only the DHCP clients serviced by a group of interfaces:

For DHCPv4:

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

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
user@host# set token token-value
```

4. For the DHCPv6 server only, you can include the **strict** statement. By default, the server accepts solicit messages from clients that do not support server-initiated reconfiguration. Including this statement causes the server to discard solicit messages from nonsupporting clients; consequently the server does not bind these clients.

For all DHCPv6 clients:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set strict
```

For only the DHCPv6 clients serviced by a group of interfaces:

```
[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]
user@host# set strict
```

5. (Optional) Configure how the server attempts reconfiguration.  
See [“Configuring Dynamic Reconfiguration Attempts for DHCP Clients” on page 1432](#).
6. (Optional) Configure the response to a failed reconfiguration.  
See [“Configuring Deletion of the Client When Dynamic Reconfiguration Fails” on page 1433](#).
7. (Optional) Configure the behavior in response to a RADIUS-initiated disconnect.  
See [“Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect” on page 1433](#).
8. (Optional) Configure a token for rudimentary server authentication.  
See [“Configuring a Token for DHCP Local Server Authentication” on page 1434](#).
9. (Optional) Initiate reconfiguration of some or all client bindings.  
See [“Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings” on page 1435](#).
10. (Optional) Prevent DHCPv6 clients from binding if they do not support reconfigure messages.  
See [“Preventing Binding of Clients That Do Not Support Reconfigure Messages” on page 1434](#).

## Configuring Dynamic Reconfiguration Attempts for DHCP Clients

You can configure how many attempts the local server makes to initiate reconfiguration of the DHCP client by sending `forcerenew` or `reconfigure` messages. You can also specify how long the server waits between attempts. By default, eight attempts are made and the initial interval is two seconds.

Each successive attempt doubles the interval between attempts. For example, if the first value is 2, the first retry is attempted 2 seconds after the first attempt fails. The second retry is attempted 4 seconds after the first retry fails. The third retry is attempted 8 seconds after the second retry fails, and so on. A group configuration takes precedence over a DHCP local server configuration.

(Optional) To configure DHCP local server reconfiguration behavior for all DHCP clients:

1. Specify the number of reconfiguration attempts.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set attempts 5
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set attempts 5
```

2. Specify the interval between reconfiguration attempts.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set timeout 8
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set timeout 8
```

To override the global configuration for a particular group of clients, include the statements at the `[edit system services dhcp-local-server group group-name reconfigure]` hierarchy level or the `[edit system services dhcpv6 dhcp-local-server group group-name reconfigure]` hierarchy level.

### Related Documentation

- [Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429](#)



## Configuring Deletion of the Client When Dynamic Reconfiguration Fails

You can configure the local server to delete the client when the maximum number of reconfiguration attempts has been made without success. By default, the client's original configuration is restored.

(Optional) To configure the DHCP local server to delete the client when reconfiguration is not successful, for all clients:

- Specify the client deletion.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set clear-on-abort
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set clear-on-abort
```

To override the global configuration for a particular group of clients, include the statement at the `[edit system services dhcp-local-server group group-name reconfigure]` hierarchy level or the `[edit system services dhcpv6 dhcp-local-server group group-name reconfigure]` hierarchy level.

### Related Documentation

- [Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429](#)
- [clear-on-abort on page 1481](#)

## Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect

You can configure the local server to reconfigure the client when the client receives a RADIUS-initiated disconnect. By default, the client is deleted when a RADIUS-initiated disconnect is received.

(Optional) To configure the DHCP local server to reconfigure the client instead of deleting the client when a RADIUS-initiated disconnect is received, for all clients:

- Specify the RADIUS-initiated disconnect trigger.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure trigger]
user@host# set radius-disconnect
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure trigger]
user@host# set radius-disconnect
```

To override the global configuration for a particular group of clients, include the statement at the `[edit system services dhcp-local-server group group-name reconfigure trigger]` hierarchy level or the `[edit system services dhcpv6 dhcp-local-server group group-name reconfigure trigger]` hierarchy level.

- Related Documentation**
- [Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429](#)
  - [radius-disconnect on page 1527](#)
  - [trigger on page 1544](#)

## Configuring a Token for DHCP Local Server Authentication

You can configure the local server to include a constant, unencoded token in the DHCP forcerenew message as part of the authentication option it sends to clients. The client compares the received token with a token already configured on the client. If the tokens do not match, the DHCP client discards the forcerenew message. Use of the token provides rudimentary protection against inadvertently instantiated DHCP servers.

(Optional) To configure the DHCP local server to include a token in the forcerenew message sent to the client, for all clients:

- Specify the token.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set token 8ysIU9E32k8r
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set token 8ysIU9E32k8r
```

To override the global configuration for a particular group of clients, include the statement at the `[edit system services dhcp-local-server group group-name reconfigure]` hierarchy level or the `[edit system services dhcpv6 dhcp-local-server group group-name reconfigure]` hierarchy level.

- Related Documentation**
- [Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429](#)
  - [token on page 1539](#)

## Preventing Binding of Clients That Do Not Support Reconfigure Messages

The DHCPv6 client and server negotiate the use of reconfigure messages. When the client can accept reconfigure messages from the server, then the client includes the Reconfigure Accept option in both solicit and request messages sent to the server.

By default, the DHCPv6 server accepts solicit messages from clients regardless of whether they support reconfiguration. You can specify that the server require clients to accept reconfigure messages. In this case, the DHCPv6 server includes the Reconfigure Accept option in both advertise and reply messages when reconfiguration is configured for the client interface. Solicit messages from nonsupporting clients are discarded and the clients are not allowed to bind.

(Optional) To configure the DHCPv6 local server to require that all clients accept reconfiguration:

- Specify strict reconfiguration.

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set strict
```

To override the global configuration for a particular group of clients, include the statement at the `[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]` hierarchy level.

The `show dhcpv6 server statistics` command displays a count of solicit messages that the server has discarded.

#### Related Documentation

- [Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429](#)
- [strict on page 1535](#)

## Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings

You can request that the DHCP local server initiate reconfiguration of all of clients or only specified clients.

To request reconfiguration of all clients:

- Specify the `all` option.

For DHCPv4:

```
user@host> request dhcp server reconfigure all
```

For DHCPv6:

```
user@host> request dhcpv6 server reconfigure all
```

You can use any of the following methods to request reconfiguration of specific clients:

- Specify the IP address of the DHCP client.

For DHCPv4:

```
user@host> request dhcp server reconfigure 192.168.27.3
```

For DHCPv6:

```
user@host> request dhcpv6 server reconfigure 2001:bd8:1111:2222::
```

- Specify the client ID of a DHCPv6 client.

```
user@host> request dhcpv6 server reconfigure
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
```

- Specify the session ID of a DHCPv6 client.

```
user@host> request dhcpv6 server reconfigure 5
```

- Specify the MAC address of a DHCPv4 client.

```
user@host> request dhcp server reconfigure 12:23:34:45:56:67
```

- Specify an interface; reconfiguration is attempted for all clients on this interface.

```
user@host> request dhcp server reconfigure interface fe-0/0/0.100
```

- Specify a logical system; reconfiguration is attempted for all clients or the specified clients in this logical system.

```
user@host> request dhcp server reconfigure all logical-system ls-bldg5
```

- Specify a routing instance; reconfiguration is attempted for all clients or the specified clients in this routing instance.

```
user@host> request dhcp server reconfigure all routing-instance ri-boston
```

**Related  
Documentation**

- [Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429](#)
- [request dhcp server reconfigure on page 1707](#)

## Configuring Detection of DHCP Local Server Client Connectivity

Liveness detection for DHCP subscriber IP sessions or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.



**NOTE:** You can also configure DHCP liveness detection for DHCP relay.

To configure liveness detection for DHCP local server:

1. Specify that you want to configure liveness detection.

- For DHCP global configuration:

```
[edit system services dhcp-local-server]  
user@host# edit liveness-detection
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name]  
user@host# edit liveness-detection
```



**NOTE:** Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the [liveness-detection](#) statement, and any subsequent configuration statements, at the `[edit system services dhcp-local-server dhcpv6]` or `[edit system services dhcp-local-server dhcpv6 group group-name]` hierarchy level.

2. Specify that you want to configure the liveness detection method.

- For DHCP global configuration:

```
[edit system services dhcp-local-server liveness-detection]
user@host# edit method
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection]
user@host# edit method
```

3. Specify the liveness detection method that you want DHCP to use.



**NOTE:** The only method supported for liveness detection is Bidirectional Forwarding Detection (BFD).

- For DHCP global configuration:

```
[edit system services dhcp-local-server liveness-detection method]
user@host# edit bfd
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection method]
user@host# edit bfd
```

4. Configure the liveness detection method as desired.

See [“Example: Configuring Group Liveness Detection for DHCP Local Server Clients” on page 1395](#) for an example of how to configure DHCPv4 groups for DHCP local server liveness detection.

5. Configure the action the router takes when a liveness detection failure occurs.

- For DHCP global configuration:

```
[edit system services dhcp-local-server liveness-detection]
user@host# edit failure-action action
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection]
user@host# edit failure-action action
```

#### Related Documentation

- [DHCP Liveness Detection Overview on page 1472](#)
- [Extended DHCP Local Server Overview on page 1356](#)
- [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 1473](#)
- [Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395](#)
- [Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients](#)

## Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server

You can configure how DHCP local server handles DHCP snooped packets. Depending on the configuration, DHCP local server either forwards or drops the snooped packets it receives.

Table 139 on page 1438 indicates the action the router takes for DHCP local server snooped packets.



**NOTE:** Configured interfaces are those interfaces that have been configured with the `group` statement in the `[edit system services dhcp-local-server]` hierarchy. Non-configured interfaces are those that are in the logical system/routing instance but have not been configured by the `group` statement.

**Table 139: Actions for DHCP Local Server Snooped Packets**

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
<code>forward-snooped-clients</code> not configured	dropped	dropped
<code>all-interfaces</code>	forwarded	forwarded
<code>configured-interfaces</code>	forwarded	dropped
<code>non-configured-interfaces</code>	dropped	forwarded

To configure DHCP snooped packet forwarding for DHCP local server:

1. Specify that you want to configure DHCP local server.

```
[edit]
user@host# edit system services dhcp-local-server
```

2. Enable DHCP snooped packet forwarding for DHCP local server.

```
[edit system services dhcp-local-server]
user@host# edit forward-snooped-clients
```

3. Specify the interfaces that are supported for snooped packet forwarding.

```
[edit system services dhcp-local-server forward-snooped-clients]
user@host# set (all-interfaces | configured-interfaces | non-configured-interfaces)
```

For example, to configure DHCP local server to forward DHCP snooped packets on only configured interfaces:

```
[edit]
system {
  services {
    dhcp-local-server {
      forward-snooped-clients configured-interfaces;
    }
  }
}
```

**Related Documentation**

- [DHCP Snooping Support on page 1367](#)

## Configuring Passwords for Usernames

You can configure an optional password that the extended DHCP application presents to the external AAA authentication service to authenticate the specified username.

To configure a password that authenticates the username:

1. Specify that you want to configure authentication options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit authentication
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit authentication
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit authentication
```

2. Configure the password. (DHCP local server, DHCPv6 local server, and DHCP relay agent all support the **password** statement.)

```
[edit system services dhcp-local-server authentication]
user@host# set password myPassword1234
```

### Related Documentation

- [Extended DHCP Local Server Overview on page 1356](#)
- [DHCPv6 Local Server Overview on page 1361](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)
- [Using External AAA Authentication Services with DHCP on page 1420](#)
- [Special Requirements for Junos OS Plain-Text Passwords](#)

## Creating Unique Usernames for DHCP Clients

You can configure the extended DHCP application to include additional information in the username that is passed to the external AAA authentication service when the DHCP client logs in. This additional information enables you to construct usernames that uniquely identify subscribers (DHCP clients).



**NOTE:** If you do not include a username in the authentication configuration, the router (or switch) does not perform authentication; however, the IP address is provided by the local pool if it is configured.

When you use the DHCPv6 local server, you must configure authentication and the client username; otherwise client login fails.

The following list describes the optional information that you can include as part of the username:

- **circuit-type**—The circuit type used by the DHCP client, for example **enet**.
- **client-id**—The client identifier option (option 1). (DHCPv6 local server DHCPv6 relay agent only)
- **delimiter**—The delimiter character that separates components that make up the concatenated username. The default delimiter is a period (.). The semicolon (;) is not supported as a delimiter character.
- **domain-name**—The client domain name as a string. The router adds the @ delimiter to the username.
- **interface-name**—The interface name, including the interface device and associated VLAN IDs.
- **logical-system-name**—The name of the logical system, if the receiving interface is in a logical system.
- **mac-address**—The client MAC address, in a string of the format *xxxx.xxxx.xxxx*. (Not supported for DHCPv6 local server)
- **option-60**—The portion of the option 60 payload that follows the length field. (Not supported for DHCPv6 local server)
- **option-82 <circuit-id> <remote-id>**—The specified contents of the option 82 payload. (Not supported for DHCPv6 local server)
  - **circuit-id**—The payload of the Agent Circuit ID suboption.
  - **remote-id**—The payload of the Agent Remote ID suboption.
  - Both **circuit-id** and **remote-id**—The payloads of both suboptions, in the format: **circuit-id[delimiter]remote-id**.
  - Neither **circuit-id** or **remote-id**—The raw payload of the option 82 from the PDU is concatenated to the username.



**NOTE:** For DHCP relay agent, the option 82 value used in creating the username is based on the option 82 value that is encoded in the outgoing (relayed) PDU.

---

- **relay-agent-interface-id**—The Interface-ID option (option 18). (DHCPv6 local server or relay agent)
- **relay-agent-remote-id**—The DHCPv6 Relay Agent Remote-ID option (option 37). (DHCPv6 local server or relay agent)
- **relay-agent-subscriber-id**—(On routers only) The DHCPv6 Relay Agent Subscriber-ID option (option 38). (DHCPv6 local server or relay agent)



- **routing-instance-name**—The name of the routing instance, if the receiving interface is in a routing instance.
- **user-prefix**—A string indicating the user prefix.

The router (switch) creates the unique username by including the specified additional information in the following order, with the fields separated by a delimiter.

For DHCP local server and DHCP relay agent:

```
user-prefix[delimiter]mac-address[delimiter]logical-system-name[delimiter]
routing-instance-name[delimiter]circuit-type[delimiter]interface-name[delimiter]option-82[delimiter]
option-60@domain-name
```

For DHCPv6 local server:

```
user-prefix[delimiter]logical-system-name[delimiter]routing-instance-name[delimiter]
circuit-type[delimiter]interface-name[delimiter]relay-agent-remote-id[delimiter]
relay-agent-subscriber-id[delimiter]relay-agent-interface-id[delimiter]client-id@domain-name
```

To configure a unique username:

1. Specify that you want to configure authentication.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit authentication
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit authentication
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit authentication
```

2. Specify that you want to include optional information in the username. (DHCP local server, DHCPv6 local server, and DHCP relay agent all support the **username-include** statement.)

```
[edit system services dhcp-local-server authentication]
user@host# set username-include
```

3. (Optional) Specify the optional information you want to include in the username.

```
[edit system services dhcp-local-server authentication username-include]
user@host# set username-include circuit-type
user@host# set username-include domain-name isp55.com
user@host# set username-include mac-address
user@host# set username-include user-prefix wallybrown
```

The previous **username-include** configuration produces this unique username:

```
wallybrown.0090.1a01.1234.enet@isp55.com
```

#### Related Documentation

- [Extended DHCP Local Server Overview on page 1356](#)
- [DHCPv6 Local Server Overview on page 1361](#)

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [Using External AAA Authentication Services with DHCP on page 1420](#)

## Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use

You can specify the match order in which the extended DHCP local server uses the client data to determine the address-assignment pool that provides the IP address and configuration for a DHCP client. You use the **pool-match-order** statement to specify the match order. If you do not specify the **pool-match-order**, the router (or switch) uses the default **ip-address-first** matching to select the address pool. After DHCP local server determines the address assignment pool to use, the server performs the matching based on the criteria you specified in the pool configuration.

In the default **ip-address-first** matching, the server selects the address-assignment pool to use by matching the IP address in the client DHCP request with the network address of the address-assignment pool. If the client request contains the gateway IP address (giaddr), the local server matches the giaddr to the address-assignment pool's address. If there is no giaddr in the request, then the DHCP local server matches the IP address of the receiving interface to the address of the address-assignment pool.

In **external-authority** matching, the DHCP local server receives the address assignment from an external authority, such as RADIUS or Diameter. If RADIUS is the external authority, the DHCP local server uses the Framed-IPv6-Pool attribute (RADIUS attribute 100) to select the pool. If Diameter is the external authority, the server uses the Diameter counterpart of the Framed-IPv6-Pool attribute to determine the pool.

For IPv4 address-assignment pools, you can optionally configure the extended DHCP local server to match the DHCP relay agent information option (option 82) in the client DHCP packets to a named range in the address-assignment pool used for the client. Named ranges are subsets within the overall address-assignment pool address range, which you can configure when you create the address-assignment pool.



**NOTE:** To use the DHCP local server option 82 matching feature with an IPv4 address-assignment pool, you must ensure that the **option-82** statement is included in the **dhcp-attributes** statement for the address-assignment pool.

---

To configure the matching order the extended DHCP local server uses to determine the address-assignment pool used for a client:

1. Access the **pool-match-order** configuration.

```
[edit system services dhcp-local-server]  
user@host# edit pool-match-order
```

2. Specify the pool matching methods in the order in which the router (switch) performs the methods. You can specify the methods in any order. All methods are optional—the router (switch) uses the **ip-address-first** method by default.

- Configure the router (switch) to use an external addressing authority.

```
[edit system services dhcp-local-server pool-match-order]
user@host# set external-authority
```

- Configure the router (switch) to use the ip-address-first method.

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

- (IPv4 address-assignment pools only) Specify the option 82 matching method.

```
[edit system services dhcp-local-server pool-match-order]
user@host# set option-82
```

#### Related Documentation

- [Address-Assignment Pools Overview on page 1370](#)
- [Configuring Address-Assignment Pools](#)
- [Extended DHCP Local Server Overview on page 1356](#)
- [Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 1394](#)

## Configuration Tasks for DHCP Relay Agent

- [Using External AAA Authentication Services with DHCP on page 1444](#)
- [Grouping Interfaces with Common DHCP Configurations on page 1445](#)
- [Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces on page 1446](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)
- [Changing the Gateway IP Address \(giaddr\) Field to the giaddr of the DHCP Relay Agent on page 1449](#)
- [Replacing the DHCP Relay Request and Release Packet Source Address on page 1450](#)
- [Overriding Option 82 Information on page 1450](#)
- [Using Layer 2 Unicast Transmission for DHCP Packets on page 1451](#)
- [Trusting Option 82 Information on page 1451](#)
- [Specifying the Maximum Number of DHCP Clients Per Interface on page 1451](#)
- [Automatically Logging Out DHCP Clients on page 1453](#)
- [How DHCP Relay Agent Uses Option 82 for Auto Logout on page 1454](#)
- [Configuring DHCP Snooping for DHCP Relay Agent on page 1455](#)
- [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 1455](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 1460](#)

- [Sending Release Messages When Clients Are Deleted on page 1462](#)
- [Disabling Automatic Binding of Stray DHCP Requests on page 1463](#)
- [Using DHCP Relay Agent Option 82 Information on page 1464](#)
- [Configuring Server Groups on page 1469](#)
- [Configuring Active Server Groups on page 1469](#)
- [Enabling DHCP Relay Proxy Mode on page 1470](#)
- [Inserting DHCPv6 Interface-ID Option \(Option 18\) In DHCPv6 Packets on page 1470](#)
- [DHCP Liveness Detection Overview on page 1472](#)
- [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 1473](#)
- [Disabling DHCP Relay on page 1474](#)

## Using External AAA Authentication Services with DHCP

The extended DHCP local server, including DHCPv6 local server, and the extended DHCP relay agent, including DHCPv6 relay agent, support the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients. When the extended DHCP local server or relay agent receives a discover PDU from a client, the extended DHCP application contacts the AAA server to authenticate the DHCP client. The extended DHCP application can obtain client addresses and DHCP configuration options from the external AAA authentication server.



**NOTE:** This section uses the term *extended DHCP application* to refer to both the extended DHCP local server and the extended DHCP relay agent.

The external authentication feature also supports AAA directed logout. If the external AAA service supports a user logout directive, the extended DHCP application honors the logout and responds as though it were requested by a CLI management command. All of the client state information and allocated resources are deleted at logout. The extended DHCP application supports directed logout using the list of configured authentication servers you specify with the **authentication-server** statement at the **[edit access profile profile-name]** hierarchy level.

You can configure either global authentication support or group-specific support.

You must configure the **username-include** statement to enable the use of authentication. The **password** statement is not required and does not cause DHCP to use authentication if the **username-include** statement is not included.

To configure DHCP local server and DHCP relay agent authentication support:

1. Specify that you want to configure authentication options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit authentication
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit authentication
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit authentication
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit authentication
```

2. (Optional) Configure a password that authenticates the username to the external authentication service.

See [“Configuring Passwords for Usernames” on page 1439](#).

3. (Optional) Configure optional features to create a unique username.

See [“Creating Unique Usernames for DHCP Clients” on page 1439](#).

#### Related Documentation

- [Extended DHCP Local Server Overview on page 1356](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)
- [DHCPv6 Local Server Overview on page 1361](#)
- [DHCPv6 Relay Agent Overview on page 1384](#)

## Grouping Interfaces with Common DHCP Configurations

You use the group feature to group a set of interfaces and then apply a common DHCP configuration to the named interface group. The extended DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent all support interface groups.

The following steps create a DHCP local server group; the steps are similar for the DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent.

To configure a DHCP local server interface group:

1. Specify that you want to configure DHCP local server.

```
[edit system services]
user@host# edit dhcp-local-server
```

2. Create the group and assign a name.

```
[edit system services dhcp-local-server]
user@host# edit group boston
```

3. Specify the names of one or more interfaces on which the extended DHCP application is enabled. You can repeat the `interface interface-name` statement to specify multiple interfaces within the group, but you cannot use the same interface in more than one group.

```
[edit system services dhcp-local-server group boston]
```

```
user@host# set interface fe-1/0/1.1
user@host# set interface fe-1/0/1.2
```

4. (Optional) You can use the **upto** option to specify a range of interfaces for a group.

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.3 upto fe-1/0/1.9
```

5. (Optional) You can use the **exclude** option to exclude a specific interface or a specified range of interfaces from the group. For example:

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.1 upto fe-1/0/1.102
user@host# set interface fe-1/0/1.6 exclude
user@host# set interface fe-1/0/1.70 upto fe-1/0/1.80 exclude
```

#### Related Documentation

- [Extended DHCP Local Server Overview on page 1356](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)
- [DHCPv6 Local Server Overview on page 1361](#)
- [DHCPv6 Relay Agent Overview on page 1384](#)
- [Configuring Group-Specific DHCP Local Server Options on page 1364](#)
- [Configuring Group-Specific DHCP Relay Options on page 1384](#)
- [Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces on page 1422](#)

## Guidelines for Configuring Interface Ranges for Groups of DHCP Interfaces

This topic describes guidelines to consider when configuring interface ranges for named interface groups for DHCP local server and DHCP relay. The guidelines refer to the following configuration statement:

```
user@host# set interface interface-name upto upto-interface-name
```

- The start subunit, **interface *interface-name***, serves as the key for the stanza. The remaining configuration settings are considered attributes.
- If the subunit is not included, an implicit **.0** subunit is enforced. The implicit subunit is applied to all interfaces when autoconfiguration is enabled. For example, **interface ge-2/2/2** is treated as **interface ge-2/2/2.0**.
- Ranged entries contain the **upto** option, and the configuration applies to all interfaces within the specified range. The start of a ranged entry must be less than the end of the range. Discrete entries apply to a single interface, except in the case of autoconfiguration, in which a **0** (zero) subunit acts as a wildcard.
- Interface stanzas defined within the same router or switch context are dependent and can constrain each other—both DHCP local server and DHCP relay are considered. Interface stanzas defined across different router (switch) contexts are independent and do not constrain one another.

- Each interface stanza, whether discrete or ranged, has a unique start subunit across a given router context. For example, the following configuration is not allowed within the same group because **ge-1/0/0.10** is the start subunit for both.

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.10
```

- Two groups cannot share interface space. For example, the following configuration is not allowed because the three stanzas share the same space and interfere with one another—interface **ge-1/0/0.26** is common to all three.

```
dhcp-relay group diamond interface ge-1/0/0.10 upto ge-1/0/0.30
dhcp-local-server group ruby interface ge-1/0/0.26
dhcp-relay group sapphire interface ge-1/0/0.25 upto ge-1/0/0.35
```

- Two ranges cannot overlap, either within a group or across groups. Overlapping occurs when two interface ranges share common subunit space but neither range is a proper subset of the other. The following ranges overlap:

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.20 upto ge-1/0/0.40
```

- A range can contain multiple nested ranges. A nested range is a proper subset of another range. When ranges are nested, the smallest matching range applies.

In the following example, the three ranges nest properly:

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.12 upto ge-1/0/0.15 exclude
interface ge-1/0/0.25 upto ge-1/0/0.29 exclude
```

- Discrete interfaces take precedence over ranges. In the following example, interface **ge-1/0/0.20** takes precedence and enforces an interface client limit of 5.

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.15 upto ge-1/0/0.25 exclude
interface ge-1/0/0.20 overrides interface-client-limit 5
```

#### Related Documentation

- [Grouping Interfaces with Common DHCP Configurations on page 1421](#)

## Overriding the Default DHCP Relay Configuration Settings

You can override the default DHCP and DHCPv6 relay agent configuration settings at the global level, for a named group of interfaces, or for a specific interface within a named group.

- To override global default DHCP relay agent configuration options, include the **overrides** statement and its subordinate statements at the **[edit forwarding-options dhcp-relay]** hierarchy level.
- To override DHCP relay configuration options for a named group of interfaces, include the statements at the **[edit forwarding-options dhcp-relay group group-name]** hierarchy level.

- To override DHCP relay configuration options for a specific interface within a named group of interfaces, include the statements at the **[edit forwarding-options dhcp-relay group *group-name* interface]** hierarchy level.
- To configure overrides for DHCPv6 relay, use the supported statements at the **[edit forwarding-options dhcp-relay dhcpv6]** hierarchy level.

To override default DHCP relay agent configuration settings:

1. Specify that you want to configure override options.

Global override:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

Group-level override:

```
[edit forwarding-options dhcp-relay]
user@host# edit group boston overrides
```

Per-interface override:

```
[edit forwarding-options dhcp-relay]
user@host# edit group boston interface fe-1/0/1.2 overrides
```

2. (DHCPv4 only) Enable DHCP relay proxy mode.  
See [“Enabling DHCP Relay Proxy Mode” on page 1470](#).
3. (DHCPv4 only) Overwrite the giaddr in DHCP packets that the DHCP relay agent forwards.  
See [“Changing the Gateway IP Address \(giaddr\) Field to the giaddr of the DHCP Relay Agent” on page 1449](#).
4. (DHCPv4 only) Replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).  
See [“Replacing the DHCP Relay Request and Release Packet Source Address” on page 1450](#).
5. (DHCPv4 only) Override the DHCP relay agent information option (option 82) in DHCP packets.  
See [“Overriding Option 82 Information” on page 1450](#).
6. (DHCPv4 only) Override the setting of the broadcast bit in DHCP request packets and use the Layer 2 unicast transmission method.  
See [“Using Layer 2 Unicast Transmission for DHCP Packets” on page 1451](#).
7. (DHCPv4 only) Trust DHCP client packets that have a giaddr of 0 and that contain option 82 information.  
See [“Trusting Option 82 Information” on page 1451](#).
8. (DHCPv4 and DHCPv6) Override the maximum number of DHCP clients allowed per interface.  
See [“Specifying the Maximum Number of DHCP Clients Per Interface” on page 1424](#).



9. (DHCPv4 only) Configure client auto logout.  
See [“DHCP Auto Logout Overview” on page 1368](#).
10. (DHCPv4 and DHCPv6) Enable or disable support for DHCP snooped clients on interfaces.  
See [“Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent” on page 1455](#).
11. (DHCPv4 and DHCPv6) Delay authentication of subscribers until the DHCP client sends a Request packet.  
See the *delay-authentication* statement.
12. (DHCPv4 and DHCPv6) Send release messages to the DHCP server when clients are deleted.  
See [“Sending Release Messages When Clients Are Deleted” on page 1462](#).
13. (DHCPv4 only) Disable the DHCP relay agent on specific interfaces.  
See [“Disabling DHCP Relay” on page 1474](#).
14. (DHCPv4 and DHCPv6) Disable automatic binding of stray DHCP requests.  
See [“Disabling Automatic Binding of Stray DHCP Requests” on page 1463](#).

- Related Documentation**
- [Configuring Group-Specific DHCP Relay Options on page 1384](#)
  - [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)

## Changing the Gateway IP Address (giaddr) Field to the giaddr of the DHCP Relay Agent

You can configure the DHCP relay agent to change the gateway IP address (giaddr) field in packets that it forwards between a DHCP client and a DHCP server.

To overwrite the giaddr of every DHCP packet with the giaddr of the DHCP relay agent before forwarding the packet to the DHCP server:

1. Specify that you want to configure override options.  

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```
2. Specify that the giaddr of DHCP packets is overwritten.  

```
[edit forwarding-options dhcp-relay overrides]
user@host# set always-write-giaddr
```

- Related Documentation**
- [Extended DHCP Relay Agent Overview on page 1379](#)
  - [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)

## Replacing the DHCP Relay Request and Release Packet Source Address

You can configure the DHCP relay agent to replace request and release packets with the gateway IP address (giaddr) before forwarding the packet to the DHCP server.

To replace the source address with giaddr:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Specify that you want to replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).

```
[edit forwarding-options dhcp-relay overrides]
user@host# set replace-ip-source-with giaddr
```

### Related Documentation

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)

## Overriding Option 82 Information

You can configure the DHCP relay agent to add or remove the DHCP relay agent information option (option 82) in DHCP packets.

This feature causes the DHCP relay agent to perform one of the following actions, depending on the configuration:

- If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server.
- If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.

To override the default option 82 information in DHCP packets destined for a DHCP server:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Specify that the option 82 information in DHCP packets is overwritten.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set always-write-option-82
```

### Related Documentation

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)

## Using Layer 2 Unicast Transmission for DHCP Packets

You can configure the DHCP relay agent to override the setting of the broadcast bit in DHCP request packets. DHCP relay agent then instead uses the Layer 2 unicast transmission method to send DHCP Offer reply packets and DHCP ACK reply packets from the DHCP server to DHCP clients during the discovery process.

To override the default setting of the broadcast bit in DHCP request packets:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Specify that the DHCP relay agent uses the Layer 2 unicast transmission method.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set layer2-unicast-replies
```

### Related Documentation

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)

## Trusting Option 82 Information

By default, the DHCP relay agent treats client packets with a giaddr of 0 (zero) and option 82 information as if the packets originated at an untrusted source, and drops them without further processing. You can override this behavior and specify that the DHCP relay agent process DHCP client packets that have a giaddr of 0 (zero) and contain option 82 information.

To configure DHCP relay agent to trust option 82 information:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Specify that the DHCP relay agent process DHCP client packets with a giaddr of 0 and that contain option 82 information.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set trust-option-82
```

### Related Documentation

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)

## Specifying the Maximum Number of DHCP Clients Per Interface

By default, there is no limit to the number of DHCP local server or DHCP relay clients allowed on an interface. However, you can override the default setting and specify the maximum number of clients allowed per interface, in the range 1 through 500,000. When the number of clients on the interface reaches the specified limit, no additional DHCP

Discover PDUs or DHCPv6 Solicit PDUs are accepted. When the number of clients subsequently drops below the limit, new clients are again accepted.



**NOTE:** The maximum number of DHCP (and DHCPv6) local server clients or DHCP (and DHCPv6) relay clients can also be specified by Juniper Networks VSA 26-143 during client login. The VSA-specified value always takes precedence if the `interface-client-limit` statement specifies a different number.

If the VSA-specified value differs with each client login, DHCP uses the largest limit set by the VSA until there are no clients on the interface.

To configure the maximum number of DHCP clients allowed per interface:

1. Specify that you want to configure override options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit overrides
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit overrides
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit overrides
```

2. Configure the maximum number of clients allowed per interface. (DHCP local server, DHCPv6 local server, DHCP relay agent and DHCPv6 relay agent all support the `interface-client-limit` statement.)

```
[edit system services dhcp-local-server overrides]
user@host# set interface-client-limit number
```



**NOTE:** For DHCP local server and DHCP relay agent, you can use either the `interface-client-limit` statement or the `client-discover-match incoming-interface` statement to set a limit of one client per interface. The `interface-client-limit` statement with a value of 1 retains the existing client and rejects any new client connections. The `client-discover-match incoming-interface` statement deletes the existing client and allows a new client to connect.

#### Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)
- [Allowing Only One DHCP Client Per Interface](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)

- [Extended DHCP Local Server Overview on page 1356](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)

## Automatically Logging Out DHCP Clients

You can configure the extended DHCP local server and extended DHCP relay to automatically log out DHCP clients. Auto logout immediately releases an existing client when DHCP receives a discover packet from a client whose identity matches an existing client. DHCP then releases the existing client IP address without waiting for the normal lease expiration.



**NOTE:** When the existing client is released, the new client undergoes the normal authentication process. The new client might not receive the same IP address as the original client.

To configure DHCP client auto logout:

1. Specify that you want to configure override options.
  - For DHCP local server:
 

```
[edit system services dhcp-local-server]
user@host# edit overrides
```
  - For DHCP relay agent:
 

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```
2. Enable auto logout and specify the secondary identification method you want to use when the primary identification method is unsuccessful.
  - For example, to configure DHCP local server to use the incoming interface method:
 

```
[edit system services dhcp-local-server overrides]
user@host# set client-discover-match incoming-interface
```
  - For example, to configure DHCP relay agent to use the option 60 and option 82 method:
 

```
[edit forwarding-options dhcp-relay overrides]
user@host# set client-discover-match option60-and-option82
```



**NOTE:** If you change the auto logout configuration, existing clients continue to use the auto logout setting that was configured when they logged in. New clients use the new setting.

### Related Documentation

- [DHCP Auto Logout Overview on page 1368](#)
- [How DHCP Relay Agent Uses Option 82 for Auto Logout on page 1454](#)
- [Allowing Only One DHCP Client Per Interface](#)

- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)
- [Extended DHCP Local Server Overview on page 1356](#)
- [Extended DHCP Relay Agent Overview on page 1379](#)

## How DHCP Relay Agent Uses Option 82 for Auto Logout

Table 140 on page 1454 indicates how the DHCP relay agent determines the option 82 value used for the client auto logout feature. Depending on the configuration settings, DHCP relay agent takes the action indicated in the right column.

**Table 140: DHCP Relay Agent Option 82 Value for Auto Logout**

DHCP Relay Agent Configuration Settings				gladdr in non-snooped packet	Action Taken
DHCP Relay Configured with Option 82	Discover Packet Contains Option 82	Override "trust-option-82"	Override "always-write-option-82"		
No	No	—	—	—	No secondary search performed
No	Yes	Yes	—	—	Use option 82 from packet
No	Yes	No	—	Zero	Drop packet
No	Yes	No	—	Non-zero	Use option 82 from packet
Yes	No	—	—	—	Use configured option 82
Yes	Yes	No	—	Zero	Drop packet
Yes	Yes	No	No	Non-zero	Use option 82 from packet
Yes	Yes	No	Yes	Non-zero	Overwrite the configured option 82
Yes	Yes	Yes	No	—	Use option 82 from packet
Yes	Yes	Yes	Yes	—	Overwrite the configured option 82

- Related Documentation**
- [DHCP Auto Logout Overview on page 1368](#)
  - [Automatically Logging Out DHCP Clients on page 1426](#)

## Configuring DHCP Snooping for DHCP Relay Agent

DHCP relay agent uses a two-part configuration to determine how to handle DHCP snooped packets. First, you enable or disable snooping support for DHCP relay agent and, optionally, override the default snooping configuration. Then you configure the forwarding action for snooped clients, which specifies whether DHCP relay agent forwards or drops snooped traffic.

To configure DHCP snooping for DHCP relay agent:

1. (DHCPv4 and DHCPv6) Enable or disable DHCP snooping. You can configure DHCP snooping globally, for a named group of interfaces, or for a specific interface.  
See [“Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent” on page 1455](#).
2. (DHCPv4 only) Configure snooped packets forwarding support.  
See [“Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent” on page 1460](#).

- Related Documentation**
- [DHCP Snooping Support on page 1367](#)
  - [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 1455](#)
  - [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 1460](#)

## Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent

DHCP relay agent uses a two-part configuration to determine how to handle DHCP snooped packets. This topic describes the first procedure, in which you enable or disable snooping support for DHCP relay agent and, optionally, override the default snooping configuration.

The second procedure, which applies only to DHCPv4 relay agent, is described in [“Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent” on page 1460](#), and configures the forwarding action for snooped clients, which specifies whether DHCP relay agent forwards or drops snooped traffic.

You can enable or disable DHCP globally for DHCP relay, for a group of interfaces, or for a specific interface in a group.

By default, DHCP snooping is disabled for DHCP relay. To enable or disable DHCP snooping support globally:

1. Specify that you want to configure DHCP relay agent.

- For DHCP relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay
```

- For DHCPv6 relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay dhcpv6
```

2. Specify that you want to override the default configuration.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit overrides
```

3. Enable or disable DHCP snooping support.

- To enable DHCP snooping:

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay overrides]
user@host# set allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 overrides]
user@host# set allow-snooped-clients
```

- To disable DHCP snooping:

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay overrides]
user@host# set no-allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 overrides]
user@host# set no-allow-snooped-clients
```

For example, to enable global DHCP snooping support :

```
forwarding-options {
  dhcp-relay {
    overrides {
      allow-snooped-clients;
    }
  }
}
```



To enable or disable DHCP snooping support for a group of interfaces:

1. Specify that you want to configure DHCP relay agent.

- For DHCP relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay
```

- For DHCPv6 relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay dhcpv6
```

2. Specify the named group.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit group group-name
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit group group-name
```

3. Specify that you want to override the default configuration.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name]
user@host# edit overrides
```

4. Enable or disable DHCP snooping support.

- To enable DHCP snooping:

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name overrides]
user@host# set allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides]
user@host# set allow-snooped-clients
```

- To disable DHCP snooping:

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name overrides]
user@host# set no-allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides]
user@host# set no-allow-snooped-clients
```

For example, to enable DHCP snooping support on all interfaces in group **boston**:

```
forwarding-options {  
  dhcp-relay {  
    group boston {  
      overrides {  
        allow-snooped-clients;  
      }  
    }  
  }  
}
```

To enable or disable DHCP snooping support on a specific interface:

1. Specify that you want to configure DHCP relay agent.

- For DHCP relay agent:

```
[edit]  
user@host# edit forwarding-options dhcp-relay
```

- For DHCPv6 relay agent:

```
[edit]  
user@host# edit forwarding-options dhcp-relay dhcpv6
```

2. Specify the named group containing the interface.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]  
user@host# edit group group-name
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]  
user@host# edit group group-name
```

3. Specify the interface for which you want to configure DHCP snooping.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name]  
user@host# edit interface interface-name
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name]  
user@host# edit interface interface-name
```

4. Specify that you want to override the default configuration on the interface.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name interface interface-name]  
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface  
  interface-name]  
user@host# edit overrides
```

5. Enable or disable DHCP snooping support.

- To enable DHCP snooping:
  - For DHCP relay agent:
 

```
[edit forwarding-options dhcp-relay group group-name interface interface-name
overrides]
user@host# set allow-snooped-clients
```
  - For DHCPv6 relay agent:
 

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface
interface-name overrides]
user@host# set allow-snooped-clients
```
- To disable DHCP snooping:
  - For DHCP relay agent:
 

```
[edit forwarding-options dhcp-relay group group-name interface interface-name
overrides]
user@host# set no-allow-snooped-clients
```
  - For DHCPv6 relay agent:
 

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface
interface-name overrides]
user@host# set no-allow-snooped-clients
```

For example, to disable DHCP snooping support on interface **ge-2/1/8.0** in group **boston**:

```
forwarding-options {
  dhcp-relay {
    group boston {
      interface ge-2/1/8.0 {
        overrides {
          no-allow-snooped-clients;
        }
      }
    }
  }
}
```

To enable DHCPv6 snooping support on interface **ge-3/2/1.1** in group **sunnyvale**:

```
forwarding-options {
  dhcp-relay {
    dhcpv6 {
      group sunnyvale {
        interface ge-3/2/1.1 {
          overrides {
            allow-snooped-clients;
          }
        }
      }
    }
  }
}
```

- Related Documentation**
- [DHCP Snooping Support on page 1367](#)
  - [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 1460](#)
  - [Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 1403](#)
  - [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)

## Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent

You can configure how DHCP relay agent handles DHCP snooped packets. Depending on the configuration, DHCP relay agent either forwards or drops the snooped packets it receives.

DHCP relay uses a two-part configuration to determine how to handle DHCP snooped packets. This topic describes how you use the **forward-snooped-clients** statement to manage whether DHCP relay agent forwards or drops snooped packets, depending on the type of interface on which the packets are snooped. In the other part of the DHCP relay agent snooping configuration, which is described in “[Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent](#)” on page 1455, you enable or disable the DHCP relay snooping feature.

[Table 141 on page 1460](#) shows the action the router or switch takes on snooped packets when DHCP snooping is enabled by the **allow-snooped-clients** statement.

[Table 142 on page 1461](#) shows the action the router (or switch) takes on snooped packets when DHCP snooping is disabled by the **no-allow-snooped-clients** statement.

The router or switch also uses the configuration of the DHCP relay agent forwarding support to determine how to handle snooped BOOTREPLY packets. [Table 143 on page 1461](#) shows the action the router (or switch) takes for the snooped BOOTREPLY packets.



**NOTE:** Configured interfaces have been configured with the **group** statement in the **[edit forwarding-options dhcp-relay]** hierarchy. Non-configured interfaces are in the logical system/routing instance but have not been configured by the **group** statement.

**Table 141: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Enabled**

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
<b>forward-snooped-clients</b> not configured	snooped packets result in subscriber (DHCP client) creation	dropped
<b>all-interfaces</b>	forwarded	forwarded
<b>configured-interfaces</b>	forwarded	dropped

Table 141: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Enabled (*continued*)

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
non-configured-interfaces	snooped packets result in subscriber (DHCP client) creation	forwarded

Table 142: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Disabled

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
forward-snooped-clients not configured	dropped	dropped
all-interfaces	dropped	forwarded
configured-interfaces	dropped	dropped
non-configured-interfaces	dropped	forwarded

Table 143: Actions for Snooped BOOTREPLY Packets

forward-snooped-clients Configuration	Action
forward-snooped-clients not configured	snooped <b>BOOTREPLY</b> packets dropped if client is not found
forward-snooped-clients all configurations	snooped <b>BOOTREPLY</b> packets forwarded if client is not found

To configure DHCP snooped packet forwarding and BOOTREPLY snooped packet forwarding for DHCP relay agent:

- Specify that you want to configure DHCP relay agent.  

```
[edit]
user@host# edit forwarding-options dhcp-relay
```
- Enable DHCP snooped packet forwarding.  

```
[edit forwarding-options dhcp-relay]
user@host# edit forward-snooped-clients
```
- Specify the interfaces that are supported for snooped packet forwarding.  

```
[edit forwarding-options dhcp-relay forward-snooped-clients]
user@host# set (all-interfaces | configured-interfaces | non-configured-interfaces)
```

For example, to configure DHCP relay agent to forward DHCP snooped packets on only configured interfaces:

```
[edit]
forwarding-options {
  dhcp-relay {
    forward-snooped-clients configured-interfaces;
  }
}
```

- Related Documentation**
- [DHCP Snooping Support on page 1367](#)
  - [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 1455](#)

## Sending Release Messages When Clients Are Deleted

By default, when DHCP relay and relay proxy delete a client, they do not send a release message to the DHCP server. You can override the default behavior and configure DHCP relay and relay proxy to send a release message whenever they delete a client. The release message sent by DHCP relay and relay proxy includes option 82 information.



**NOTE:** You must include the `send-release-on-delete` statement to configure DHCP relay and relay proxy to send the release message when the `client-discover-match` statement is included.

You can use the `[edit forwarding-options dhcp-relay dhcpv6]` hierarchy level to override the default behavior for DHCPv6 relay agent.

To send a release message:

1. Specify that you want to configure override options.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit overrides
```

2. Specify that you want DHCP relay and relay proxy (or DHCPv6 relay agent) to send a release message when clients are deleted.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set send-release-on-delete
```

- Related Documentation**
- [Extended DHCP Relay Agent Overview on page 1379](#)
  - [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)

## Disabling Automatic Binding of Stray DHCP Requests

DHCP requests that are received but have no entry in the database are known as stray requests. By default, DHCP relay, DHCP relay proxy, and DHCPv6 relay agent attempt to bind the requesting client by creating a database entry and forwarding the request to the DHCP server. If the server responds with an ACK, the client is bound and the ACK is forwarded to the client. If the server responds with a NAK, the database entry is deleted and the NAK is forwarded to the client. This behavior occurs regardless of whether authentication is configured.

You can override the default configuration at the global level, for a named group of interfaces, or for a specific interface within a named group. Overriding the default causes DHCP relay, DHCP relay proxy, and DHCPv6 relay agent to drop all stray requests instead of attempting to bind the clients.



**NOTE:** Automatic binding of stray requests is enabled by default.

- To disable automatic binding behavior, include the **no-bind-on-request** statement when you configure DHCP overrides at the global, group, or interface level.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set no-bind-on-request
```

- To override the default behavior for DHCPv6 relay agent, configure the override at the **[edit forwarding-options dhcp-relay dhcpv6]** hierarchy level.

```
[edit forwarding-options dhcp-relay dhcpv6 overrides]
user@host# set no-bind-on-request
```

The following two examples show a configuration that disables automatic binding of stray requests for a group of interfaces and a configuration that disables automatic binding on a specific interface.

To disable automatic binding of stray requests on a group of interfaces:

- Specify the named group.

```
[edit forwarding-options dhcp-relay]
user@host# edit group boston
```

- Specify that you want to configure overrides.

```
[edit forwarding-options dhcp-relay group boston]
user@host# edit overrides
```

- Disable automatic binding for the group.

```
[edit forwarding-options dhcp-relay group boston overrides]
user@host# set no-bind-on-request
```

To disable automatic binding of stray requests on a specific interface:

- Specify the named group of which the interface is a member.

```
[edit forwarding-options dhcp-relay]
```

```
user@host# edit group boston
```

2. Specify the interface on which you want to disable automatic binding.

```
[edit forwarding-options dhcp-relay group boston]
```

```
user@host# edit interface fe-1/0/1.2
```

3. Specify that you want to configure overrides.

```
[edit forwarding-options dhcp-relay group boston interface fe-1/0/1.2]
```

```
user@host# edit overrides
```

4. Disable automatic binding on the interface.

```
[edit forwarding-options dhcp-relay group boston interface fe-1/0/1.2 overrides]
```

```
user@host# set no-bind-on-request
```

**Related  
Documentation**

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)

## Using DHCP Relay Agent Option 82 Information

Subscriber management enables you to configure DHCP relay agent to include additional option 82 information in the DHCP packets that the relay receives from clients and forwards to a DHCP server. The DHCP server uses the additional information to determine the IP address to assign to the client. The server might also use the information for other purposes; for example, to determine which services to grant the client, or to provide additional security against threats such as address spoofing. The DHCP server sends its reply back to the DHCP relay agent, and the agent removes the option 82 information from the message and forwards the packet to the client.

To configure support for the DHCP relay agent information option 82, you use the **relay-option-82** statement. You can configure the DHCP relay agent to include the following suboptions in the packet the relay sends to the DHCP server:

- Agent Circuit ID (suboption 1)—An ASCII string that identifies the interface on which the client DHCP packet is received.
- Agent Remote ID (suboption 2)—An ASCII string assigned by the DHCP relay agent that securely identifies the client.

You can configure the option 82 support globally or for a named group of interfaces.

To restore the default behavior, in which option 82 information is not inserted into DHCP packets, you use the **delete relay-option-82** statement.





**NOTE:** The DHCPv6 relay agent provides similar Agent Circuit ID and Agent Remote ID support for DHCPv6 clients. For DHCPv6, subscriber management uses DHCPv6 option 18 to include the circuit ID in the packets that the relay sends to a DHCPv6 server, and option 37 to include the remote ID in the packets. See *DHCPv6 Relay Agent Options*.)

The following sections describe the option 82 operations you can configure:

- [Configuring Option 82 Information on page 1465](#)
- [Including a Prefix in DHCP Options on page 1466](#)
- [Including a Textual Description in DHCP Options on page 1468](#)

### Configuring Option 82 Information

You use the **relay-option-82** statement to configure DHCP relay agent to insert option 82 information in DHCP packets that the relay receives from clients and forwards to a DHCP server. When you configure option 82, you also include at least one of the suboption statements to specify the type of information you want to include in the DHCP packets. Use the **circuit-id** statement to include the Agent Circuit ID (suboption 1) in the packets, or the **remote-id** statement to include the Agent Remote ID (suboption 2).

You can optionally configure DHCP relay agent to include a prefix or the interface description as part of the suboption information. If you specify the **circuit-id** or **remote-id** statement without including either of the optional **prefix** or **use-interface-description** statements, the format of the Agent Circuit ID or Agent Remote ID information for Fast Ethernet (**fe**) or Gigabit Ethernet (**ge**) interfaces is one of the following, depending on your network configuration:

- For Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual local area networks (VLANs) or stacked VLANs (S-VLANs):

`(fe | ge)-fpc/pic/port`

- For Fast Ethernet or Gigabit Ethernet interfaces that use VLANs:

`(fe | ge)-fpc/pic/port:vlan-id`

- For Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs:

`(fe | ge)-fpc/pic/port:svlan-id-vlan-id`

To enable insertion of option 82 information:

1. Specify that you want to configure option 82 support.

```
[edit forwarding-options dhcp-relay]
user@host# edit relay-option-82
```

2. Configure DHCP relay agent to insert the Agent Circuit ID suboption, the Agent Remote ID suboption, or both.

```
[edit forwarding-options dhcp-relay relay-option-82]
user@host# set circuit-id
```

3. (Optional) Configure a prefix that is used in the option 82 information in the DHCP packets

See [“Including a Prefix in DHCP Options” on page 1466](#).

4. (Optional) Configure DHCP relay agent to include the interface’s textual description instead of interface identifier in the option 82 information.

See [“Including a Textual Description in DHCP Options” on page 1468](#).

---

### Including a Prefix in DHCP Options

When you configure DHCPv4 or DHCPv6 relay agent to include DHCP options in the packets the relay agent sends to a DHCP server, you can specify that the relay agent add a prefix to the DHCP option. You can add a prefix to the following DHCP options:

- DHCPv4 Option 82 Agent Circuit ID (suboption 1)
- DHCPv4 Option 82 Agent Remote ID (suboption 2)
- DHCPv6 Relay Agent Interface-ID (option 18)
- DHCPv6 Relay Agent Remote-ID (option 37)

The prefix is separated from the DHCP option information by a colon (:), and can include any combination of the **host-name**, **logical-system-name**, and **routing-instance-name** options. The DHCP relay agent obtains the values for the **host-name**, **logical-system-name**, and **routing-instance-name** as follows:

- If you include the **host-name** option, the DHCP relay agent uses the hostname of the router configured with the **host-name** statement at the **[edit system]** hierarchy level.
- If you include the **logical-system-name** option, the DHCP relay agent uses the logical system name configured with the **logical-system** statement at the **[edit logical-system]** hierarchy level.
- If you include the **routing-instance-name** option, the DHCP relay agent uses the routing instance name configured with the **routing-instance** statement at the **[edit routing-instances]** hierarchy level or at the **[edit logical-system logical-system-name routing-instances]** hierarchy level.

If you include the hostname and either or both of the logical system name and the routing instance name in the prefix, the hostname is followed by a forward slash (/). If you include both the logical system name and the routing instance name in the prefix, these values are separated by a semicolon (;).

The following examples show several possible formats for the DHCP option information when you specify the **prefix** statement for Fast Ethernet (**fe**) or Gigabit Ethernet (**ge**) interfaces with S-VLANs.

- If you include only the hostname in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

*hostname:(fe | ge)-fpc/pic/port:svlan-id-vlan-id*

- If you include only the logical system name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

*logical-system-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id*

- If you include only the routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

*routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id*

- If you include both the hostname and the logical system name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

*host-name/logical-system-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id*

- If you include both the logical system name and the routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

*logical-system-name;routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id*

- If you include the hostname, logical system name, and routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

*host-name/logical-system-name;routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id*

For Fast Ethernet or Gigabit Ethernet interfaces that use VLANs but not S-VLANs, only the **vlan-id** value appears in the DHCP option format. For Fast Ethernet or Gigabit Ethernet interfaces that do not use VLANs or S-VLANs, neither the **vlan-id** value nor the **svlan-id** value appears.

(DHCPv4) To configure a prefix with the option 82 information:

1. Specify that you want to configure option 82 support.

```
[edit forwarding-options dhcp-relay]
user@host# edit relay-option-82
```

2. Configure DHCP relay agent to insert the Agent Circuit ID, the Agent Remote ID, or both.

```
[edit forwarding-options dhcp-relay relay-option-82]
user@host# edit circuit-id
```

3. Specify that the prefix is included in the option 82 information. In this example, the prefix includes the hostname and logical system name

```
[edit forwarding-options dhcp-relay relay-option-82 circuit-id]
user@host# set prefix host-name logical-system-name
```

(DHCPv6) To use a prefix with the DHCPv6 option 18 or option 37 information:

1. Specify that you want to configure DHCPv6 relay agent support.

```
[edit forwarding-options dhcp-relay]
user@host# edit dhcpv6
```

2. Configure DHCPv6 relay agent to insert option 18 (Relay Agent Interface-ID), option 37 (Relay Agent Remote-ID), or both.

```
[edit forwarding-options dhcp-relay dhcpv6]
```

```
user@host# edit relay-agent-interface-id
```

3. Specify that the prefix is included in the option information. In this example, the prefix includes the hostname and logical system name

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]  
user@host# set prefix host-name logical-system-name
```

---

### Including a Textual Description in DHCP Options

By default, when DHCPv4 or DHCPv6 relay agent inserts agent circuit ID or agent remote ID option information in the packets sent to a DHCP server, the options include the interface identifier. However, you can configure the DHCP relay agent to include the textual description that is configured for the interface instead of the interface identifier. You can use the textual description for either the logical interface or the device interface.

You can include the textual interface description in the agent circuit ID or agent remote ID options for static interfaces.

The textual description is configured separately, using the **description** statement at the **[edit interfaces interface-name]** hierarchy level. If you specify that the textual description is used and no description is configured for the interface, DHCP relay defaults to using the interface identifier.

You can use the textual description with the following DHCP options:

- DHCPv4 Option 82 Agent Circuit ID (suboption 1)
- DHCPv4 Option 82 Agent Remote ID (suboption 2)
- DHCPv6 Relay Agent Interface-ID (option 18)
- DHCPv6 Relay Agent Remote-ID (option 37)

(DHCPv4) To configure the DHCP relay option 82 suboption to include the textual interface description:

1. Specify that you want to configure option 82 support.

```
[edit forwarding-options dhcp-relay]  
user@host# edit relay-option-82
```

2. Configure DHCP relay agent to insert the Agent Circuit ID, Agent Remote ID, or both.

```
[edit forwarding-options dhcp-relay relay-option-82]  
user@host# edit circuit-id
```

3. Specify that the textual description is included in the option 82 information. In this example, the option 82 information includes the description used for the device interface.

```
[edit forwarding-options dhcp-relay relay-option-82 circuit-id]  
user@host# set use-interface-description device
```

(DHCPv6) To use a textual description with the DHCPv6 option 18 or option 37 information:

1. Specify that you want to configure DHCPv6 relay agent support.

```
[edit forwarding-options dhcp-relay]
user@host# edit dhcpv6
```

2. Configure DHCPv6 relay agent to insert option 18 (Relay Agent Interface-ID), option 37 (Relay Agent Remote-ID), or both.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit relay-agent-interface-id
```

3. Specify that the textual description is included in the option information. In this example, the option 37 information includes the description used for the device interface.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
user@host# set use-interface-description device
```

## Configuring Server Groups

You can configure a named group of DHCP servers for use by the extended DHCP relay agent on the router or switch.

You specify the name of the DHCP server group and the IP addresses of one or more DHCP servers that belong to this group. You can configure a maximum of five IP addresses per named server group.

To configure a named server group:

1. Specify the name of the server group.

```
[edit forwarding-options dhcp-relay]
user@host# set server-group myServerGroup
```

2. Add the IP addresses of the DHCP servers belonging to the group.

```
[edit forwarding-options dhcp-relay server-group myServerGroup]
user@host# set 192.168.100.50
user@host# set 192.168.100.75
```

**Related Documentation**

- [Extended DHCP Relay Agent Overview on page 1379](#)

## Configuring Active Server Groups

You can configure an active server group. Using an active server group enables you to apply a common DHCP relay agent configuration to a named group of DHCP server addresses.

Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

To configure an active server group:

- Specify the name of the active server group.

```
[edit forwarding-options dhcp-relay]
user@host# set active-server-group myServerGroup
```

To create an active server group as a global DHCP relay agent configuration option, include the **active-server-group** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level. To have the group apply only to a named group of interfaces, include the **active-server-group** statement at the **[edit forwarding-options dhcp-relay group group-name]** hierarchy level.

Including the **active-server-group** statement at the **[edit forwarding-options dhcp-relay group group-name]** hierarchy level (as a group-specific option) overrides the effect of including the **active-server-group** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level as a global option.

**Related  
Documentation**

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [Grouping Interfaces with Common DHCP Configurations on page 1421](#)

## Enabling DHCP Relay Proxy Mode

You can enable DHCP relay proxy mode on all interfaces or a group of interfaces.

To enable DHCP relay proxy mode:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Enable DHCP relay proxy mode.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set proxy-mode
```

**Related  
Documentation**

- [DHCP Relay Proxy Overview on page 1382](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)

## Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets

You can configure DHCPv6 relay agent to insert the DHCPv6 Interface-ID (option 18) in the packets that the relay sends to a DHCPv6 server. You can configure the option 18 support at either the DHCPv6 global or group level.

When you configure option 18 support, you can optionally include the following additional information:

- Prefix—Specify the **prefix** option to add a prefix to the interface identifier. The prefix can be any combination of hostname, logical system name, and routing instance name.
- Interface description—Specify the **use-interface-description** option to include the textual interface description instead of the interface identifier. You can include either the device interface description or the logical interface description.

- Option 82 Agent Circuit ID suboption (suboption 1)—Specify the **use-option-82** option to include the DHCPv4 Option 82 Agent Circuit ID suboption (suboption 1). This configuration is useful in a dual-stack environment, which has both DHCPv4 and DHCPv6 subscribers that reside over the same underlying logical interface. The router checks for the option 82 suboption 1 value and inserts it into the outgoing packets. If no DHCPv4 binding exists or if the binding does not have an option 82 suboption 1 value, the router sends the packets without adding an option 18.



**NOTE:** If you specify one of the optional configurations, and the specified information does not exist (for example, there is no interface description), DHCPv6 relay ignores the optional configuration and inserts the default interface identifier in the packets.

To insert the DHCPv6 Interface-ID option (option 18) in DHCPv6 packets:

1. Configure the DHCPv6 relay to include option 18.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit relay-agent-interface-id
```

2. (Optional) Specify the prefix to include in option 18.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
user@host# set prefix prefix
```

3. (Optional) Specify that option 18 include the textual description of the interface. You can specify either the **logical** interface description or the **device** interface description.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
user@host# set use-interface-description (logical | device)
```

4. (Optional) Specify that option 18 use the DHCPv4 Option 82 Agent Circuit ID suboption (suboption 1) value.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
user@host# set use-option-82
```

#### Related Documentation

- [DHCPv6 Relay Agent Options](#)
- [Configuring DHCPv6 Relay Agent Options](#)
- [Including a Prefix in DHCP Options on page 1466](#)
- [Including a Textual Description in DHCP Options on page 1468](#)

## DHCP Liveness Detection Overview

Unlike PPP, DHCP does not define a native keepalive mechanism as part of either the DHCPv4 or DHCPv6 protocols. Without a keepalive mechanism, DHCP local server, DHCP relay, or DHCP relay proxy is unable to quickly detect if it has lost connectivity with a subscriber or a DHCP client; and it must rely on standard DHCP subscriber session or DHCP client session termination messages.

DHCP clients often do not send DHCP release messages prior to exiting the network. The discovery of their absence is dependent on existing DHCP lease time and release request mechanisms. These mechanisms are often considered insufficient when serving as session health checks for clients in a DHCP subscriber access or a DHCP-managed network. Because DHCP lease times are typically too long to provide an adequate response time for a session health failure, and configuring short DHCP lease times can pose an undue burden on control plane processing, implementing a DHCP liveness detection mechanism enables better monitoring of bound DHCP clients. When configured with a liveness detection protocol, if a given subscriber (or client) fails to respond to a configured number of consecutive liveness detection requests, the subscriber (or client) binding is deleted and its resources released.

DHCP liveness detection for DHCP subscriber IP or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.

Using DHCP liveness detection, IP sessions are acted upon as soon as liveness detection checks fail. This faster response time serves to:

- Provide more accurate time-based accounting of subscriber (or DHCP client) sessions.
- Better preserve router (switch) resources.
- Help to reduce the window of vulnerability to some security attacks.

Examples of liveness detection protocols include Bidirectional Forwarding Detection (BFD) for both DHCPv4 and DHCPv6 subscribers, IPv4 Address Resolution Protocol (ARP) for DHCPv4 subscribers, and IPv6 Neighbor Unreachability Detection for DHCPv6 subscribers.



**NOTE:** Only BFD for DHCPv4 and DHCPv6 liveness detection is supported.

---

When configuring BFD liveness detection, keep the following in mind:

- You can configure DHCPv4 and DHCPv6 liveness detection either globally or per DHCPv4 or DHCPv6 group.
- DHCPv4 or DHCPv6 subscriber access clients that do not support BFD are not affected by the liveness detection configuration. These clients can continue to access the



network (once validated) even if BFD liveness detection is enabled on the router (or switch).

- When configured, DHCPv4 or DHCPv6 initiates liveness detection checks for relevant clients (that is, clients that support BFD) when those clients enter a bound state.
- After protocol-specific messages are initiated for a BFD client, they are periodically sent to the subscriber (or client) IP address of the client and responses to those liveness detection requests are expected within a configured amount of time.
- If liveness detection responses are not received from clients that support BFD within the configured amount of time for a configured number of consecutive attempts, the liveness detection check is deemed to have failed and a configured failure action is implemented.

#### Related Documentation

- [Configuring Detection of DHCP Local Server Client Connectivity on page 1436](#)
- [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 1473](#)

## Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity

Liveness detection for DHCP subscriber IP or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.

To configure liveness detection for DHCP relay:

1. Specify that you want to configure liveness detection.

- For DHCP global configuration:

```
[edit forwarding-options dhcp-relay]
user@host# edit liveness-detection
```

- For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name]
user@host# edit liveness-detection
```



**NOTE:** Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the **liveness-detection** statement, and any subsequent configuration statements, at the **[edit forwarding-options dhcp-relay dhcpv6]** or **[edit forwarding-options dhcp-relay dhcpv6 group *group-name*]** hierarchy level.

2. (Optional) Specify that you want to use DHCP relay proxy mode.

```
[edit forwarding-options dhcp-relay group group-name]
user@host# set overrides proxy-mode
```

3. Specify that you want to configure the liveness detection method.

- For DHCP global configuration:

```
[edit forwarding-options dhcp-relay liveness-detection]  
user@host# edit method
```

- For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name liveness-detection]  
user@host# edit method
```

4. Specify the liveness detection method that you want DHCP to use.



**NOTE:** The only method supported for liveness detection is Bidirectional Forwarding Detection (BFD).

- For DHCP global configuration:

```
[edit forwarding-options dhcp-relay liveness-detection method]  
user@host# edit bfd
```

- For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name liveness-detection method]  
user@host# edit bfd
```

5. Configure the liveness detection method as desired.

See *Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients* for an example of how to globally configure DHCP relay liveness detection.

6. Configure the action the router takes when a liveness detection failure occurs.

- For DHCP global configuration:

```
[edit forwarding-options dhcp-relay liveness-detection]  
user@host# edit failure-action action
```

- For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name liveness-detection]  
user@host# edit failure-action action
```

**Related  
Documentation**

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [DHCP Liveness Detection Overview on page 1472](#)
- [Configuring Detection of DHCP Local Server Client Connectivity on page 1436](#)
- [Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395](#)
- [Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients](#)

## Disabling DHCP Relay

You can disable DHCP relay on all interfaces or a group of interfaces.

To disable DHCP relay agent:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Disable the DHCP relay agent.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set disable-relay
```

**Related  
Documentation**

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)

## DHCP Local Server Configuration Statements

- [attempts \(DHCP Local Server\) on page 1477](#)
- [authentication \(DHCP Local Server\) on page 1478](#)
- [bfd on page 1479](#)
- [circuit-type \(DHCP Local Server\) on page 1480](#)
- [clear-on-abort \(DHCP Local Server\) on page 1481](#)
- [client-discover-match \(DHCP Local Server\) on page 1482](#)
- [client-id \(DHCP Local Server\) on page 1483](#)
- [delegated-pool \(DHCP Local Server\) on page 1484](#)
- [delimiter \(DHCP Local Server\) on page 1485](#)
- [detection-time on page 1486](#)
- [dhcp \(DHCP Client\) on page 1487](#)
- [dhcp-local-server on page 1488](#)
- [dhcpx6 \(DHCP Local Server\) on page 1493](#)
- [domain-name \(DHCP Local Server\) on page 1496](#)
- [dynamic-profile \(DHCP Local Server\) on page 1497](#)
- [external-authority on page 1498](#)
- [failure-action on page 1499](#)
- [forward-snooped-clients \(DHCP Local Server\) on page 1500](#)
- [group \(DHCP Local Server\) on page 1501](#)
- [holddown-interval on page 1503](#)
- [interface \(DHCP Local Server\) on page 1504](#)
- [interface-client-limit \(DHCP Local Server\) on page 1506](#)
- [interface-delete \(Subscriber Management or DHCP Client Management\) on page 1507](#)
- [interface-name \(DHCP Local Server\) on page 1508](#)
- [ip-address-first on page 1509](#)

- [liveness-detection](#) on page 1510
- [mac-address \(DHCP Local Server\)](#) on page 1511
- [method](#) on page 1512
- [minimum-interval](#) on page 1513
- [minimum-receive-interval](#) on page 1514
- [multiplier](#) on page 1515
- [no-adaptation](#) on page 1516
- [option-60 \(DHCP Local Server\)](#) on page 1517
- [option-82 \(DHCP Local Server Authentication\)](#) on page 1518
- [option-82 \(DHCP Local Server Pool Matching\)](#) on page 1519
- [overrides \(DHCP Local Server\)](#) on page 1520
- [password \(DHCP Local Server\)](#) on page 1522
- [pool \(DHCP Local Server Overrides\)](#) on page 1523
- [pool-match-order](#) on page 1524
- [process-inform](#) on page 1525
- [radius-disconnect \(DHCP Local Server\)](#) on page 1527
- [rapid-commit \(DHCPv6 Local Server\)](#) on page 1528
- [reconfigure \(DHCP Local Server\)](#) on page 1529
- [relay-agent-interface-id \(DHCP Local Server\)](#) on page 1530
- [relay-agent-remote-id \(DHCP Local Server\)](#) on page 1531
- [routing-instance-name \(DHCP Local Server\)](#) on page 1532
- [service-profile \(DHCP Local Server\)](#) on page 1533
- [session-mode](#) on page 1534
- [strict \(DHCP Local Server\)](#) on page 1535
- [threshold \(detection-time\)](#) on page 1536
- [threshold \(transmit-interval\)](#) on page 1537
- [timeout \(DHCP Local Server\)](#) on page 1538
- [token \(DHCP Local Server\)](#) on page 1539
- [traceoptions \(DHCP Server\)](#) on page 1540
- [transmit-interval](#) on page 1543
- [trigger \(DHCP Local Server\)](#) on page 1544
- [use-primary \(DHCP Local Server\)](#) on page 1545
- [user-prefix \(DHCP Local Server\)](#) on page 1546
- [username-include \(DHCP Local Server\)](#) on page 1547
- [version \(BFD\)](#) on page 1548

## attempts (DHCP Local Server)

<b>Syntax</b>	<code>attempts <i>attempt-count</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 10.4.</p>
<b>Description</b>	Configure how many attempts are made to reconfigure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces before reconfiguration is considered to have failed. A group configuration takes precedence over a DHCP local server configuration.
<b>Options</b>	<p><b><i>attempt-count</i></b>—Maximum number of attempts.</p> <p><b>Range:</b> 1 through 10</p> <p><b>Default:</b> 8</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429</a></li> <li>• <a href="#">Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 1432</a></li> </ul>

## authentication (DHCP Local Server)

---

Syntax	<pre>authentication {   password <i>password-string</i>;   username-include {     circuit-type;     client-id;     delimiter <i>delimiter-character</i>;     domain-name <i>domain-name-string</i>;     interface-name;     logical-system-name;     mac-address;     option-60;     option-82 &lt;circuit-id&gt; &lt;remote-id&gt;;     relay-agent-interface-id;     relay-agent-remote-id;     relay-agent-subscriber-id;     routing-instance-name;     user-prefix <i>user-prefix-string</i>;   } }</pre>
Hierarchy Level	<pre>[edit system services <a href="#">dhcp-local-server</a>], [edit system services dhcp-local-server <a href="#">dhcpv6</a>], [edit system services dhcp-local-server dhcpv6 <a href="#">group group-name</a>], [edit system services dhcp-local-server <a href="#">group group-name</a>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system   services <a href="#">dhcp-local-server</a> ...], [edit logical-systems <i>logical-system-name</i> system services <a href="#">dhcp-local-server</a> ...], [edit routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server</a> ...]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay or DHCP local server configuration.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li></ul>

## bfd

<b>Syntax</b>	<pre> bfd {   version (0   1   automatic);   minimum-interval <i>milliseconds</i>;   minimum-receive-interval <i>milliseconds</i>;   multiplier <i>number</i>;   no-adaptation;   transmit-interval {     minimum-interval <i>milliseconds</i>;     threshold <i>milliseconds</i>;   }   detection-time {     threshold <i>milliseconds</i>;   }   session-mode (automatic   multihop   singlehop);   holddown-interval <i>milliseconds</i>; } </pre>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection <a href="#">method</a>],  [edit system services dhcp-local-server dhcpv6 liveness-detection <a href="#">method</a>],  [edit forwarding-options dhcp-relay liveness-detection <a href="#">method</a>],  [edit forwarding-options dhcp-relay dhcpv6 liveness-detection <a href="#">method</a>],  [edit system services dhcp-local-server group <i>group-name</i> liveness-detection <a href="#">method</a>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection <a href="#">method</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection <a href="#">method</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection <a href="#">method</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.  Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure Bidirectional Forwarding Detection (BFD) as the liveness detection method.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.  routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## circuit-type (DHCP Local Server)


<b>Syntax</b>	circuit-type;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit system services dhcp-local-server <b>group group-name authentication username-include</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify that the circuit type is concatenated with the username during the subscriber authentication or client authentication process.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> </ul>



## clear-on-abort (DHCP Local Server)

<b>Syntax</b>	clear-on-abort;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 10.4.</p>
<b>Description</b>	Delete all DHCP clients or only the DHCP clients serviced by the specified group of interfaces when reconfiguration fails; that is, when the maximum number of retry attempts have been made without success. A group configuration takes precedence over a DHCP local server configuration.
<b>Default</b>	Restores the original client configuration when reconfiguration fails.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429</a></li> <li>• <a href="#">Configuring Deletion of the Client When Dynamic Reconfiguration Fails on page 1433</a></li> </ul>

## client-discover-match (DHCP Local Server)

<b>Syntax</b>	client-discover-match (option60-and-option82   incoming-interface);
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>]</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ... <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...<a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...<a href="#">overrides</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Option <b>incoming-interface</b> introduced in Junos OS Release 13.3.</p>
<b>Description</b>	Configure the match criteria DHCP local server uses to uniquely identify DHCP subscribers or clients when primary identification fails. The options are mutually exclusive.
<b>Options</b>	<p><b>incoming-interface</b>—Allow only one client device to connect on the interface. If the client device changes, the router deletes the existing client binding and creates a binding for the newly connected device.</p>
<div style="display: flex; align-items: flex-start;"> <div style="flex: 1; text-align: center; margin-right: 10px;">  </div> <div> <p><b>NOTE:</b> The <b>overrides client-discover-match incoming-interface</b> configuration deletes and replaces the existing binding when a new device connects. This action differs from the <b>overrides interface-client-limit 1</b> statement, which retains the existing binding and rejects the newly connected client.</p> </div> </div>	
<p><b>option60-and-option82</b>—Use option 60 and option 82 information to identify subscribers.</p>	
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Local Server Overview on page 1356</a></li> <li>• <a href="#">Overriding Default DHCP Local Server Configuration Settings on page 1423</a></li> <li>• <a href="#">DHCP Auto Logout Overview on page 1368</a></li> <li>• <a href="#">Allowing Only One DHCP Client Per Interface</a></li> </ul>

## client-id (DHCP Local Server)

<b>Syntax</b>	client-id;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify that the DHCPv6 Client-ID option (option 1) in the client PDU name is concatenated with the username during the subscriber authentication or client authentication process.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li> </ul>

## delegated-pool (DHCP Local Server)

---

<b>Syntax</b>	<code>delegated-pool <i>pool-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit system services dhcp-local-server <a href="#">dhcpv6 overrides</a>],</code> <code>[edit system services dhcp-local-server dhcpv6 <a href="#">group group-name overrides</a>],</code> <code>[edit system services dhcp-local-server dhcpv6 <a href="#">group group-name</a> interface <i>interface-name</i></code> <code><a href="#">overrides</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system</code> <code>services dhcp-local-server <a href="#">dhcpv6</a> ...],</code> <code>[edit logical-systems <i>logical-system-name</i> system services system services dhcp-local-server</code> <code><a href="#">dhcpv6</a> ...],</code> <code>[edit routing-instances <i>routing-instance-name</i> system services system services</code> <code>dhcp-local-server <a href="#">dhcpv6</a> ...]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Specify the address pool that assigns the IA_PD address. A pool specified by RADIUS VSA 26-161 takes precedence over the pool specified by this <b>delegated-pool</b> statement.
<b>Options</b>	<b><i>pool-name</i></b> —Name of the address-assignment pool.
<b>Required Privilege Level</b>	<b>system</b> —To view this statement in the configuration. <b>system-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Specifying the Delegated Address Pool for IPv6 Prefix Assignment on page 1428</a></li><li>• <a href="#">Overriding Default DHCP Local Server Configuration Settings on page 1423</a></li></ul>

## delimiter (DHCP Local Server)

<b>Syntax</b>	<code>delimiter <i>delimiter-character</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server dhcpv6 group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server dhcpv6 group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server dhcpv6 group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server group group-name authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server dhcpv6 authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server dhcpv6 group group-name authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server group group-name authentication username-include</b>],</p> <p>[edit system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit system services <b>dhcp-local-server dhcpv6 authentication username-include</b>],</p> <p>[edit system services <b>dhcp-local-server dhcpv6 group group-name authentication username-include</b>],</p> <p>[edit system services <b>dhcp-local-server group group-name authentication username-include</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify the character used as the delimiter between the concatenated components of the username.
<b>Options</b>	<b><i>delimiter-character</i></b> —Character that separates components that make up the concatenated username. You cannot use the semicolon (;) as a delimiter.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Using External AAA Authentication Services with DHCP on page 1420](#)

---

## detection-time

---

<b>Syntax</b>	<pre>detection-time {     threshold milliseconds; }</pre>
<b>Hierarchy Level</b>	<pre>[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group group-name liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd], [edit forwarding-options dhcp-relay group group-name liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Enable failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li><li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li></ul>

## dhcp (DHCP Client)

<b>Syntax</b>	<pre>dhcp {   client-identifier (ascii <i>ascii</i>   hexadecimal <i>hexadecimal</i>);   lease-time (<i>seconds</i>   infinite);   retransmission-attempt <i>number</i>;   retransmission-interval <i>seconds</i>;   server-address <i>ip-address</i>;   update-server;   vendor-id <i>vendor-id</i>; }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <i>unit</i> <i>logical-unit-number</i> <i>family</i> inet]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	<p>Configure a DHCP client for an IPv4 interface.</p> <p>The remaining statements are described separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a DHCP Client (CLI Procedure) on page 1413</a></li> </ul>

## dhcp-local-server

---

```
Syntax  dhcp-local-server {
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dhcpv6 {
            authentication {
                ...
            }
            group group-name {
                authentication {
                    ...
                }
                interface interface-name {
                    exclude;
                    liveness-detection {
                        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                        method {
                            bfd {
                                version (0 | 1 | automatic);
                                minimum-interval milliseconds;
                                minimum-receive-interval milliseconds;
                                multiplier number;
                                no-adaptation;
                                transmit-interval {
                                    minimum-interval milliseconds;
                                    threshold milliseconds;
                                }
                                detection-time {
                                    threshold milliseconds;
                                }
                            }
                            session-mode (automatic | multihop | singlehop);
                            holddown-interval milliseconds;
                        }
                    }
                }
            }
        }
        overrides {
            interface-client-limit number;
            multi-address-embedded-option-response;
            process-inform {
                pool pool-name;
            }
        }
    }
```



```

    }
    rapid-commit;
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
overrides {
    delegated-pool;
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
}
overrides {
    delegated-pool;
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {

```

```
        radius-disconnect;
    }
}
route-suppression;
service-profile dynamic-profile-name;
}
duplicate-clients-in-subnet (incoming-interface | option-82);
dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
primary-profile-name>;
forward-snooped-clients (all-interfaces | configured-interfaces |
non-configured-interfaces);
group group-name {
    authentication {
        ...
    }
}
dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
primary-profile-name>;
interface interface-name {
    exclude;
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
}
overrides {
    client-discover-match (option60-and-option82 | incoming-interface);
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
overrides {
    client-discover-match (option60-and-option82 | incoming-interface);
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
}
```

```

    }
  }
  requested-ip-network-match subnet-mask
  route-suppression;
  service-profile dynamic-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode (automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
overrides {
  client-discover-match (option60-and-option82 | incoming-interface);
  interface-client-limit number;
  process-inform {
    pool pool-name;
  }
}
pool-match-order {
  external-authority;
  ip-address-first;
  option-82;
}
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}
requested-ip-network-match subnet-mask;
route-suppression;
service-profile dynamic-profile-name;
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services],  
[edit logical-systems *logical-system-name* system services],  
[edit routing-instances *routing-instance-name* system services],  
[edit system services]

**Release Information** Statement introduced in Junos OS Release 9.0.  
Statement introduced in Junos OS Release 12.1 for EX Series switches.  
Statement introduced in Junos OS Release 13.2X51 for the QFX Series.

**Description** Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch and enable the router or switch to function as an extended DHCP local server. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client.

The DHCP local server and the DHCP/BOOTP relay server, which are configured under the **[edit forwarding-options helpers]** hierarchy level, cannot both be enabled on the router or switch at the same time. The extended DHCP local server is fully compatible with the extended DHCP relay feature.

The **dhcpx6** stanza configures the router or switch to support Dynamic Host Configuration Protocol for IPv6 (DHCPv6). The DHCPv6 local server is fully compatible with the extended DHCP local server and the extended DHCP relay feature.



**NOTE:** When you configure the **dhcp-local-server** statement at the routing instance hierarchy level, you must use a routing instance type of **virtual-router**.

---

The remaining statements are explained separately.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Extended DHCP Local Server Overview on page 1356](#)
- [DHCPv6 Local Server Overview on page 1361](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\)](#)

## dhcpv6 (DHCP Local Server)

```
Syntax  dhcpv6 {
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            logical-system-name;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    group group-name {
        authentication {
            ...
        }
        interface interface-name {
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                    }
                    session-mode (automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
            }
        }
    }
    overrides {
        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
        }
        rapid-commit;
    }
    service-profile dynamic-profile-name;
    trace;
```

```
    upto upto-interface-name;
  }
  overrides {
    delegated-pool;
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
      pool pool-name;
    }
    rapid-commit;
  }
  route-suppression;
  service-profile dynamic-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode (automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
overrides {
  delegated-pool;
  interface-client-limit number;
  multi-address-embedded-option-response;
  process-inform {
    pool pool-name;
  }
  rapid-commit;
  reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
      radius-disconnect;
    }
  }
}
reconfigure {
  attempts attempt-count;
```

```

clear-on-abort;
strict;
timeout timeout-value;
token token-value;
trigger {
    radius-disconnect;
}
}
requested-ip-network-match subnet-mask;
route-suppression;
service-profile dynamic-profile-name;
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services [dhcp-local-server](#)],  
 [edit logical-systems *logical-system-name* system services [dhcp-local-server](#)],  
 [edit routing-instances *routing-instance-name* system services [dhcp-local-server](#)],  
 [edit system services [dhcp-local-server](#)]

**Release Information** Statement introduced in Junos OS Release 9.6.  
 Statement introduced in Junos OS Release 12.3 for EX Series switches.

**Description** Configure DHCPv6 local server options on the router or switch and enable the router or switch to function as a server for the DHCP protocol for IP version 6 (IPv6). The DHCPv6 local server sends and receives packets using the IPv6 protocol and informs IPv6 of the routing requirements of router clients. The local server works together with the AAA service framework to control subscriber access (or DHCP client access) and accounting.

The DHCPv6 local server is fully compatible with the extended DHCP local server and DHCP relay agent.

The remaining statements are explained separately.

**Required Privilege Level** system—To view this statement in the configuration.  
 system-control—To add this statement to the configuration.

**Related Documentation** • [DHCPv6 Local Server Overview on page 1361](#)

## domain-name (DHCP Local Server)

<b>Syntax</b>	<code>domain-name <i>domain-name-string</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit system services dhcp],</p> <p>[edit system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit system services dhcp-local-server <b>group group-name authentication username-include</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify the domain name that is concatenated with the username during the subscriber authentication or DHCP client authentication process.
<b>Options</b>	<b><i>domain-name-string</i></b> —Domain name formatted string.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>



- Related Documentation**
- [Using External AAA Authentication Services with DHCP on page 1420](#)

## dynamic-profile (DHCP Local Server)

<b>Syntax</b>	dynamic-profile <i>profile-name</i> { aggregate-clients (merge   replace); use-primary <i>primary-profile-name</i> ; }
<b>Hierarchy Level</b>	[edit system services <a href="#">dhcp-local-server</a> ], [edit system services dhcp-local-server <a href="#">dhcpv6</a> ], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">interface</a> <i>interface-name</i> ], [edit system services dhcp-local-server <a href="#">group</a> <i>group-name</i> ], [edit system services dhcp-local-server group <i>group-name</i> <a href="#">interface</a> <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> system services <a href="#">dhcp-local-server</a> ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server</a> ...], [edit routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server</a> ...]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Options <b>aggregate-clients</b> and <b>use-primary</b> introduced in Junos OS Release 9.3. Support at the [edit ... <a href="#">interface</a> ] hierarchy levels introduced in Junos OS Release 11.2.
<b>Description</b>	Specify the dynamic profile that is attached to all interfaces, a named group of interfaces, or a specific interface.
<b>Options</b>	<b><i>profile-name</i></b> —Name of the dynamic profile.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</a></li> <li>• <a href="#">Configuring a Default Subscriber Service</a></li> </ul>

## external-authority

---

<b>Syntax</b>	external-authority;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">pool-match-order</a> ], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">pool-match-order</a> ], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">pool-match-order</a> ], [edit system services dhcp-local-server <a href="#">pool-match-order</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	<p>Specify that an external authority (for example, RADIUS or Diameter) provides the address assignment.</p> <p>When RADIUS is the external authority, the router uses the Framed-IPv6-Pool attribute (RADIUS attribute 100) to select the pool. When Diameter is the external authority, the router uses the Diameter counterpart of RADIUS Framed-IPv6-Pool attribute.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 1442</a></li><li>• <a href="#">Extended DHCP Local Server Overview on page 1356</a></li><li>• <a href="#">Address-Assignment Pools Overview on page 1370</a></li></ul>

## failure-action

<b>Syntax</b>	failure-action (clear-binding   clear-binding-if-interface-up   log-only);
<b>Hierarchy Level</b>	[edit system services dhcp-local-server <a href="#">liveness-detection</a> ], [edit system services dhcp-local-server dhcpv6 <a href="#">liveness-detection</a> ], [edit forwarding-options dhcp-relay <a href="#">liveness-detection</a> ], [edit forwarding-options dhcp-relay dhcpv6 <a href="#">liveness-detection</a> ], [edit system services dhcp-local-server group <i>group-name</i> <a href="#">liveness-detection</a> ], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">liveness-detection</a> ], [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">liveness-detection</a> ], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">liveness-detection</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure the action the router (or switch) takes when a liveness detection failure occurs.
<b>Options</b>	<p><b>clear-binding</b>—The client session is cleared when a liveness detection failure occurs.</p> <p><b>clear-binding-if-interface-up</b>—The client session is cleared only when a liveness detection failure occurs and the local interface is detected as being up.</p> <p><b>log-only</b>—A message is logged to indicate the event; no action is taken and DHCP is left to manage the failure.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">DHCP Liveness Detection Overview on page 1472</a></li> <li>• <a href="#">Configuring Detection of DHCP Local Server Client Connectivity on page 1436</a></li> <li>• <a href="#">Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 1473</a></li> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## forward-snooped-clients (DHCP Local Server)

---

<b>Syntax</b>	forward-snooped-clients (all-interfaces   configured-interfaces   non-configured-interfaces);
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server</a> ], [edit logical-systems <i>logical-system-name</i> system services <a href="#">dhcp-local-server</a> ], [edit routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server</a> ], [edit system services <a href="#">dhcp-local-server</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure how the DHCP local server handles DHCP snooped packets on specific interfaces.
<b>Options</b>	<b>all-interfaces</b> —Perform the action on all interfaces.  <b>configured-interfaces</b> —Perform the action only on configured interfaces.  <b>non-configured-interfaces</b> —Perform the action only on nonconfigured interfaces.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">DHCP Snooping Support on page 1367</a></li><li>• <a href="#">Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 1437</a></li></ul>

## group (DHCP Local Server)

```

Syntax  group group-name {
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                relay-agent-interface-id
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
            primary-profile-name>;
        interface interface-name {
            exclude;
            overrides {
                client-discover-match (option60-and-option82 | incoming-interface);
                interface-client-limit number;
                process-inform {
                    pool pool-name;
                }
                rapid-commit;
            }
            service-profile dynamic-profile-name;
            trace;
            upto upto-interface-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                }
                detection-time {
                    threshold milliseconds;
                }
            }
            session-mode (automatic | multihop | singlehop);
        }
    }

```

```
        holddown-interval milliseconds;
    }
}
overrides {
    client-discover-match (option60-and-option82 | incoming-interface);
    delegated-pool;
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
route-suppression;
service-profile dynamic-profile-name;
}
```

<b>Hierarchy Level</b>	[edit system services <b>dhcp-local-server</b> ], [edit system services <b>dhcp-local-server dhcpv6</b> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server</b> ...], [edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server</b> ...], [edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server</b> ...]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	Configure a group of interfaces that have a common configuration, such as authentication parameters. A group must contain at least one interface.
<b>Options</b>	<b>group-name</b> —Name of the group.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

- Related Documentation**
- [Extended DHCP Local Server Overview on page 1356](#)
  - [Grouping Interfaces with Common DHCP Configurations on page 1421](#)
  - [Using External AAA Authentication Services with DHCP on page 1420](#)
  - *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
  - *Configuring a DHCP Server on Switches (CLI Procedure)*

## holddown-interval

<b>Syntax</b>	<code>holddown-interval <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure the time (in milliseconds) for which Bidirectional Forwarding Detection (BFD) holds a session up notification.
<b>Options</b>	<b><i>milliseconds</i></b> —Interval specifying how long a BFD session must remain up before a state change notification is sent. <b>Range:</b> 0 through 255,000 <b>Default:</b> 0
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <i>Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</i></li> </ul>

## interface (DHCP Local Server)

**Syntax** interface *interface-name* {  
 exclude;  
 overrides {  
   client-discover-match (option60-and-option82 | incoming-interface);  
   interface-client-limit *number*;  
   rapid-commit;  
 }  
 service-profile *dynamic-profile-name*;  
 trace;  
 upto *upto-interface-name*;  
}

**Hierarchy Level** [edit system services dhcp-local-server *group group-name*],  
 [edit system services dhcp-local-server *dhcpv6 group group-name*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system  
 services *dhcp-local-server ...*],  
 [edit logical-systems *logical-system-name* system services *dhcp-local-server ...*],  
 [edit routing-instances *routing-instance-name* system services *dhcp-local-server ...*]

**Release Information** Statement introduced in Junos OS Release 9.0.  
 Statement introduced in Junos OS Release 12.3R2 for EX Series switches.  
 Options **upto** and **exclude** introduced in Junos OS Release 9.1.

**Description** Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP local server is enabled. You can repeat the **interface *interface-name*** statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group. Also, you cannot use an interface that is being used by the DHCP relay agent.



**NOTE:** DHCP values are supported in Integrated Routing and Bridging (IRB) configurations. When you configure an IRB interface in a network that is using DHCP, the DHCP information (for example, authentication, address assignment, and so on) is propagated in the associated bridge domain. This enables the DHCP server to configure client IP addresses residing within the bridge domain. IRB currently only supports static DHCP configurations. For additional information about how to configure IRB, see *Configuring Integrated Routing and Bridging for Bridge Domains*.

**Options** **exclude**—Exclude an interface or a range of interfaces from the group. This option and the **overrides** option are mutually exclusive.

***interface-name***—Name of the interface. You can repeat this option multiple times.

***upto-interface-name***—Upper end of the range of interfaces; the lower end of the range is the *interface-name* entry. The interface device name of the ***upto-interface-name*** must be the same as the device name of the ***interface-name***.



The remaining statements are explained separately.

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Extended DHCP Local Server Overview on page 1356</a></li><li>• <a href="#">Grouping Interfaces with Common DHCP Configurations on page 1421</a></li><li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li></ul>

## interface-client-limit (DHCP Local Server)

<b>Syntax</b>	<code>interface-client-limit <i>number</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">dhcpv6 overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">group group-name overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">dhcpv6 overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <a href="#">group group-name overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">dhcpv6 overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">group group-name overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">dhcpv6 overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">group group-name overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server <a href="#">dhcpv6 overrides</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 <a href="#">group group-name overrides</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 <a href="#">group</a> interface <i>interface-name</i> <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Set the maximum number of DHCP subscribers or DHCP clients per interface allowed for a specific group or for all groups. A group specification takes precedence over a global specification for the members of that group.
<b>Default</b>	No limit
<b>Options</b>	<i>number</i> —Maximum number of clients allowed.

**Range:** 1 through 500,000

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Specifying the Maximum Number of DHCP Clients Per Interface on page 1424](#)
- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)

## interface-delete (Subscriber Management or DHCP Client Management)

**Syntax** interface-delete;

**Hierarchy Level** [edit system services subscriber-management maintain-subscriber]

**Release Information** Statement introduced in Junos OS Release 11.1.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** On router—Configure the router to maintain, rather than log out, subscribers when the subscriber interface is deleted. By default, the router logs out subscribers when the subscriber interface is deleted.

On switch—Configure the switch to maintain rather than log out DHCP clients when the client interface is deleted. By default, the switch logs out DHCP clients when the client interface is deleted.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events](#)

## interface-name (DHCP Local Server)

---

<b>Syntax</b>	interface-name;
<b>Hierarchy Level</b>	[edit system services <b>dhcp-local-server authentication username-include</b> ], [edit system services dhcp-local-server <b>dhcpv6 authentication username-include</b> ], [edit system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b> ], [edit system services dhcp-local-server <b>group group-name authentication username-include</b> ] [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server ...</b> ], [edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server ...</b> ], [edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server ...</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Specify that the interface name is concatenated with the username during the subscriber authentication or DHCP client authentication process. Use the statement at the <b>[edit ... dhcpv6]</b> hierarchy levels to configure DHCPv6 support.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li></ul>

## ip-address-first

<b>Syntax</b>	ip-address-first;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">pool-match-order</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">pool-match-order</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">pool-match-order</a>],</p> <p>[edit system services <a href="#">dhcp-local-server pool-match-order</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Configure the extended DHCP local server to use the IP address method to determine which address-assignment pool to use. The local server uses the IP address in the gateway IP address if one is present in the DHCP client PDU. If no gateway IP address is present, the local server uses the IP address of the receiving interface to find the address-assignment pool. The DHCP local server uses this method by default when no method is explicitly specified.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 1442</a></li> <li>• <a href="#">Extended DHCP Local Server Overview on page 1356</a></li> <li>• <a href="#">Address-Assignment Pools Overview on page 1370</a></li> <li>• <i>Configuring a DHCP Server on Switches (CLI Procedure)</i></li> </ul>

## liveness-detection

<b>Syntax</b>	<pre> liveness-detection {   failure-action (clear-binding   clear-binding-if-interface-up   log-only);   method {     bfd {       version (0   1   automatic);       minimum-interval milliseconds;       minimum-receive-interval milliseconds;       multiplier number;       no-adaptation;       transmit-interval {         minimum-interval milliseconds;         threshold milliseconds;       }       detection-time {         threshold milliseconds;       }       session-mode (automatic   multihop   singlehop);       holddown-interval milliseconds;     }   } } </pre>
<b>Hierarchy Level</b>	<pre> [edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6], [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit system services dhcp-local-server group group-name], [edit system services dhcp-local-server dhcpv6 group group-name], [edit forwarding-options dhcp-relay group group-name], [edit forwarding-options dhcp-relay dhcpv6 group group-name] </pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure bidirectional failure detection timers and authentication criteria for static routes.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">DHCP Liveness Detection Overview on page 1472</a></li> <li>• <a href="#">Configuring Detection of DHCP Local Server Client Connectivity on page 1436</a></li> <li>• <a href="#">Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 1473</a></li> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> </ul>

- *Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients*

## mac-address (DHCP Local Server)

<b>Syntax</b>	mac-address;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit system services dhcp-local-server <b>group group-name authentication username-include</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify that the MAC address from the client PDU be concatenated with the username during the subscriber authentication or DHCP client authentication process.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> </ul>

## method

---

<b>Syntax</b>	<pre>method {   bfd {     version (0   1   automatic);     minimum-interval milliseconds;     minimum-receive-interval milliseconds;     multiplier number;     no-adaptation;     transmit-interval {       minimum-interval milliseconds;       threshold milliseconds;     }     detection-time {       threshold milliseconds;     }     session-mode (automatic   multihop   singlehop);     holddown-interval milliseconds;   } }</pre>
<b>Hierarchy Level</b>	<pre>[edit system services dhcp-local-server <a href="#">liveness-detection</a>], [edit system services dhcp-local-server dhcpv6 <a href="#">liveness-detection</a>], [edit forwarding-options dhcp-relay <a href="#">liveness-detection</a>], [edit forwarding-options dhcp-relay dhcpv6 <a href="#">liveness-detection</a>], [edit system services dhcp-local-server group <i>group-name</i> <a href="#">liveness-detection</a>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">liveness-detection</a>], [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">liveness-detection</a>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">liveness-detection</a>]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure the liveness detection method.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li><li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li></ul>



## minimum-interval

<b>Syntax</b>	<code>minimum-interval <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the <a href="#">transmit-interval</a> <a href="#">minimal-interval</a> and <a href="#">minimum-receive-interval</a> statements.</p>
<b>Options</b>	<p><i>milliseconds</i> — Specify the minimum interval value for BFD liveliness detection.</p> <p><b>Range:</b> 1 through 255,000</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## minimum-receive-interval

---

<b>Syntax</b>	<code>minimum-receive-interval <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a>],</code> <code>[edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a>],</code> <code>[edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a>], [edit forwarding-options</code> <code>  dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a>],</code> <code>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],</code> <code>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method</code> <code>  <a href="#">bfd</a>],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],</code> <code>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method</code> <code>  <a href="#">bfd</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure the minimum interval at which the local routing device (or switch) must receive a reply from a neighbor with which it has established a BFD session.
<b>Options</b>	<i>milliseconds</i> — Specify the minimum receive interval value. <b>Range:</b> 1 through 255,000
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li><li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li></ul>

## multiplier

<b>Syntax</b>	<code>multiplier <i>number</i>;</code>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.  Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Configure the number of hello packets not received by the neighbor before Bidirectional Forwarding Detection (BFD) declares the neighbor down.
<b>Options</b>	<p><b>number</b>—Maximum allowable number of hello packets missed by the neighbor.  <b>Range:</b> 1 through 255  <b>Default:</b> 3</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.  routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## no-adaptation

---

<b>Syntax</b>	no-adaptation;
<b>Hierarchy Level</b>	[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li><li>• <i>Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</i></li></ul>

## option-60 (DHCP Local Server)

<b>Syntax</b>	option-60;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit system services dhcp-local-server authentication <a href="#">username-include</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify that the payload of Option 60 (Vendor Class Identifier) from the client PDU be concatenated with the username during the subscriber authentication or DHCP client authentication process.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> </ul>

## option-82 (DHCP Local Server Authentication)

---

<b>Syntax</b>	<code>option-82 &lt;circuit-id&gt; &lt;remote-id&gt;;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit system services dhcp-local-server authentication <a href="#">username-include</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Specify the type of Option 82 information from the client PDU that is concatenated with the username during the subscriber authentication or DHCP client authentication process. You can specify either, both, or neither of the Agent Circuit ID and Agent Remote ID suboptions. If you specify both, the Agent Circuit ID is supplied first, followed by a delimiter, and then the Agent Remote ID. If you specify that neither suboption is supplied, the raw payload of Option 82 from the PDU is concatenated to the username.</p>
<b>Options</b>	<p><b>circuit-id</b>—(Optional) Agent Circuit ID suboption (suboption 1).</p> <p><b>remote-id</b>—(Optional) Agent Remote ID suboption (suboption 2).</p>
<b>Required Privilege Level</b>	<p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li></ul>

## option-82 (DHCP Local Server Pool Matching)

<b>Syntax</b>	option-82;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">pool-match-order</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">pool-match-order</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">pool-match-order</a>],</p> <p>[edit system services dhcp-local-server <a href="#">pool-match-order</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure the extended DHCP local server to use the option 82 value in the DHCP client DHCP PDU together with the ip-address-first method to determine which address-assignment pool to use. You must configure the <b>ip-address-first</b> statement before configuring the <b>option-82</b> statement. The DHCP local server first determines which address-assignment pool to use based on the ip-address-first method. Then, the local server matches the option 82 value in the client PDU with the option 82 configuration in the address-assignment pool. This statement is supported for IPv4 address-assignment pools only.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 1442</a></li> <li>• <a href="#">Extended DHCP Local Server Overview on page 1356</a></li> <li>• <a href="#">Address-Assignment Pools Overview on page 1370</a></li> </ul>

## overrides (DHCP Local Server)

---

**Syntax** `overrides {  
 client-discover-match (option60-and-option82 | incoming-interface);  
 delegated-pool;  
 interface-client-limit number;  
 multi-address-embedded-option-response;  
 process-inform {  
 pool pool-name;  
 }  
 rapid-commit;  
}`

**Hierarchy Level** `[edit system services dhcp-local-server],  
[edit system services dhcp-local-server dhcpv6],  
[edit system services dhcp-local-server dhcpv6 group group-name],  
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],  
[edit system services dhcp-local-server group group-name],  
[edit system services dhcp-local-server group group-name interface interface-name],  
[edit logical-systems logical-system-name routing-instances routing-instance-name system  
services dhcp-local-server ...],  
[edit logical-systems logical-system-name system services dhcp-local-server ...],  
[edit routing-instances routing-instance-name system services dhcp-local-server ...]`

**Release Information** Statement introduced in Junos OS Release 9.2.  
Statement introduced in Junos OS Release 12.1 for EX Series switches.

**Description** Override the default configuration settings for the extended DHCP local server. Specifying the **overrides** statement with no subordinate statements removes all DHCP local server overrides at that hierarchy level.

- To override global DHCP local server configuration options, include the **overrides** statement and its subordinate statements at the `[edit system services dhcp-local-server]` hierarchy level.
- To override configuration options for a named group of interfaces, include the statements at the `[edit system services dhcp-local-server group group-name]` hierarchy level.
- To override configuration options for a specific interface within a named group of interfaces, include the statements at the `[edit system services dhcp-local-server group group-name interface interface-name]` hierarchy level.
- Use the `[edit system services dhcp-local-server dhcpv6]` hierarchy level to override DHCPv6 configuration options.

The remaining statements are explained separately.

The **interface-client-limit** statement is not supported in the `[edit system services dhcp-local-server dhcpv6]` hierarchy level.



The [delegated-pool](#), [multi-address-embedded-option-response](#), and the [rapid-commit](#) statements are supported in the `[edit system services dhcp-local-server dhcpv6 ...]` hierarchy level only.

<b>Required Privilege</b>	system—To view this statement in the configuration.
<b>Level</b>	system-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Extended DHCP Local Server Overview on page 1356</a></li><li>• <a href="#">Overriding Default DHCP Local Server Configuration Settings on page 1423</a></li><li>• <a href="#">Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429</a></li><li>• <a href="#">Configuring a DHCP Server on Switches (CLI Procedure)</a></li></ul>
------------------------------	--

## password (DHCP Local Server)

<b>Syntax</b>	<code>password password-string;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server authentication</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">dhcpv6 authentication</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">group group-name authentication</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">group group-name authentication</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <a href="#">dhcp-local-server authentication</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">dhcpv6 authentication</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <a href="#">group group-name authentication</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">group group-name authentication</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server authentication</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">dhcpv6 authentication</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">group group-name authentication</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">group group-name authentication</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server authentication</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">dhcpv6 authentication</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">group group-name authentication</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">group group-name authentication</a>],</p> <p>[edit system services <a href="#">dhcp-local-server authentication</a>],</p> <p>[edit system services dhcp-local-server <a href="#">dhcpv6</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 <a href="#">group group-name authentication</a>],</p> <p>[edit system services dhcp-local-server <a href="#">group group-name authentication</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Configure the password that is sent to the external AAA authentication server for subscriber authentication or DHCP client authentication.
<b>Options</b>	<i>password-string</i> —Authentication password.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> </ul>

## pool (DHCP Local Server Overrides)

**Syntax** `pool pool-name;`

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server overrides [process-inform](#)],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 overrides [process-inform](#)],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* overrides [process-inform](#)],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides [process-inform](#)],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* overrides [process-inform](#)],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* interface *interface-name* overrides [process-inform](#)],  
 [edit logical-systems *logical-system-name* system services dhcp-local-server overrides [process-inform](#)],  
 [edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 overrides [process-inform](#)],  
 [edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 group *group-name* overrides [process-inform](#)],  
 [edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides [process-inform](#)],  
 [edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* overrides [process-inform](#)],  
 [edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* interface *interface-name* overrides [process-inform](#)],  
 [edit routing-instances *routing-instance-name* system services dhcp-local-server overrides [process-inform](#)],  
 [edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 overrides [process-inform](#)],  
 [edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* overrides [process-inform](#)],  
 [edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides [process-inform](#)],  
 [edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* overrides [process-inform](#)],  
 [edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* interface *interface-name* overrides [process-inform](#)],  
 [edit system services dhcp-local-server overrides [process-inform](#)],  
 [edit system services dhcp-local-server dhcpv6 overrides [process-inform](#)],  
 [edit system services dhcp-local-server dhcpv6 group *group-name* overrides [process-inform](#)],  
 [edit system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides [process-inform](#)],  
 [edit system services dhcp-local-server group *group-name* overrides [process-inform](#)],  
 [edit system services dhcp-local-server group *group-name* interface *interface-name* overrides [process-inform](#)]

**Release Information** Statement introduced in Junos OS Release 11.4.  
 Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

<b>Description</b>	Configure DHCP or DHCPv6 local server to reply to DHCP information request messages (DHCPINFORM for DHCPv4 and INFORMATION-REQUEST for DHCPv6) with information taken from the specified pool without interacting with AAA.
<b>Options</b>	<b>pool-name</b> —Name of the address pool, which must be configured within <b>family inet</b> for DHCP local server and within <b>family inet6</b> for DHCPv6 local server.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling Processing of Client Information Requests on page 1427</a></li><li>• <a href="#">Overriding Default DHCP Local Server Configuration Settings on page 1423</a></li></ul>

---

## pool-match-order

---

<b>Syntax</b>	<pre>pool-match-order {   external-authority;   ip-address-first;   option-82; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server</b> ], [edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server</b> ], [edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server</b> ], [edit system services <b>dhcp-local-server</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Configure the order in which the DHCP local server uses information in the DHCP client PDU to determine how to obtain an address for the client.  The remaining statements are explained separately.
<b>Default</b>	DHCP local server uses the <b>ip-address-first</b> method to determine which address pool to use.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 1442</a></li><li>• <a href="#">Extended DHCP Local Server Overview on page 1356</a></li><li>• <a href="#">Configuring a DHCP Server on Switches (CLI Procedure)</a></li></ul>

## process-inform

<b>Syntax</b>	<pre>process-inform {     pool pool-name; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Enable the processing of DHCP information request messages (DHCPINFORM for DHCPv4 and INFORMATION-REQUEST for DHCPv6) sent from the client to request DHCP options. For DHCP local servers, the messages are also passed to the configured server list.</p>

The remaining statement is explained separately.

**Default** Information request messages are not processed.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Enabling Processing of Client Information Requests on page 1427](#)
- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\)](#)

## radius-disconnect (DHCP Local Server)

<b>Syntax</b>	radius-disconnect;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure <a href="#">trigger</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure <a href="#">trigger</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure <a href="#">trigger</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure <a href="#">trigger</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server reconfigure <a href="#">trigger</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 reconfigure <a href="#">trigger</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure <a href="#">trigger</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure <a href="#">trigger</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure <a href="#">trigger</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure <a href="#">trigger</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure <a href="#">trigger</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure <a href="#">trigger</a>],</p> <p>[edit system services dhcp-local-server reconfigure <a href="#">trigger</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 reconfigure <a href="#">trigger</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> reconfigure <a href="#">trigger</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure <a href="#">trigger</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <b>[edit ... dhcpv6 ...]</b> hierarchy levels introduced in Junos OS Release 10.4.</p>
<b>Description</b>	Configure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces to be reconfigured when a RADIUS-initiated disconnect is received by the DHCP client or group of clients. A group configuration takes precedence over a DHCP local server configuration.
<b>Default</b>	The client is deleted when a RADIUS-initiated disconnect is received.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429</a></li> <li>• <a href="#">Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 1433</a></li> </ul>

## rapid-commit (DHCPv6 Local Server)

---

<b>Syntax</b>	rapid-commit;
<b>Hierarchy Level</b>	[edit system services dhcp-local-server <a href="#">dhcpv6 overrides</a> ], [edit system services dhcp-local-server dhcpv6 <a href="#">group group-name overrides</a> ], [edit system services dhcp-local-server dhcpv6 <a href="#">group group-name interface interface-name overrides</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server dhcpv6 ...</a> ], [edit logical-systems <i>logical-system-name</i> system services <a href="#">dhcp-local-server dhcpv6 ...</a> ], [edit routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server dhcpv6 ...</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure DHCPv6 local server to recognize the Rapid Commit option (DHCPv6 option 14) in DHCPv6 solicit messages sent from the DHCPv6 client. When rapid commit is enabled for both DHCPv6 local server and the DHCPv6 client, a two-message handshake is used instead of the standard four-message handshake. You can enable rapid commit support on DHCPv6 local server globally, for a named group, or for a specific interface.
<b>Default</b>	Rapid commit support is not enabled.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling DHCPv6 Rapid Commit Support on page 1428</a></li><li>• <a href="#">Overriding Default DHCP Local Server Configuration Settings on page 1423</a></li></ul>



## reconfigure (DHCP Local Server)

<b>Syntax</b>	<pre>reconfigure {     attempts <i>attempt-count</i>;     clear-on-abort;     strict;     timeout <i>timeout-value</i>;     token <i>token-value</i>;     trigger {         radius-disconnect;     } }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>dhcpv6</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group</b> <i>group-name</i>],</p> <p>[edit system services <b>dhcp-local-server</b>],</p> <p>[edit system services dhcp-local-server <b>dhcpv6</b>],</p> <p>[edit system services dhcp-local-server <b>group</b> <i>group-name</i>],</p> <p>[edit system services dhcp-local-server dhcpv6 <b>group</b> <i>group-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Support at the [edit ... <b>dhcpv6</b> ...] hierarchy levels introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Enable dynamic reconfiguration triggered by the DHCP local server of all DHCP clients or only the DHCP clients serviced by the specified group of interfaces. A group configuration takes precedence over a DHCP local server configuration. The <b>strict</b> statement is available only for DHCPv6.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429](#)
  - [Configuring a DHCP Server on Switches \(CLI Procedure\)](#)

## relay-agent-interface-id (DHCP Local Server)

---

<b>Syntax</b>	relay-agent-interface-id;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify that the DHCPv6 Relay Agent Interface-ID option (option 18) in the client PDU name is concatenated with the username during the subscriber authentication or DHCP client authentication process.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li></ul>

## relay-agent-remote-id (DHCP Local Server)

<b>Syntax</b>	relay-agent-remote-id;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>For MX Series routers only, <b>enterprise-id</b> and <b>remote-id</b> options introduced in Junos OS Release 12.3R3.</p> <p>For MX Series routers only, the <b>enterprise-id</b> and <b>remote-id</b> options are obsoleted starting in Junos OS Releases 12.3R7, 13.2R4, 13.3R2, and 14.1R1.</p>
<b>Description</b>	Specify that the DHCPv6 Relay Agent Remote-ID option (option 37) in the client PDU name is concatenated with the username during the subscriber authentication or DHCP client authentication process. In order to generate an ASCII version of the username, the router concatenates only the remote-id portion of option 37 to the username, and ignores the enterprise number.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li> </ul>

## routing-instance-name (DHCP Local Server)

<b>Syntax</b>	routing-instance-name;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit system services dhcp-local-server <b>group group-name authentication username-include</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify that the routing instance name be concatenated with the username during the subscriber authentication or DHCP client authentication process. No routing instance name is concatenated if the configuration is in the default routing instance.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Using External AAA Authentication Services with DHCP on page 1420](#)

## service-profile (DHCP Local Server)

<b>Syntax</b>	<code>service-profile <i>dynamic-profile-name</i>;</code>
<b>Hierarchy Level</b>	<p>[edit system services <a href="#">dhcp-local-server</a>],  [edit system services dhcp-local-server <a href="#">dhcpv6</a>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">interface</a> <i>interface-name</i>],  [edit system services dhcp-local-server <a href="#">group</a> <i>group-name</i>],  [edit system services dhcp-local-server group <i>group-name</i> <a href="#">interface</a> <i>interface-name</i>],  [edit logical-systems <i>logical-system-name</i> system services <a href="#">dhcp-local-server</a> ...],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server</a> ...],  [edit routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server</a> ...]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Specify the default subscriber service or DHCP client management service, which is activated when the subscriber or client logs in and no other service is activated by a RADIUS server or a provisioning server.</p> <ul style="list-style-type: none"> <li>• To specify the default service for all DHCP local server clients, include the <b>service-profile</b> statement at the <b>[edit system services dhcp-local-server]</b> hierarchy level.</li> <li>• To specify the default service for a named group of interfaces, include the <b>service-profile</b> statement at the <b>[edit system services dhcp-local-server group <i>group-name</i>]</b> hierarchy level.</li> <li>• To specify the default service for a particular interface within a named group of interfaces, include the <b>service-profile</b> statement at the <b>[edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>]</b> hierarchy level.</li> <li>• For DHCPv6 clients, use the <b>service-profile</b> statement at the <b>[edit system services dhcp-local-server dhcpv6]</b> hierarchy level.</li> </ul>
<b>Options</b>	<b><i>dynamic-profile-name</i></b> —Name of the dynamic profile that defines the service.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Local Server Overview on page 1356</a></li> <li>• <a href="#">Default Subscriber Service Overview</a></li> <li>• <a href="#">Configuring a Default Subscriber Service</a></li> </ul>

## session-mode


---

<b>Syntax</b>	session-mode (automatic   multihop   singlehop);
<b>Hierarchy Level</b>	[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure the session mode.
<b>Options</b>	<b>automatic</b> —Configure single-hop BFD sessions if the peer is directly connected to the router interface and multihop BFD sessions if the peer is not directly connected to the router interface.  <b>multihop</b> —Configure multihop BFD sessions.  <b>single-hop</b> —Configure single hop BFD sessions.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li><li>• <i>Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</i></li></ul>

## strict (DHCP Local Server)


<b>Syntax</b>	strict;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify whether the server denies a client to bind when the client does not indicate that it accepts reconfigure messages. This feature is available only for DHCPv6.
<b>Default</b>	Accept solicit messages from clients that do not support reconfiguration and permit them to bind.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429</a></li> <li>• <a href="#">Preventing Binding of Clients That Do Not Support Reconfigure Messages on page 1434</a></li> </ul>

## threshold (detection-time)

<b>Syntax</b>	<code>threshold <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection method bfd <a href="#">detection-time</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd <a href="#">detection-time</a>],</p> <p>[edit forwarding-options dhcp-relay liveness-detection method bfd <a href="#">detection-time</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd <a href="#">detection-time</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd <a href="#">detection-time</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd <a href="#">detection-time</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd <a href="#">detection-time</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd <a href="#">detection-time</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify the threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.
<div>  <p><b>NOTE:</b> The threshold time must be greater than or equal to the <code>minimum-interval</code> or the <code>minimum-receive-interval</code>.</p> </div>	
<b>Options</b>	<p><i>milliseconds</i>— Value for the detection time adaptation threshold.</p> <p><b>Range:</b> 1 through 255,000</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>



## threshold (transmit-interval)

<b>Syntax</b>	<code>threshold <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit forwarding-options dhcp-relay liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Specify the threshold for detecting the adaptation of the transmit interval. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.
<b>Options</b>	<p><i>milliseconds</i> — Threshold value.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p>
<div>  <p><b>NOTE:</b> The threshold value specified in the <code>threshold</code> statement must be greater than the value specified in the <code>minimum-interval</code> statement for the <code>transmit-interval</code> statement.</p> </div>	
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## timeout (DHCP Local Server)

<b>Syntax</b>	<code>timeout <i>timeout-value</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 10.4.</p>
<b>Description</b>	Configure the initial value in seconds between attempts to reconfigure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces.
<b>Options</b>	<p><b><i>timeout-value</i></b>—Initial retry timeout value.</p> <p><b>Range:</b> 1 through 10 seconds</p> <p><b>Default:</b> 2 seconds</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429</a></li> <li><a href="#">Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 1432</a></li> </ul>

## token (DHCP Local Server)

<b>Syntax</b>	<code>token <i>token-value</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 10.4.</p>
<b>Description</b>	<p>Configure a plain-text token for all DHCP clients or only the DHCP clients serviced by the specified group of interfaces. The token enables rudimentary entity authentication to protect against inadvertently instantiated DHCP servers. A null token (empty string) indicates that the configuration token functionality is not enabled. A group configuration takes precedence over a DHCP local server configuration. For more information about tokens, see RFC 3118, <i>Authentication for DHCP Messages</i>, section 4.</p>
<b>Options</b>	<p><b><i>token-value</i></b>—Plain-text alphanumeric string.</p> <p><b>Default:</b> null (empty string)</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429</a></li> <li>• <a href="#">Configuring a Token for DHCP Local Server Authentication on page 1434</a></li> </ul>

## traceoptions (DHCP Server)

---

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regex</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i>; }</pre>
<b>Hierarchy Level</b>	[edit system services dhcp]
<b>Release Information</b>	Statement for tracing J Series Services Router DHCP processes introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Define tracing operations for DHCP processes for J Series Services Routers and EX Series switches.
<b>Options</b>	<p><b>file <i>filename</i></b>—Name of the file that receives the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory <i>/var/log</i>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>trace-file</i></b> reaches its maximum size, it is renamed <b><i>trace-file.0</i></b>, then <b><i>trace-file.1</i></b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option and a filename.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—All tracing operations</li><li>• <b>binding</b>—Trace binding operations</li><li>• <b>config</b>—Log reading of configuration</li><li>• <b>conflict</b>—Trace user-detected conflicts for IP addresses</li><li>• <b>event</b>—Trace important events</li><li>• <b>ifdb</b>—Trace interface database operations</li><li>• <b>io</b>— Trace I/O operations</li><li>• <b>lease</b>—Trace lease operations</li><li>• <b>main</b>—Trace main loop operations</li><li>• <b>misc</b>— Trace miscellaneous operations</li><li>• <b>packet</b>—Trace DHCP packets</li></ul>

- **options**—Trace DHCP options
- **pool**—Trace address pool operations
- **protocol**—Trace protocol operations
- **rtsock**—Trace routing socket operations
- **scope**—Trace scope operations
- **signal**—Trace DHCP signal operations
- **trace**—All tracing operations
- **ui**—Trace user interface operations

**match *regex***—(Optional) Refine the output to include lines that contain the regular expression.

- **all**—All tracing operations
- **binding**—Trace binding operations
- **config**—Log reading of configuration
- **conflict**—Trace user-detected conflicts for IP addresses
- **event**—Trace important events
- **ifdb**—Trace interface database operations
- **io**—Trace I/O operations
- **lease**—Trace lease operations
- **main**—Trace main loop operations
- **match *regex***—Refine the output to include lines that contain the regular expression.
- **misc**—Trace miscellaneous operations
- **packet**—Trace DHCP packets
- **options**—Trace DHCP options
- **pool**—Trace address pool operations
- **protocol**—Trace protocol operations
- **rtsock**—Trace routing socket operations
- **scope**—Trace scope operations
- **signal**—Trace DHCP signal operations
- **trace**—All tracing operations
- **ui**—Trace user interface operations

**no-world-readable**—(Optional) Disable unrestricted file access.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
---------------------------------	---

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Tracing Operations for DHCP Processes</i></li><li>• <i>System Management Configuration Statements</i></li></ul>
------------------------------	--

## transmit-interval

<b>Syntax</b>	<pre>transmit-interval {     threshold milliseconds;     minimum-interval milliseconds; }</pre>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a>],          [edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a>],          [edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a>], [edit forwarding-options          dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a>],          [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],          [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method  <a href="#">bfd</a>],          [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],          [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method  <a href="#">bfd</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.          Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure the Bidirectional Forwarding Detection (BFD) transmit interval.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.          routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## trigger (DHCP Local Server)

<b>Syntax</b>	<pre>trigger {     radius-disconnect; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> <a href="#">reconfigure</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">reconfigure</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <b>[edit ... dhcpv6 ...]</b> hierarchy levels introduced in Junos OS Release 10.4.</p>
<b>Description</b>	<p>Configure behavior in response to a trigger for all DHCP clients or only the DHCP clients serviced by the specified group of interfaces.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429</a></li> <li>• <a href="#">Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 1433</a></li> <li>• <a href="#">radius-disconnect on page 1527</a></li> </ul>



## use-primary (DHCP Local Server)

<b>Syntax</b>	<code>use-primary <i>primary-profile-name</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit system services dhcp-local-server <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify the dynamic profile to configure as the primary dynamic profile. The primary dynamic profile is instantiated when the first subscriber or DHCP client logs in. Subsequent subscribers (or clients) are not assigned the primary dynamic profile; instead, they are assigned the dynamic profile specified for the interface. When the first subscriber (or client) logs out, the next subscriber (or client) that logs in is assigned the primary dynamic profile.
<b>Options</b>	<b><i>primary-profile-name</i></b> —Name of the dynamic profile to configure as the primary dynamic profile
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</i></li> </ul>

## user-prefix (DHCP Local Server)

<b>Syntax</b>	<code>user-prefix <i>user-prefix-string</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <b>group group-name authentication username-include</b>],</p> <p>[edit system services <b>dhcp-local-server authentication username-include</b>],</p> <p>[edit system services dhcp-local-server <b>dhcpv6 authentication username-include</b>],</p> <p>[edit system services dhcp-local-server dhcpv6 <b>group group-name authentication username-include</b>],</p> <p>[edit system services dhcp-local-server <b>group group-name authentication username-include</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify the user prefix that is concatenated with the username during the subscriber authentication or DHCP client authentication process.
<b>Options</b>	<i>user-prefix-string</i> —User prefix string.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Using External AAA Authentication Services with DHCP on page 1420](#)

## username-include (DHCP Local Server)

<b>Syntax</b>	<pre>username-include {   circuit-type;   client-id;   delimiter <i>delimiter-character</i>;   domain-name <i>domain-name-string</i>;   interface-name;   logical-system-name;   mac-address;   option-60;   option-82 &lt;circuit-id&gt; &lt;remote-id&gt;;   relay-agent-interface-id;   relay-agent-remote-id;   relay-agent-subscriber-id;   routing-instance-name;   user-prefix <i>user-prefix-string</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit system services <a href="#">dhcp-local-server authentication</a>],  [edit system services dhcp-local-server <a href="#">dhcpv6 authentication</a>],  [edit system services dhcp-local-server dhcpv6 <a href="#">group group-name authentication</a>],  [edit system services dhcp-local-server <a href="#">group group-name authentication</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server ...</a>],  [edit logical-systems <i>logical-system-name</i> system services <a href="#">dhcp-local-server ...</a>],  [edit routing-instances <i>routing-instance-name</i> system services <a href="#">dhcp-local-server ...</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.  Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure the username that the router or switch passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router (or switch) accesses the local authentication service only and does not use external authentication services, such as RADIUS.</p> <p>The statements are explained separately. The <a href="#">option-60</a> and <a href="#">option-82</a> statements are not supported in the DHCPv6 hierarchy levels. The <a href="#">client-id</a>, <a href="#">relay-agent-interface-id</a>, <a href="#">relay-agent-remote-id</a> and <a href="#">relay-agent-subscriber-id</a> statements are supported in the DHCPv6 hierarchy levels only.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.  system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> <li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li> </ul>

## version (BFD)

---

<b>Syntax</b>	version (0   1   automatic);
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-liveness-detection], [edit system services dhcp-local-server liveness-detection method <i>bfd</i>], [edit system services dhcp-local-server dhcpv6 liveness-detection method <i>bfd</i>], [edit forwarding-options dhcp-relay liveness-detection method <i>bfd</i>], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <i>bfd</i>], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <i>bfd</i>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <i>bfd</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <i>bfd</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <i>bfd</i>], [edit protocols ldp oam bfd-liveness-detection], [edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Configure the BFD protocol version to detect.
<b>Options</b>	<p>0—Use BFD protocol version 0.</p> <p>1—Use BFD protocol version 1.</p> <p><b>automatic</b>—Autodetect the BFD protocol version.</p> <p><b>Default:</b> automatic</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li><li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li><li>• <a href="#">Configuring BFD for LDP LSPs</a></li></ul>

## DHCP Relay Agent Configuration Statements

---

- [\[edit forwarding-options dhcp-relay\] Configuration Statement Hierarchy for EX Series Switches on page 1550](#)
- [access \(Dynamic Access Routes\) on page 1556](#)
- [access-internal \(Dynamic Access-Internal Routes\) on page 1557](#)
- [active-server-group on page 1558](#)
- [allow-snooped-clients on page 1559](#)

- [always-write-giaddr](#) on page 1560
- [always-write-option-82](#) on page 1561
- [authentication \(DHCP Relay Agent\)](#) on page 1562
- [bfd](#) on page 1563
- [circuit-id \(DHCP Relay Agent\)](#) on page 1564
- [circuit-type \(DHCP Relay Agent\)](#) on page 1565
- [client-discover-match \(DHCP Relay Agent\)](#) on page 1566
- [client-id \(DHCP Relay Agent\)](#) on page 1567
- [delete-binding-on-renegotiation](#) on page 1567
- [delimiter \(DHCP Relay Agent\)](#) on page 1568
- [detection-time](#) on page 1569
- [dhcp-relay](#) on page 1570
- [dhcpcv6 \(DHCP Relay Agent\)](#) on page 1576
- [disable-relay](#) on page 1579
- [domain-name \(DHCP Relay Agent\)](#) on page 1580
- [drop \(DHCP Relay Agent Option\)](#) on page 1581
- [dynamic-profile \(DHCP Relay Agent\)](#) on page 1582
- [failure-action](#) on page 1583
- [forward-snooped-clients \(DHCP Relay Agent\)](#) on page 1584
- [group \(DHCP Relay Agent\)](#) on page 1585
- [holddown-interval](#) on page 1588
- [interface \(DHCP Relay Agent\)](#) on page 1589
- [interface-client-limit \(DHCP Relay Agent\)](#) on page 1591
- [interface-delete \(Subscriber Management or DHCP Client Management\)](#) on page 1592
- [interface-name \(DHCP Relay Agent\)](#) on page 1593
- [layer2-unicast-replies](#) on page 1594
- [liveness-detection](#) on page 1595
- [local-server-group \(DHCP Relay Agent Option\)](#) on page 1596
- [mac-address \(DHCP Relay Agent\)](#) on page 1597
- [method](#) on page 1598
- [minimum-interval](#) on page 1599
- [minimum-receive-interval](#) on page 1600
- [multiplier](#) on page 1601
- [next-hop \(Dynamic Access-Internal Routes\)](#) on page 1602
- [no-adaptation](#) on page 1603
- [no-allow-snooped-clients](#) on page 1604
- [no-bind-on-request \(DHCP Relay Agent\)](#) on page 1605

- [option-60 \(DHCP Relay Agent\) on page 1606](#)
- [option-82 \(DHCP Relay Agent\) on page 1607](#)
- [option-number \(DHCP Relay Agent Option\) on page 1608](#)
- [overrides \(DHCP Relay Agent\) on page 1609](#)
- [password \(DHCP Relay Agent\) on page 1611](#)
- [preference \(Subscriber Management\) on page 1612](#)
- [prefix \(DHCP Relay Agent\) on page 1613](#)
- [proxy-mode on page 1614](#)
- [relay-agent-interface-id \(DHCPv6 Relay Agent\) on page 1615](#)
- [relay-agent-remote-id \(DHCPv6 Relay Agent Username\) on page 1616](#)
- [relay-option \(DHCP Relay Agent\) on page 1617](#)
- [relay-option-82 on page 1618](#)
- [relay-server-group \(DHCP Relay Agent Option\) on page 1619](#)
- [replace-ip-source-with on page 1620](#)
- [routing-instance-name \(DHCP Relay Agent\) on page 1621](#)
- [send-release-on-delete \(DHCP Relay Agent\) on page 1622](#)
- [server-group on page 1623](#)
- [service-profile \(DHCP Relay Agent\) on page 1624](#)
- [session-mode on page 1625](#)
- [threshold \(detection-time\) on page 1626](#)
- [threshold \(transmit-interval\) on page 1627](#)
- [trace \(DHCP Relay Agent\) on page 1628](#)
- [transmit-interval on page 1629](#)
- [trust-option-82 on page 1630](#)
- [use-interface-description on page 1631](#)
- [use-primary \(DHCP Relay Agent\) on page 1633](#)
- [user-prefix \(DHCP Relay Agent\) on page 1634](#)
- [username-include \(DHCP Relay Agent\) on page 1635](#)
- [version \(BFD\) on page 1636](#)

## **[edit forwarding-options dhcp-relay] Configuration Statement Hierarchy for EX Series Switches**

This topic lists supported and unsupported configuration statements in the **[edit forwarding-options dhcp-relay]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.

- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit forwarding-options dhcp-relay\] Hierarchy Level on page 1551](#)
- [Unsupported Statements in the \[edit forwarding-options dhcp-relay\] Hierarchy Level on page 1555](#)

### Supported Statements in the [edit forwarding-options dhcp-relay] Hierarchy Level

The following hierarchy shows the **[edit forwarding-options dhcp-relay]** configuration statements supported on EX Series switches:

```
forwarding-options {
  dhcp-relay {
    active-server-group server-group-name;
    arp-inspection;
    authentication {
      ...
    }
  }
  authentication {
    password password-string;
  }
  username-include {
    circuit-type;
    client-id;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-name;
    mac-address;
    option-60;
    option-82 <circuit-id> <remote-id>;
    relay-agent-interface-id;
    relay-agent-remote-id;
    relay-agent-subscriber-id;
    routing-instance-name;
    user-prefix user-prefix-string;
  }
  dhcpv6 {
    active-server-group server-group-name;
    authentication {
      ...
    }
    dynamic-profile profile-name {
      aggregate-clients (merge | replace);
      use-primary primary-profile-name;
    }
    group group-name {
      ...
    }
  }
}
```

```
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    ...
  }
}
}
overrides {
  ...
}
relay-agent-interface-id;
relay-option {
  ...
}
server-group {
  server-group-name {
    server-ip-address;
  }
}
service-profile dynamic-profile-name;
OBSOLETE – duplicate-clients-on-interface;
dynamic-profile profile-name {
  aggregate-clients (merge | replace);
  use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
non-configured-interfaces);
group group-name {
  active-server-group server-group-name;
  authentication {
    ...
  }
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  interface interface-name {
    exclude;
    liveness-detection {
      ...
    }
  }
  overrides {
    ...
  }
  service-profile dynamic-profile-name;
  trace;
  upto upto-interface-name;
}
overrides {
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  client-discover-match <option60-and-option82>;
```



```

disable-relay;
interface-client-limit number;
layer2-unicast-replies;
no-allow-snooped-clients;
no-bind-on-request;
proxy-mode;
replace-ip-source-with;
send-release-on-delete;
trust-option-82;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
}
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
}
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
        }
    }
}

```

```
        detection-time {
            threshold milliseconds;
        }
        session-mode (automatic | multihop | singlehop);
        holddown-interval milliseconds;
    }
}
}
overrides {
    allow-snooped-clients;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
}
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
service-profile dynamic-profile-name;
}
```

### Unsupported Statements in the [edit forwarding-options dhcp-relay] Hierarchy Level

All statements in the [edit forwarding-options dhcp-relay] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 144: Unsupported [edit forwarding-options] Configuration Statements on EX Series Switches**

Statement	Hierarchy Level
-----------	-----------------


**NOTE:** Variables, such as *filename*, are not shown in the statements or hierarchies.

logical-system-name	[edit forwarding-options dhcp-relay authentication]
---------------------	---

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
  - [\[edit forwarding-options\] Configuration Statement Hierarchy on EX Series Switches on page 2698](#)

## access (Dynamic Access Routes)

---

Syntax	<pre>access {   route <i>prefix</i> {     <i>next-hop</i> <i>next-hop</i>;     metric <i>route-cost</i>;     <i>preference</i> <i>route-distance</i>;     tag <i>route-tag</i>;   } }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options], [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> ], [edit dynamic-profiles <i>profile-name</i> routing-options]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options] and [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> ] hierarchy levels introduced in Junos OS Release 10.1.
Description	Dynamically configure access routes.
<div> <b>BEST PRACTICE:</b> We recommend that you always include the <code>access-internal</code> stanza in the dynamic-profile when the <code>access</code> stanza is present for framed-route support.</div>	
Options	The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li><i>Configuring Dynamic Access Routes for Subscriber Management</i></li></ul>

## access-internal (Dynamic Access-Internal Routes)

<b>Syntax</b>	<pre>access-internal {   route <i>subscriber-ip-address</i> {     qualified-next-hop <i>underlying-interface</i> {       mac-address <i>address</i>;     }   } }</pre>
<b>Hierarchy Level</b>	<p>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options],</p> <p>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i>],</p> <p>[edit dynamic-profiles routing-options]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Support at the [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options] and [edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i>] hierarchy levels introduced in Junos OS Release 10.1.</p>
<b>Description</b>	<p>Dynamically configure access-internal routes. Access-internal routes are optional, but are used instead of access routes if the next-hop address is not specified in the Framed-Route Attribute [22] for IPv4 or the Framed-IPv6-Route attribute [99] for IPv6.</p>



**BEST PRACTICE:** We recommend that you always include the `access-internal` stanza in the dynamic-profile when the `access` stanza is present for framed-route support.


The remaining statements are explained separately.

<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management</i></li> <li>• <i>Configuring Dynamic Access-Internal Routes for PPP Subscriber Management</i></li> </ul>

## active-server-group

<b>Syntax</b>	<code>active-server-group <i>server-group-name</i>;</code>
<b>Hierarchy Level</b>	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay <b>dhcpv6</b>], [edit forwarding-options dhcp-relay <b>group</b> <i>group-name</i>], [edit forwarding-options dhcp-relay <b>group</b> <i>group-name</i> <b>dhcpv6</b>], [edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay</b>], [edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay</b> <b>dhcpv6</b>], [edit logical-systems <i>logical-system-name</i> forwarding-options <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options <b>group</b> <i>group-name</i> <b>dhcpv6</b>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   forwarding-options <b>dhcp-relay</b>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   forwarding-options <b>dhcp-relay</b> <b>dhcpv6</b>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   forwarding-options dhcp-relay <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   forwarding-options dhcp-relay <b>group</b> <i>group-name</i> <b>dhcpv6</b>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay] [edit routing-instances <i>routing-instance-name</i> forwarding-options <b>dhcp-relay</b> <b>dhcpv6</b>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <b>group</b>   <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <b>group</b>   <i>group-name</i>] </pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Apply a DHCP relay agent configuration to the named group of DHCP server addresses. Use the statement at the <b>[edit ... dhcpv6]</b> hierarchy levels to configure DHCPv6 support.</p> <p>A group-specific configuration overrides a global option.</p>
<b>Options</b>	<b><i>server-group-name</i></b> —Name of the group of DHCP or DHCPv6 server addresses to which the DHCP or DHCPv6 relay agent configuration applies.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> <li>• <a href="#">Configuring Active Server Groups on page 1469</a></li> <li>• <a href="#">Configuring Group-Specific DHCP Relay Options on page 1384</a></li> <li>• <a href="#">dhcp-relay on page 1570</a></li> </ul>

## allow-snooped-clients

<b>Syntax</b>	allow-snooped-clients;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay <b>dhcpv6</b> group <i>group-name</i> interface <i>interface-name</i> <b>overrides</b>],</p> <p>[edit forwarding-options dhcp-relay <b>dhcpv6</b> group <i>group-name</i> <b>overrides</b>],</p> <p>[edit forwarding-options dhcp-relay <b>dhcpv6</b> <b>overrides</b>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <b>overrides</b>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> <b>overrides</b>],</p> <p>[edit forwarding-options dhcp-relay <b>overrides</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay</b> ...],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <b>dhcp-relay</b> ...],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options <b>dhcp-relay</b> ...]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Support at the [edit ... <b>dhcpv6</b>] hierarchy levels introduced in Junos OS Release 12.1.</p>
<b>Description</b>	<p>Explicitly enable DHCP snooping support on the router.</p> <p>Use the statement at the [edit ... <b>dhcpv6</b>] hierarchy levels to explicitly enable snooping support on the router for DHCPv6 relay agent.</p>
<div>  <b>NOTE:</b> DHCP snooping is <i>disabled</i> by default.         </div>	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> <li>• <a href="#">Overriding the Default DHCP Relay Configuration Settings on page 1447</a></li> <li>• <a href="#">DHCP Snooping Support on page 1367</a></li> </ul>

## always-write-giaddr

---

<b>Syntax</b>	<code>always-write-giaddr;</code>
<b>Hierarchy Level</b>	<code>[edit forwarding-options dhcp-relay <a href="#">overrides</a>],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	Overwrite the gateway IP address (giaddr) of every DHCP packet with the giaddr of the DHCP relay agent before forwarding the packet to the DHCP server.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li><li>• <a href="#">dhcp-relay on page 1570</a></li></ul>



## always-write-option-82

<b>Syntax</b>	<code>always-write-option-82;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay <a href="#">overrides</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.  Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Override the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server. The use of this option causes the DHCP relay agent to perform one of the following actions, depending on how it is configured:</p> <ul style="list-style-type: none"> <li>• If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server.</li> <li>• If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.</li> </ul>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.  interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> </ul>

## authentication (DHCP Relay Agent)

---

Syntax	<pre>authentication {   password <i>password-string</i>;   username-include {     circuit-type;     client-id;     delimiter <i>delimiter-character</i>;     domain-name <i>domain-name-string</i>;     interface-name;     logical-system-name;     mac-address;     option-60;     option-82 &lt;circuit-id&gt; &lt;remote-id&gt;;     relay-agent-interface-id;     relay-agent-remote-id;     relay-agent-subscriber-id;     routing-instance-name;     user-prefix <i>user-prefix-string</i>;   } }</pre>
Hierarchy Level	<pre>[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay <b>dhcpv6</b>], [edit forwarding-options dhcp-relay dhcpv6 <b>group</b> <i>group-name</i>], [edit forwarding-options dhcp-relay <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.</p>
Description	<p>Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay configuration. Use the statement at the <b>[edit...dhcpv6]</b> hierarchy levels to configure DHCPv6 support.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">dhcp-relay on page 1570</a></li><li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li></ul>

## bfd

<b>Syntax</b>	<pre> bfd {   version (0   1   automatic);   minimum-interval <i>milliseconds</i>;   minimum-receive-interval <i>milliseconds</i>;   multiplier <i>number</i>;   no-adaptation;   transmit-interval {     minimum-interval <i>milliseconds</i>;     threshold <i>milliseconds</i>;   }   detection-time {     threshold <i>milliseconds</i>;   }   session-mode (automatic   multihop   singlehop);   holddown-interval <i>milliseconds</i>; } </pre>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection <a href="#">method</a>],  [edit system services dhcp-local-server dhcpv6 liveness-detection <a href="#">method</a>],  [edit forwarding-options dhcp-relay liveness-detection <a href="#">method</a>],  [edit forwarding-options dhcp-relay dhcpv6 liveness-detection <a href="#">method</a>],  [edit system services dhcp-local-server group <i>group-name</i> liveness-detection <a href="#">method</a>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection <a href="#">method</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection <a href="#">method</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection <a href="#">method</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.  Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure Bidirectional Forwarding Detection (BFD) as the liveness detection method.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.  routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>


## circuit-id (DHCP Relay Agent)

<b>Syntax</b>	<pre> circuit-id {   prefix <i>prefix</i>;   use-interface-description (logical   device);   use-vlan-id; } </pre>
<b>Hierarchy Level</b>	<pre> [edit forwarding-options dhcp-relay <a href="#">relay-option-82</a>], [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">relay-option-82</a>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">relay-option-82</a>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">relay-option-82</a>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">relay-option-82</a>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">relay-option-82</a>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">relay-option-82</a>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">relay-option-82</a>] </pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
<b>Description</b>	<p>Specify the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option (option 82) to include in DHCP packets destined for a DHCP server. Optionally specify that the suboption includes a prefix or textual description, or VLAN tag.</p> <p>The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual LANs (VLANs) or stacked VLANs (S-VLANs) is as follows:</p> <pre>(fe   ge)-<i>fpc/pic/port</i></pre> <p>The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that use VLANs is as follows:</p> <pre>(fe   ge)-<i>fpc/pic/port:vlan-id</i></pre> <p>The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs is as follows:</p> <pre>(fe   ge)-<i>fpc/pic/port:svlan-id-vlan-id</i></pre> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using DHCP Relay Agent Option 82 Information on page 1464</a></li> <li>• <a href="#">Configuring Option 82 Information on page 1465</a></li> </ul>

## circuit-type (DHCP Relay Agent)

<b>Syntax</b>	circuit-type;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay ...</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay ...</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay ...</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels introduced in Junos OS Release 11.4.</p>
<b>Description</b>	Specify that the circuit type is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels to configure DHCPv6 support.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> <li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li> </ul>

## client-discover-match (DHCP Relay Agent)

<b>Syntax</b>	client-discover-match (option60-and-option82   incoming-interface);
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ... <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ... <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group ... <a href="#">overrides</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Option <b>incoming-interface</b> introduced in Junos OS Release 13.3.</p>
<b>Description</b>	Configure the match criteria DHCP relay uses to uniquely identify DHCP subscribers or clients when primary identification fails. The options are mutually exclusive.
<b>Options</b>	<p><b>incoming-interface</b>—Allow only one client device to connect on the interface. If the client device changes, the router deletes the existing client binding and creates a binding for the newly connected device.</p>
<div>  <p><b>NOTE:</b> The <b>overrides client-discover-match incoming-interface</b> configuration deletes and replaces the existing binding when a new device connects. This action differs from the <b>overrides interface-client-limit 1</b> statement, which retains the existing binding rejects the newly connected client.</p> </div>	
<p><b>option60-and-option82</b>—Use option 60 and option 82 information to identify subscribers.</p>	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> <li>• <a href="#">Overriding the Default DHCP Relay Configuration Settings on page 1447</a></li> <li>• <a href="#">DHCP Auto Logout Overview on page 1368</a></li> <li>• <a href="#">Allowing Only One DHCP Client Per Interface</a></li> </ul>

## client-id (DHCP Relay Agent)

<b>Syntax</b>	client-id;
<b>Hierarchy Level</b>	[edit forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a> ], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a> ], [edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay dhcpv6</a> ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay dhcpv6</a> ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay dhcpv6</a> ...]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Specify that the client ID is concatenated with the username during the subscriber authentication or client authentication process.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> <li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li> </ul>

## delete-binding-on-renegotiation

<b>Syntax</b>	delete-binding-on-renegotiation;
<b>Hierarchy Level</b>	[edit forwarding-options dhcp-relay <a href="#">overrides</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2 for EX Series switches.
<b>Description</b>	Configure the DHCP relay agent to delete binding information for a specific client when a DHCP DISCOVER packet is received from the client while the client already has a binding on the relay that is in BOUND state. A DHCP client sends discover messages to renegotiate the lease for an IP address.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> </ul>

## delimiter (DHCP Relay Agent)

<b>Syntax</b>	<code>delimiter <i>delimiter-character</i>;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels introduced in Junos OS Release 11.4.</p>
<b>Description</b>	Specify the character used as the delimiter between the concatenated components of the username. Use the statement at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels to configure DHCPv6 support.
<b>Options</b>	<b><i>delimiter-character</i></b> —Character that separates components that make up the concatenated username. You cannot use the semicolon (;) as a delimiter.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> <li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li> </ul>



## detection-time

<b>Syntax</b>	<pre> detection-time {     threshold milliseconds; } </pre>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server <a href="#">liveness-detection</a> method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 <a href="#">liveness-detection</a> method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay <a href="#">liveness-detection</a> method <a href="#">bfd</a>], [edit forwarding-options  dhcp-relay dhcpv6 <a href="#">liveness-detection</a> method <a href="#">bfd</a>],  [edit system services dhcp-local-server group <i>group-name</i> <a href="#">liveness-detection</a> method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">liveness-detection</a> method  <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">liveness-detection</a> method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">liveness-detection</a> method  <a href="#">bfd</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.  Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Enable failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.  routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## dhcp-relay

---

```
Syntax  dhcp-relay {
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
    }
    dhcpv6 {
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        group group-name {
            active-server-group server-group-name;
            authentication {
                ...
            }
            dynamic-profile profile-name {
                ...
            }
            interface interface-name {
                exclude;
                liveness-detection {
                    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                    method {
```

```

bfd {
  version (0 | 1 | automatic);
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  detection-time {
    threshold milliseconds;
  }
  session-mode (automatic | multihop | singlehop);
  holddown-interval milliseconds;
}
}
}
overrides {
  ...
}
relay-option {
  ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
route-suppression:
service-profile dynamic-profile-name;
overrides {
  ...
}
relay-agent-interface-id {
  ...
}
relay-agent-remote-id {
  ...
}
relay-option {
  ...
}
route-suppression;
server-response-time seconds;
service-profile dynamic-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {

```

```
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    detection-time {
        threshold milliseconds;
    }
    session-mode (automatic | multihop | singlehop);
    holddown-interval milliseconds;
}
}
}
overrides {
    allow-snooped-clients;
    delay-authentication;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    prefix prefix;
    use-interface-description (logical | device);
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
duplicate-clients-in-subnet (incoming-interface | option-82);
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    active-server-group server-group-name;
    authentication {
        ...
    }
    dynamic-profile profile-name {
        ...
    }
}
interface interface-name {
    exclude;
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
            }
        }
    }
}
```

```

        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        detection-time {
            threshold milliseconds;
        }
        session-mode(automatic | multihop | singlehop);
        holddown-interval milliseconds;
    }
}
overrides {
    ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
overrides {
    ...
}
relay-option {
    ...
}
relay-option-82 {
    ...
}
route-suppression:
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {

```

```
allow-snooped-clients;
always-write-giaddr;
always-write-option-82;
client-discover-match (option60-and-option82 | incoming-interface);
delay-authentication;
disable-relay;
interface-client-limit number;
layer2-unicast-replies;
no-allow-snooped-clients;
no-bind-on-request;
proxy-mode;
replace-ip-source-with;
send-release-on-delete;
trust-option-82;
}
relay-option {
  option-number option-number;
  default-action {
    drop;
    forward-only;
    relay-server-group group-name;
  }
  equals (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    relay-server-group relay-server-group;
  }
  starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    local-server-group local-server-group;
    relay-server-group relay-server-group;
  }
}
relay-option-82 {
  circuit-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
  remote-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
}
server-group {
  server-group-name {
    server-ip-address;
  }
}
route-suppression:
server-response-time seconds;
service-profile dynamic-profile-name;
}
```

<b>Hierarchy Level</b>	<p>[edit forwarding-options],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51 for the QFX Series.</p>
<b>Description</b>	<p>Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch and enable the router (or switch) to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.</p> <p>DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.</p> <p>The extended DHCP and DHCPv6 relay agent options configured with the <b>dhcpr-relay</b> and <b>dhcprv6</b> statements are incompatible with the DHCP/BOOTP relay agent options configured with the <b>bootp</b> statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router (or switch) at the same time.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> <li>• <a href="#">DHCPv6 Relay Agent Overview on page 1384</a></li> <li>• <a href="#">DHCP Relay Proxy Overview on page 1382</a></li> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> </ul>

## dhcpcv6 (DHCP Relay Agent)

```
Syntax  dhcpcv6 {
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            logical-system-name;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    group group-name {
        active-server-group server-group-name;
        authentication {
            ...
        }
        dynamic-profile profile-name {
            ...
        }
    }
    interface interface-name {
        exclude;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
        }
    }
}
```



```

    overrides {
        ...
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
}
}
overrides {
    ...
}
relay-agent-interface-id {
    ...
}
relay-agent-remote-id {
    ...
}
relay-option {
    ...
}
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {
    ...
}
}
overrides {
    allow-snooped-clients;
    delay-authentication;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    prefix prefix;
    use-interface-description (logical | device);
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
}
equals (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    relay-server-group relay-server-group;
}
starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {

```

```
    drop;
    forward-only;
    relay-server-group relay-server-group;
  }
}
server-group {
  server-group-name {
    server-ip-address;
  }
}
route-suppression;
server-response-time seconds;
service-profile dynamic-profile-name;
}
```

**Hierarchy Level** [edit forwarding-options dhcp-relay],  
[edit logical-systems *logical-system-name* forwarding-options **dhcp-relay**],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*  
forwarding-options dhcp-relay],  
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay]

**Release Information** Statement introduced in Junos OS Release 11.4.  
Statement introduced in Junos OS Release 12.3 for EX Series switches.

**Description** Configure DHCPv6 relay options on the router or switch and enable the router or switch to function as a DHCPv6 relay agent. A DHCPv6 relay agent forwards DHCPv6 request and reply packets between a DHCPv6 client and a DHCPv6 server.

The DHCPv6 relay agent server is fully compatible with the extended DHCP local server and DHCP relay agent. However, the options configured with the **dhcpv6** statement are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, the DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot be enabled on the router or switch at the same time.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [dhcp-relay on page 1570](#)
- [DHCPv6 Relay Agent Overview on page 1384](#)
- [Using External AAA Authentication Services with DHCP on page 1420](#)

## disable-relay

<b>Syntax</b>	disable-relay;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Disable DHCP relay on specific interfaces in a group.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> </ul>

## domain-name (DHCP Relay Agent)

<b>Syntax</b>	<code>domain-name <i>domain-name-string</i>;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels introduced in Junos OS Release 11.4.</p>
<b>Description</b>	Specify the domain name that is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels to configure DHCPv6 support.
<b>Options</b>	<i>domain-name-string</i> —Domain name formatted string.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> <li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li> </ul>

## drop (DHCP Relay Agent Option)

<b>Syntax</b>	drop;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay relay-option (default-action   equals   starts-with)],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action   equals   starts-with)],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> relay-option (default-action   equals   starts-with)],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> relay-option (default-action   equals   starts-with)],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay</a> ...],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a> ...],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a> ...]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3.
<b>Description</b>	Drop (discard) specified DHCP client packets when you use DHCP relay agent selective processing. You can configure the drop operation globally or for a group of interfaces, and for either DHCP or DHCPv6 relay agent.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Using DHCP Option Information to Selectively Process DHCP Client Traffic</i></li> </ul>

## dynamic-profile (DHCP Relay Agent)

---

<b>Syntax</b>	<pre>dynamic-profile <i>profile-name</i> {     aggregate-clients (merge   replace);     use-primary <i>primary-profile-name</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit forwarding-options dhcp-relay], [edit forwarding-options <b>dhcp-relay dhcpv6</b>], [edit forwarding-options dhcp-relay dhcpv6 <b>group group-name</b>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <b>interface interface-name</b>], [edit forwarding-options dhcp-relay <b>group group-name</b>], [edit forwarding-options dhcp-relay group <i>group-name</i> <b>interface interface-name</b>], [edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay ...</b>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>     forwarding-options <b>dhcp-relay ...</b>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Specify the dynamic profile that is attached to all interfaces, to a named group of interfaces, or to a specific interface.</p> <p>M120 and M320 routers do not support DHCPv6.</p>
<b>Options</b>	<p><b><i>profile-name</i></b>—Name of the dynamic profile.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">dhcp-relay on page 1570</a></li><li>• <i>Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</i></li><li>• <a href="#">Grouping Interfaces with Common DHCP Configurations on page 1421</a></li><li>• <i>Configuring a Default Subscriber Service</i></li></ul>

## failure-action

<b>Syntax</b>	failure-action (clear-binding   clear-binding-if-interface-up   log-only);
<b>Hierarchy Level</b>	[edit system services dhcp-local-server <a href="#">liveness-detection</a> ], [edit system services dhcp-local-server dhcpv6 <a href="#">liveness-detection</a> ], [edit forwarding-options dhcp-relay <a href="#">liveness-detection</a> ], [edit forwarding-options dhcp-relay dhcpv6 <a href="#">liveness-detection</a> ], [edit system services dhcp-local-server group <i>group-name</i> <a href="#">liveness-detection</a> ], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <a href="#">liveness-detection</a> ], [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">liveness-detection</a> ], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">liveness-detection</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure the action the router (or switch) takes when a liveness detection failure occurs.
<b>Options</b>	<p><b>clear-binding</b>—The client session is cleared when a liveness detection failure occurs.</p> <p><b>clear-binding-if-interface-up</b>—The client session is cleared only when a liveness detection failure occurs and the local interface is detected as being up.</p> <p><b>log-only</b>—A message is logged to indicate the event; no action is taken and DHCP is left to manage the failure.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">DHCP Liveness Detection Overview on page 1472</a></li> <li>• <a href="#">Configuring Detection of DHCP Local Server Client Connectivity on page 1436</a></li> <li>• <a href="#">Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 1473</a></li> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## forward-snooped-clients (DHCP Relay Agent)

---

<b>Syntax</b>	forward-snooped-clients (all-interfaces   configured-interfaces   non-configured-interfaces);
<b>Hierarchy Level</b>	[edit forwarding-options <a href="#">dhcp-relay</a> ], [edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	<p>Configure how DHCP relay agent handles DHCP snooped packets on specific interfaces. The router or switch determines the DHCP snooping action to perform based on a combination of the <b>forward-snooped-clients</b> configuration and the configuration of either the <b>allow-snooped-clients</b> statement or the <b>no-allow-snooped-clients</b> statement.</p> <p>The router (or switch) also uses this statement to determine how to handle snooped BOOTREPLY packets received on nonconfigured interfaces.</p>
<b>Options</b>	<p><b>all-interfaces</b>—Perform the action on all interfaces.</p> <p><b>configured-interfaces</b>—Perform the action only on configured interfaces.</p> <p><b>non-configured-interfaces</b>—Perform the action only on nonconfigured interfaces.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">DHCP Snooping Support on page 1367</a></li><li>• <a href="#">Configuring DHCP Snooping for DHCP Relay Agent on page 1455</a></li></ul>



## group (DHCP Relay Agent)

```
Syntax  group group-name {
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                logical-system-name;
                mac-address;
                option-60;
                option-82 [circuit-id] [remote-id];
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        interface interface-name {
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                    }
                    session-mode (automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
            }
        }
        overrides {
            ...
        }
        service-profile dynamic-profile-name;
        trace;
```

```
    upto upto-interface-name;
}
overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match (option60-and-option82 | incoming-interface);
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    prefix prefix;
    use-interface-description (logical | device);
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82;
    }
    remote-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
}
```

```

    }
  }
  route-suppression;
  service-profile dynamic-profile-name;
}

```

<b>Hierarchy Level</b>	[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay <b>dhcpv6</b> ], [edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay</b> ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <b>dhcp-relay</b> ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Support at the [edit ... <b>dhcpv6</b> ] hierarchy levels introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	Specify the name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration. A group must contain at least one interface. Use the statement at the [edit ... <b>dhcpv6</b> ] hierarchy levels to configure DHCPv6 support.
<b>Options</b>	<p><b>group-name</b>—Name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">dhcp-relay on page 1570</a></li> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> <li>• <a href="#">Configuring Group-Specific DHCP Relay Options on page 1384</a></li> <li>• <a href="#">Grouping Interfaces with Common DHCP Configurations on page 1421</a></li> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> <li>• <a href="#">Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</a></li> </ul>

## holddown-interval

---

<b>Syntax</b>	<code>holddown-interval <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a>],</code> <code>[edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a>],</code> <code>[edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a>], [edit forwarding-options</code> <code>  dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a>],</code> <code>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],</code> <code>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method</code> <code>  <a href="#">bfd</a>],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],</code> <code>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method</code> <code>  <a href="#">bfd</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure the time (in milliseconds) for which Bidirectional Forwarding Detection (BFD) holds a session up notification.
<b>Options</b>	<b><i>milliseconds</i></b> —Interval specifying how long a BFD session must remain up before a state change notification is sent. <b>Range:</b> 0 through 255,000 <b>Default:</b> 0
<b>Required Privilege Level</b>	<b>routing</b> —To view this statement in the configuration. <b>routing-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li><li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li></ul>

## interface (DHCP Relay Agent)

<b>Syntax</b>	<pre> interface <i>interface-name</i> {   exclude;   overrides {     allow-snooped-clients;     always-write-giaddr;     always-write-option-82;     client-discover-match (option60-and-option82   incoming-interface);     disable-relay;     interface-client-limit <i>number</i>;     layer2-unicast-replies;     no-allow-snooped-clients;     proxy-mode;     replace-ip-source-with;     send-release-on-delete;     trust-option-82;   }   service-profile <i>dynamic-profile-name</i>;   trace;   upto <i>upto-interface-name</i>; } </pre>
<b>Hierarchy Level</b>	<pre> [edit forwarding-options dhcp-relay dhcpv6 <b>group</b> <i>group-name</i>], [edit forwarding-options dhcp-relay <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay</b> ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>  forwarding-options <b>dhcp-relay</b> ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...] </pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Options <b>upto</b> and <b>exclude</b> introduced in Junos OS Release 9.1.</p> <p>Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP or DHCPv6 relay agent is enabled. You can repeat the <b>interface <i>interface-name</i></b> statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group. Also, you cannot use an interface that is being used by the DHCP local server. Use the statement at the <b>[edit ... dhcpv6]</b> hierarchy levels to configure DHCPv6 support.</p> <p>EX Series switches do not support DHCPv6.</p>



**NOTE:** DHCP values are supported in Integrated Routing and Bridging (IRB) configurations. When you configure an IRB interface in a network that is using DHCP, the DHCP information (for example, authentication, address assignment, and so on) is propagated in the associated bridge domain. This enables the DHCP server to configure client IP addresses residing within the bridge domain. IRB currently only supports static DHCP configurations. For

additional information about how to configure IRB, see *Configuring Integrated Routing and Bridging for Bridge Domains*.

.....

**Options**    **exclude**—Exclude an interface or a range of interfaces from the group. This option and the **overrides** option are mutually exclusive.

**interface-name**—Name of the interface. You can repeat this option multiple times.

**overrides**—Override the specified default configuration settings for the interface. The **overrides** statement is described separately.

**upto-interface-name**—Upper end of the range of interfaces; the lower end of the range is the interface-name entry. The interface device name of the **upto-interface-name** must be the same as the device name of the **interface-name**.

The remaining statements are explained separately.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                     interface-control—To add this statement to the configuration.

**Related Documentation**

- [Extended DHCP Relay Agent Overview on page 1379](#)
- [dhcp-relay on page 1570](#)
- *dhcp-relay (EX Series Switches only)*
- *Understanding the Extended DHCP Relay Agent for EX Series Switches*
- [Configuring an Extended DHCP Relay Server on EX Series Switches \(CLI Procedure\) on page 1416](#)
- [Grouping Interfaces with Common DHCP Configurations on page 1421](#)
- [Using External AAA Authentication Services with DHCP on page 1420](#)

## interface-client-limit (DHCP Relay Agent)

<b>Syntax</b>	<code>interface-client-limit <i>number</i>;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],</p> <p>[edit forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Set the maximum number of DHCP (or DHCPv6) subscribers or clients per interface allowed for a specific group or for all groups. A group specification takes precedence over a global specification for the members of that group. Use the statement at the <b>[edit ... dhcpv6]</b> hierarchy levels to configure DHCPv6 support.</p> <p>M120 and M320 routers do not support DHCPv6.</p>
<b>Default</b>	No limit
<b>Options</b>	<p><i>number</i>—Maximum number of clients allowed.</p> <p><b>Range:</b> 1 through 500,000</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

- Related Documentation**
- [dhcp-relay on page 1570](#)
  - [Extended DHCP Relay Agent Overview on page 1379](#)
  - [Configuring Group-Specific DHCP Relay Options on page 1384](#)
  - [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)

---

## interface-delete (Subscriber Management or DHCP Client Management)

---

<b>Syntax</b>	interface-delete;
<b>Hierarchy Level</b>	[edit system services subscriber-management maintain-subscriber]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	<p>On router—Configure the router to maintain, rather than log out, subscribers when the subscriber interface is deleted. By default, the router logs out subscribers when the subscriber interface is deleted.</p> <p>On switch—Configure the switch to maintain rather than log out DHCP clients when the client interface is deleted. By default, the switch logs out DHCP clients when the client interface is deleted.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <a href="#">Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events</a>



## interface-name (DHCP Relay Agent)

<b>Syntax</b>	interface-name;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay authentication <a href="#">username-include</a>],  [edit forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],  [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.4</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify that the interface name is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li> </ul>

## layer2-unicast-replies

---

<b>Syntax</b>	layer2-unicast-replies;
<b>Hierarchy Level</b>	[edit forwarding-options dhcp-relay <a href="#">overrides</a> ], [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a> ], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">overrides</a> ], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	Override the setting of the broadcast bit in DHCP request packets and instead use the Layer 2 unicast transmission method to transmit DHCP Offer reply packets and DHCP ACK reply packets from the DHCP server to DHCP clients during the discovery process.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li><li>• <a href="#">dhcp-relay on page 1570</a></li></ul>

## liveness-detection

<b>Syntax</b>	<pre> liveness-detection {   failure-action (clear-binding   clear-binding-if-interface-up   log-only);   method {     bfd {       version (0   1   automatic);       minimum-interval <i>milliseconds</i>;       minimum-receive-interval <i>milliseconds</i>;       multiplier <i>number</i>;       no-adaptation;       transmit-interval {         minimum-interval <i>milliseconds</i>;         threshold <i>milliseconds</i>;       }       detection-time {         threshold <i>milliseconds</i>;       }       session-mode (automatic   multihop   singlehop);       holddown-interval <i>milliseconds</i>;     }   } } </pre>
<b>Hierarchy Level</b>	<pre> [edit system services <a href="#">dhcp-local-server</a>], [edit system services dhcp-local-server <a href="#">dhcpv6</a>], [edit forwarding-options <a href="#">dhcp-relay</a>], [edit forwarding-options dhcp-relay <a href="#">dhcpv6</a>], [edit system services dhcp-local-server <a href="#">group group-name</a>], [edit system services dhcp-local-server dhcpv6 <a href="#">group group-name</a>], [edit forwarding-options dhcp-relay <a href="#">group group-name</a>], [edit forwarding-options dhcp-relay dhcpv6 <a href="#">group group-name</a>] </pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure bidirectional failure detection timers and authentication criteria for static routes.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">DHCP Liveness Detection Overview on page 1472</a></li> <li>• <a href="#">Configuring Detection of DHCP Local Server Client Connectivity on page 1436</a></li> <li>• <a href="#">Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 1473</a></li> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> </ul>

- *Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients*

## local-server-group (DHCP Relay Agent Option)

---

<b>Syntax</b>	<code>local-server-group <i>local-server-group</i>;</code>
<b>Hierarchy Level</b>	[edit forwarding-options dhcp-relay relay-option (default-action   equals   starts-with)], [edit forwarding-options dhcp-relay group <i>group-name</i> relay-option (default-action   equals   starts-with)], [edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay</a> ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a> ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a> ...]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	<p>Forward DHCP client packets to the specified group of DHCP local servers when you use the DHCP relay selective processing feature. You can configure the forwarding operation globally or for a group of interfaces.</p> <p>The <b>local-server-group</b> option is not supported for DHCPv6 relay agent.</p>
<b>Options</b>	<b><i>local-server-group</i></b> —Name of DHCP local server group.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Using DHCP Option Information to Selectively Process DHCP Client Traffic</i></li></ul>

## mac-address (DHCP Relay Agent)

<b>Syntax</b>	mac-address;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify that the MAC address from the client PDU be concatenated with the username during the subscriber authentication or client authentication process.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> </ul>

## method

---

<b>Syntax</b>	<pre>method {   bfd {     version (0   1   automatic);     minimum-interval <i>milliseconds</i>;     minimum-receive-interval <i>milliseconds</i>;     multiplier <i>number</i>;     no-adaptation;     transmit-interval {       minimum-interval <i>milliseconds</i>;       threshold <i>milliseconds</i>;     }     detection-time {       threshold <i>milliseconds</i>;     }     session-mode (automatic   multihop   singlehop);     holddown-interval <i>milliseconds</i>;   } }</pre>
<b>Hierarchy Level</b>	<pre>[edit system services dhcp-local-server <i>liveness-detection</i>], [edit system services dhcp-local-server dhcpv6 <i>liveness-detection</i>], [edit forwarding-options dhcp-relay <i>liveness-detection</i>], [edit forwarding-options dhcp-relay dhcpv6 <i>liveness-detection</i>], [edit system services dhcp-local-server group <i>group-name</i> <i>liveness-detection</i>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> <i>liveness-detection</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> <i>liveness-detection</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <i>liveness-detection</i>]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure the liveness detection method.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li><li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li></ul>

## minimum-interval

<b>Syntax</b>	<code>minimum-interval <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the <a href="#">transmit-interval</a> <a href="#">minimal-interval</a> and <a href="#">minimum-receive-interval</a> statements.</p>
<b>Options</b>	<p><i>milliseconds</i> — Specify the minimum interval value for BFD liveliness detection.</p> <p><b>Range:</b> 1 through 255,000</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## minimum-receive-interval

---

<b>Syntax</b>	<code>minimum-receive-interval <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a>],</code> <code>[edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a>],</code> <code>[edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a>], [edit forwarding-options</code> <code>  dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a>],</code> <code>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],</code> <code>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method</code> <code>  <a href="#">bfd</a>],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],</code> <code>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method</code> <code>  <a href="#">bfd</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure the minimum interval at which the local routing device (or switch) must receive a reply from a neighbor with which it has established a BFD session.
<b>Options</b>	<i>milliseconds</i> — Specify the minimum receive interval value. <b>Range:</b> 1 through 255,000
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li><li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li></ul>



## multiplier

<b>Syntax</b>	<code>multiplier <i>number</i>;</code>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.  Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Configure the number of hello packets not received by the neighbor before Bidirectional Forwarding Detection (BFD) declares the neighbor down.
<b>Options</b>	<p><b>number</b>—Maximum allowable number of hello packets missed by the neighbor.  <b>Range:</b> 1 through 255  <b>Default:</b> 3</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.  routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## next-hop (Dynamic Access-Internal Routes)


---

<b>Syntax</b>	<code>next-hop <i>next-hop</i>;</code>
<b>Hierarchy Level</b>	<code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access route <i>prefix</i>],</code> <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access route <i>prefix</i>],</code> <code>[edit dynamic-profiles <i>profile-name</i> routing-options <b>access</b> route <i>prefix</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5. Support at the <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access route <i>prefix</i>]</code> and <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access route <i>prefix</i>]</code> hierarchy levels introduced in Junos OS Release 10.1.
<b>Description</b>	Dynamically configure the next-hop address for an access route. Access routes are typically unnumbered interfaces.
<b>Options</b>	<i>next-hop</i> —Either the specific next-hop address you want to assign to the access route or one of the following next-hop address predefined variables. <ul style="list-style-type: none"><li>For IPv4 access routes, use the variable, <b>\$junos-framed-route-nexthop</b>. The route prefix variable is dynamically replaced with the value in Framed-Route RADIUS attribute [22].</li><li>For IPv6 access routes, use the variable, <b>\$junos-framed-route-ipv6-nexthop</b>. The variable is dynamically replaced with the value in Framed-IPv6-Route RADIUS attribute [99].</li></ul>
<b>Required Privilege Level</b>	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li><i>Configuring Dynamic Access Routes for Subscriber Management</i></li></ul>


## no-adaptation

<b>Syntax</b>	no-adaptation;
<b>Hierarchy Level</b>	[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <i>Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</i></li> </ul>

## no-allow-snooped-clients

<b>Syntax</b>	no-allow-snooped-clients;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay <b>dhcpv6</b> group <i>group-name</i> interface <i>interface-name</i> <b>overrides</b>],</p> <p>[edit forwarding-options dhcp-relay <b>dhcpv6</b> group <i>group-name</i> <b>overrides</b>],</p> <p>[edit forwarding-options dhcp-relay <b>dhcpv6</b> <b>overrides</b>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <b>overrides</b>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> <b>overrides</b>],</p> <p>[edit forwarding-options dhcp-relay <b>overrides</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay</b> ...],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <b>dhcp-relay</b> ...],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options <b>dhcp-relay</b> ...]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Support at the [edit ... <b>dhcpv6</b>] hierarchy levels introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
<b>Description</b>	<p>Explicitly disable DHCP snooping support on the router or switch.</p> <p>Use the statement at the [edit ... <b>dhcpv6</b>] hierarchy levels to explicitly disable snooping support on the router or switch for DHCPv6 relay agent.</p>
<div>  <p><b>NOTE:</b> In Junos OS Release 10.0 and earlier, DHCP snooping is <i>enabled</i> by default. In Release 10.1 and later, DHCP snooping is <i>disabled</i> by default.</p> </div>	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> <li>• <a href="#">Overriding the Default DHCP Relay Configuration Settings on page 1447</a></li> <li>• <a href="#">DHCP Snooping Support on page 1367</a></li> </ul>

## no-bind-on-request (DHCP Relay Agent)

<b>Syntax</b>	no-bind-on-request;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],  [edit forwarding-options dhcp-relay <a href="#">overrides</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.4.  Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.  Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
<b>Description</b>	<p>Explicitly disable automatic binding of received DHCP request messages that have no entry in the database (<i>stray</i> requests). Use the statement at the <b>[edit ... dhcpv6]</b> hierarchy levels to configure DHCPv6 support.</p> <p>M120 and M320 routers do not support DHCPv6.</p>
<div>  <p><b>NOTE:</b> Beginning with Junos OS Release 10.4, automatic binding of stray requests is enabled by default. In Junos OS Release 10.3 and earlier releases, automatic binding of stray requests is disabled by default.</p> </div>	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.  interface-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Extended DHCP Relay Agent Overview on page 1379](#)
  - [Overriding the Default DHCP Relay Configuration Settings on page 1447](#)
  - [Disabling Automatic Binding of Stray DHCP Requests on page 1463](#)


---

## option-60 (DHCP Relay Agent)


---

<b>Syntax</b>	option-60;
<b>Hierarchy Level</b>	[edit forwarding-options dhcp-relay authentication <a href="#">username-include</a> ], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a> ], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a> ], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Specify that the payload of the Option 60 (Vendor Class Identifier) from the client PDU is concatenated with the username during the subscriber authentication or client authentication process.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a>

## option-82 (DHCP Relay Agent)

<b>Syntax</b>	<code>option-82 &lt;circuit-id&gt; &lt;remote-id&gt;;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Specify the option 82 that is concatenated with the username during the subscriber authentication or client authentication process. You can specify either, both, or neither the Agent Circuit ID and the Agent Remote ID suboptions. If you specify both, the Agent Circuit ID is supplied first, followed by a delimiter, and then the Agent Remote ID. If neither suboption is supplied, the raw payload of option 82 is concatenated to the username.
<div>  <b>NOTE:</b> The option 82 value used in creating the username is based on the option 82 value that is encoded in the outgoing (relayed) PDU. </div>	
<b>Options</b>	<p><b>circuit-id</b>—(Optional) The string for the Agent Circuit ID suboption (suboption 1).</p> <p><b>remote-id</b>—(Optional) The string for the Agent Remote ID suboption (suboption 2).</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> </ul>

## option-number (DHCP Relay Agent Option)

<b>Syntax</b>	<code>option-number option-number;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options <a href="#">dhcp-relay relay-option</a>],</p> <p>[edit forwarding-options <a href="#">dhcp-relay dhcpv6 relay-option</a>],</p> <p>[edit forwarding-options dhcp-relay <a href="#">group group-name relay-option</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 <a href="#">group group-name relay-option</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay ...</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay ...</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay ...</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.3.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
<b>Description</b>	<p>Specify the DHCP option DHCP relay agent uses for selective processing of client traffic. You can configure support globally or for a named group of interfaces. You can also configure support for the extended DHCP relay agent on a per logical system and per routing instance basis.</p> <p>Use the <a href="#">[edit forwarding-options dhcp-relay dhcpv6]</a> hierarchy level to configure the DHCPv6 relay agent support.</p>
<b>Options</b>	<i>option-number</i> —The DHCP or DHCPv6 option in the incoming traffic.
<div>  <b>NOTE:</b> EX Series switches do not support the User Class Options.         </div>	
<ul style="list-style-type: none"> <li>• 15 (DHCPv6 only)—Use DHCPv6 option 15 (User Class Option) in packets</li> <li>• 16 (DHCPv6 only)—(MX Series routers and EX Series switches only) Use DHCPv6 option 16 (Vendor Class Option) in packets</li> <li>• 60 (DHCPv4 only)—(MX Series routers and EX Series switches only) Use DHCP option 60 (Vendor Class Identifier) in DHCP packets</li> <li>• 77 (DHCPv4 only)—Use DHCP option 77 (User Class Identifier) in packets</li> </ul>	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using DHCP Option Information to Selectively Process DHCP Client Traffic</a></li> <li>• <a href="#">Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)</a> on page 1416</li> </ul>



## overrides (DHCP Relay Agent)

<b>Syntax</b>	<pre> overrides {   allow-snooped-clients;   allow-no-end-options;   always-write-giaddr;   always-write-option-82;   client-discover-match (option60-and-option82   incoming-interface);   delay-authentication;   delete-binding-on-renegotiation;   disable-relay;   interface-client-limit <i>number</i>;   layer2-unicast-replies;   no-allow-snooped-clients;   no-bind-on-request;   proxy-mode;   replace-ip-source-with;   send-release-on-delete;   trust-option-82; } </pre>
<b>Hierarchy Level</b>	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay <b>dhcpv6</b>], [edit forwarding-options dhcp-relay <b>group</b> <i>group-name</i>], [edit forwarding-options dhcp-relay <b>group</b> <i>group-name</i> <b>interface</b> <i>interface-name</i>], [edit forwarding-options dhcp-relay <b>dhcpv6</b> <b>group</b> <i>group-name</i>], [edit forwarding-options dhcp-relay <b>dhcpv6</b> <b>group</b> <i>group-name</i> <b>interface</b> <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay</b> ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   forwarding-options <b>dhcp-relay</b> ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...] </pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Support for the <b>delete-binding-on-renegotiation</b> statement introduced in Junos OS Release 13.2 for EX Series switches.</p> <p>Support for the <b>allow-no-end-options</b> statement introduced in Junos OS Release 14.1X53 for EX Series switches.</p>
<b>Description</b>	<p>Override the default configuration settings for the extended DHCP relay agent. Specifying the <b>overrides</b> statement with no subordinate statements removes all DHCP relay agent overrides at that hierarchy level. Use the statement at the <b>[edit ... dhcpv6]</b> hierarchy levels to configure DHCPv6 support.</p> <p>M120 and M320 routers do not support DHCPv6.</p> <p>The following statements are supported at both the <b>[edit ... dhcp-relay]</b> and <b>[edit ... dhcpv6]</b> hierarchy levels. All other statements are supported at the <b>dhcp-relay</b> hierarchy levels only.</p> <ul style="list-style-type: none"> <li>• <b>allow-snooped-clients</b></li> </ul>

- **interface-client-limit**
- **no-allow-snooped-clients**
- **no-bind-on-request**
- **send-release-on-delete**

The remaining statements are explained separately.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
---------------------------------	---

- |                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li><li>• <a href="#">Overriding the Default DHCP Relay Configuration Settings on page 1447</a></li><li>• <a href="#">Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429</a></li><li>• <a href="#">dhcp-relay on page 1570</a></li></ul> |
|------------------------------|--|

## password (DHCP Relay Agent)

<b>Syntax</b>	<code>password password-string;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay <a href="#">authentication</a>],  [edit forwarding-options dhcp-relay dhcpv6 <a href="#">authentication</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">authentication</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">authentication</a>],  [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">authentication</a>],  [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">authentication</a>],  [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">authentication</a>],  [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">authentication</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">authentication</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">authentication</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">authentication</a>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">authentication</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">authentication</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">authentication</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">authentication</a>],  [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">authentication</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels introduced in Junos OS Release 11.4.</p>
<b>Description</b>	<p>Configure the password that is sent to the external AAA authentication server for subscriber authentication or client authentication. Use the statement at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels to configure DHCPv6 support.</p>
<b>Options</b>	<i>password-string</i> —Authentication password.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> <li>• <a href="#">Configuring Passwords for Usernames on page 1439</a></li> </ul>

## preference (Subscriber Management)

---

<b>Syntax</b>	<code>preference route-distance</code>
<b>Hierarchy Level</b>	<code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access route <i>prefix</i>]</code> , <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access route <i>prefix</i>]</code> , <code>[edit dynamic-profiles <i>profile-name</i> routing-options <b>access</b> route <i>prefix</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5. Support at <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options access route <i>prefix</i>]</code> and <code>[edit dynamic-profiles <i>profile-name</i> routing-instances \$junos-routing-instance routing-options rib <i>routing-table-name</i> access route <i>prefix</i>]</code> hierarchy levels introduced in Junos OS Release 10.1.
<b>Description</b>	Dynamically configure the distance for an access route.
<b>Options</b>	<b><i>route-distance</i></b> —Either the specific distance you want to assign to the access route or either of the following distance variables: <ul style="list-style-type: none"><li>• <b><i>\$junos-framed-route-distance</i></b>—Distance of an IPv4 access route; the variable is dynamically replaced with the preference value (Subattribute 5) from the RADIUS Framed-Route attribute [22].</li><li>• <b><i>\$junos-framed-route-ipv6-distance</i></b>—Distance of an IPv6 access route; the variable is dynamically replaced with the preference value (Subattribute 5) from the RADIUS Framed-IPv6-Route attribute [99].</li></ul>
<b>Required Privilege Level</b>	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Dynamic Access Routes for Subscriber Management</i></li></ul>

## prefix (DHCP Relay Agent)

<b>Syntax</b>	<code>prefix <i>prefix</i>;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay <b>dhcpv6</b> (<b>relay-agent-interface-id</b>   relay-agent-remote-id)],</p> <p>[edit forwarding-options dhcp-relay <b>dhcpv6</b> group <i>group-name</i> (<b>relay-agent-interface-id</b>   relay-agent-remote-id)],</p> <p>[edit forwarding-options dhcp-relay relay-option-82 (<b>circuit-id</b>   remote-id)],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> relay-option-82 (<b>circuit-id</b>   remote-id)],</p> <p>[edit logical-systems <i>logical-system-name</i> ... forwarding-options dhcp-relay <b>dhcpv6</b> (<b>relay-agent-interface-id</b>   relay-agent-remote-id)],</p> <p>[edit logical-systems <i>logical-system-name</i> ... forwarding-options dhcp-relay ... relay-option-82 (<b>circuit-id</b>   remote-id)],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <b>dhcpv6</b> (<b>relay-agent-interface-id</b>   relay-agent-remote-id)],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ... relay-option-82 (<b>circuit-id</b>   remote-id)]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the [edit ... <b>dhcpv6</b>] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p> <p>Support at the [edit ... <b>relay-agent-remote-id</b>] and [edit ... <b>remote-id</b>] hierarchy levels introduced in Junos OS Release 14.1.</p>
<b>Description</b>	<p>Add a prefix to the DHCP base option 82 Agent Circuit ID (suboption 1) or Agent Remote ID (suboption 2) information, or to the DHCPv6 option 18 (Relay Agent Interface-ID) or option 37 (Relay Agent Remote-ID) information in DHCP packets that DHCP relay agent sends to a DHCP server. The prefix can consist of any combination of the hostname, logical system name, and routing instance name.</p>
<b>Options</b>	<p><i>prefix</i>—Any of the following:</p> <ul style="list-style-type: none"> <li>• <b>host-name</b>—Prepend the hostname of the router configured with the <b>host-name</b> statement at the [edit <b>system</b>] hierarchy level to the DHCP option information.</li> <li>• <b>logical-system-name</b>—Prepend the name of the logical system to the option information.</li> <li>• <b>routing-instance-name</b>—Prepend the name of the routing instance to the option information.</li> </ul>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Including a Prefix in DHCP Options on page 1466</a></li> <li>• <a href="#">Using DHCP Relay Agent Option 82 Information on page 1464</a></li> <li>• <a href="#">Configuring DHCPv6 Relay Agent Options</a></li> </ul>

## proxy-mode

---

<b>Syntax</b>	proxy-mode;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay <a href="#">overrides</a>], [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Enable DHCP relay proxy mode on the extended DHCP relay. Proxy mode supports all extended DHCP relay functionality.</p> <p>You cannot configure both the DHCP relay proxy and the extended DHCP local server on the same interface.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">DHCP Relay Proxy Overview on page 1382</a></li><li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li><li>• <a href="#">Enabling DHCP Relay Proxy Mode on page 1470</a></li></ul>

## relay-agent-interface-id (DHCPv6 Relay Agent)

<b>Syntax</b>	<pre> relay-agent-interface-id {   prefix <i>prefix</i>;   use-interface-description (logical   device);   use-option-82; } </pre>
<b>Hierarchy Level</b>	<pre> [edit forwarding-options dhcp-relay <a href="#">dhcpv6</a>], [edit forwarding-options dhcp-relay <a href="#">dhcpv6 group</a> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay dhcpv6</a> ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>  forwarding-options dhcp-relay <a href="#">dhcpv6</a> ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">dhcpv6</a> ...] </pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
<b>Description</b>	<p>Insert the DHCPv6 Relay Agent Interface-ID option (option 18) in DHCPv6 packets destined for the DHCPv6 server.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">dhcp-relay on page 1570</a></li> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> <li>• <a href="#">DHCPv6 Relay Agent Overview on page 1384</a></li> <li>• <a href="#">Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 1470</a></li> </ul>

## relay-agent-remote-id (DHCPv6 Relay Agent Username)

---

<b>Syntax</b>	relay-agent-remote-id;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>For MX Series routers only, <b>enterprise-id</b> and <b>remote-id</b> options introduced in Junos OS Release 12.3R3.</p> <p>For MX Series routers only, the <b>enterprise-id</b> and <b>remote-id</b> options are obsoleted starting in Junos OS Releases 12.3R7, 13.2R4, 13.3R2, and 14.1R1.</p>
<b>Description</b>	<p>Specify that the DHCPv6 Relay Agent Remote-ID option (option 37) in the client PDU name is concatenated with the username during the subscriber authentication or client authentication process. In order to generate an ASCII version of the username, the router concatenates only the remote-id portion of option 37 to the username, and ignores the enterprise number.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">DHCPv6 Relay Agent Overview on page 1384</a></li><li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li></ul>



## relay-option (DHCP Relay Agent)

<b>Syntax</b>	<pre> relay-option {   option-number option-number;   default-action {     drop;     forward-only;     local-server-group local-server-group;     relay-server-group relay-server-group;   }   equals (ascii <i>ascii-string</i>   hexadecimal <i>hexadecimal-string</i>) {     drop;     forward-only;     local-server-group local-server-group;     relay-server-group relay-server-group;   }   starts-with (ascii <i>ascii-string</i>   hexadecimal <i>hexadecimal-string</i>) {     drop;     forward-only;     local-server-group local-server-group;     relay-server-group relay-server-group;   } } </pre>
<b>Hierarchy Level</b>	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay group group-name], [edit forwarding-options dhcp-relay dhcpv6 group group-name], [edit logical-systems logical-system-name forwarding-options dhcp-relay ...], [edit logical-systems logical-system-name routing-instances routing-instance-name   forwarding-options dhcp-relay ...], [edit routing-instances routing-instance-name forwarding-options dhcp-relay ...] </pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.3.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
<b>Description</b>	<p>Configure the extended DHCP relay agent selective processing that is based on DHCP options in DHCP client packets and specify the action to perform on client traffic. You can configure support globally or for a named group of interfaces, and for either DHCP or DHCPv6 relay agent.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Using DHCP Option Information to Selectively Process DHCP Client Traffic</i></li> </ul>

## relay-option-82

---

Syntax	<pre>relay-option-82 {   circuit-id {     prefix <i>prefix</i>;     use-interface-description (logical   device);   }   remote-id {     prefix <i>prefix</i>;     use-interface-description (logical   device);   } }</pre>
Hierarchy Level	<pre>[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   forwarding-options dhcp-relay <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <i>group</i>   <i>group-name</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
Description	<p>Enable or disable the insertion of the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server.</p> <p>To enable insertion of option 82 information in DHCP packets, you must specify at least one of the <b>circuit-id</b> or <b>remote-id</b> statements.</p> <p>You can use the <b>relay-option-82</b> statement and its subordinate statements at the <b>[edit forwarding-options dhcp-relay]</b> hierarchy level to control insertion of option 82 information globally, or at the <b>[edit forwarding-options dhcp-relay <i>group group-name</i>]</b> hierarchy level to control insertion of option 82 information for a named group of interfaces.</p> <p>To restore the default behavior (option 82 information is not inserted into DHCP packets), use the <b>delete relay-option-82</b> statement.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Using DHCP Relay Agent Option 82 Information on page 1464</a></li><li>• <a href="#">dhcp-relay on page 1570</a></li></ul>

## relay-server-group (DHCP Relay Agent Option)

<b>Syntax</b>	<code>relay-server-group <i>relay-server-group</i>;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay relay-option (default-action   equals   starts-with),  [edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action   equals   starts-with),  [edit forwarding-options dhcp-relay group <i>group-name</i> relay-option (default-action   equals   starts-with),  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> relay-option (default-action   equals   starts-with),  [edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay</a> ...],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a> ...],  [edit routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a> ...]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.3.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Relay DHCP client packets to the specified group of DHCP servers when you use the DHCP relay selective processing feature. You can configure the relay operation globally or for a group of interfaces, and for either DHCP or DHCPv6 relay agent.
<b>Options</b>	<i>relay-server-group</i> —Name of DHCP server group.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Using DHCP Option Information to Selectively Process DHCP Client Traffic</i></li> </ul>

## replace-ip-source-with

---

<b>Syntax</b>	replace-ip-source-with giaddr;
<b>Hierarchy Level</b>	[edit forwarding-options dhcp-relay <a href="#">overrides</a> ], [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a> ], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">overrides</a> ], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li><li>• <a href="#">Replacing the DHCP Relay Request and Release Packet Source Address on page 1450</a></li></ul>

## routing-instance-name (DHCP Relay Agent)

<b>Syntax</b>	routing-instance-name;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels introduced in Junos OS Release 11.4.</p>
<b>Description</b>	Specify that the routing instance name is concatenated with the username during the subscriber authentication or client authentication process. No routing instance name is concatenated if the configuration is in the default routing instance. Use the statement at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels to configure DHCPv6 support.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> <li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li> </ul>

## send-release-on-delete (DHCP Relay Agent)

<b>Syntax</b>	send-release-on-delete;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],</p> <p>[edit forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
<b>Description</b>	<p>Send a release message to the DHCP (or DHCPv6) server whenever DHCP relay or relay proxy deletes a client. Use the statement at the <b>[edit ... dhcpv6]</b> hierarchy levels to configure DHCPv6 support.</p> <p>M120 and M320 routers do not support DHCPv6.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> <li>• <a href="#">Overriding the Default DHCP Relay Configuration Settings on page 1447</a></li> <li>• <a href="#">Sending Release Messages When Clients Are Deleted on page 1462</a></li> </ul>

## server-group

<b>Syntax</b>	<pre>server-group {     server-group-name {         server-ip-address;     } }</pre>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay],          [edit forwarding-options <a href="#">dhcp-relay dhcpv6</a>],          [edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay</a>],          [edit logical-systems <i>logical-system-name</i> forwarding-options <a href="#">dhcp-relay dhcpv6</a>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay dhcpv6</a>],          [edit routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay</a>],          [edit routing-instances <i>routing-instance-name</i> forwarding-options <a href="#">dhcp-relay dhcpv6</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.          Support at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels introduced in Junos OS Release 11.4.          Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Specify the name of a group of DHCP server addresses for use by the extended DHCP relay agent. Use the statement at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels to configure DHCPv6 support.</p>
<b>Options</b>	<p><b><i>server-group-name</i></b>—Name of the group of DHCP or DHCPv6 server addresses.</p> <p><b><i>server-ip-address</i></b>—IP address of the DHCP server belonging to this named server group. Use IPv6 addresses when configuring DHCPv6 support. You can configure a maximum of five IP addresses in each named server group.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.          interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">dhcp-relay on page 1570</a></li> <li>• <a href="#">Extended DHCP Relay Agent Overview on page 1379</a></li> <li>• <a href="#">Configuring Server Groups on page 1469</a></li> </ul>

## service-profile (DHCP Relay Agent)

---

Syntax	<code>service-profile <i>dynamic-profile-name</i>;</code>
Hierarchy Level	<code>[edit forwarding-options dhcp-relay],</code> <code>[edit forwarding-options dhcp-relay <b>dhcpv6</b>],</code> <code>[edit forwarding-options dhcp-relay <b>group</b> <i>group-name</i>],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> <b>interface</b> <i>interface-name</i>],</code> <code>[edit forwarding-options dhcp-relay dhcpv6 <b>group</b><i>group-name</i>],</code> <code>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <b>interface</b> <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> forwarding-options <b>dhcp-relay</b> ...],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i></code> <code>forwarding-options <b>dhcp-relay</b> ...],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</code>
Release Information	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Support at the <code>[edit ... <b>dhcpv6</b> ...]</code> hierarchy levels introduced in Junos OS Release 11.4.
Description	<p>Specify the default subscriber service (or the default DHCP client management service), which is activated when the subscriber (or client) logs in and no other service is activated by a RADIUS server or a provisioning server.</p> <ul style="list-style-type: none"><li>• To specify the default service for all DHCP relay agent clients, include the <b>service-profile</b> statement at the <code>[edit forwarding-options dhcp relay]</code> hierarchy level.</li><li>• To specify the default service for a named group of interfaces, include the <b>service-profile</b> statement at the <code>[edit forwarding-options dhcp relay group <i>group-name</i>]</code> hierarchy level.</li><li>• To specify the default service for a particular interface within a named group of interfaces, include the <b>service-profile</b> statement at the <code>[edit forwarding-options dhcp relay group <i>group-name</i> interface <i>interface-name</i>]</code> hierarchy level.</li></ul>
Options	<b><i>dynamic-profile-name</i></b> —Name of the dynamic profile.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">dhcp-relay on page 1570</a></li><li>• <a href="#">Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</a></li><li>• <a href="#">Grouping Interfaces with Common DHCP Configurations on page 1421</a></li><li>• <a href="#">Default Subscriber Service Overview</a></li><li>• <a href="#">Configuring a Default Subscriber Service</a></li></ul>




## session-mode


<b>Syntax</b>	<code>session-mode (automatic   multihop   singlehop);</code>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],  [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <a href="#">bfd</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Configure the session mode.
<b>Options</b>	<p><b>automatic</b>—Configure single-hop BFD sessions if the peer is directly connected to the router interface and multihop BFD sessions if the peer is not directly connected to the router interface.</p> <p><b>multihop</b>—Configure multihop BFD sessions.</p> <p><b>single-hop</b>—Configure single hop BFD sessions.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## threshold (detection-time)

---

<b>Syntax</b>	<code>threshold <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit system services dhcp-local-server liveness-detection method bfd <i>detection-time</i>],</code> <code>[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd</code> <code>    <i>detection-time</i>],</code> <code>[edit forwarding-options dhcp-relay liveness-detection method bfd <i>detection-time</i>],</code> <code>[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd <i>detection-time</i>],</code> <code>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd</code> <code>    <i>detection-time</i>],</code> <code>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method</code> <code>    bfd <i>detection-time</i>],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd</code> <code>    <i>detection-time</i>],</code> <code>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method</code> <code>    bfd <i>detection-time</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Specify the threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.
<div> <b>NOTE:</b> The threshold time must be greater than or equal to the minimum-interval or the minimum-receive-interval.</div>	
<b>Options</b>	<i>milliseconds</i> — Value for the detection time adaptation threshold. <b>Range:</b> 1 through 255,000
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li><li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li></ul>

## threshold (transmit-interval)

<b>Syntax</b>	<code>threshold <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit forwarding-options dhcp-relay liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd <a href="#">transmit-interval</a>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Specify the threshold for detecting the adaptation of the transmit interval. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.
<b>Options</b>	<p><i>milliseconds</i> — Threshold value.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p>
<div>  <p><b>NOTE:</b> The threshold value specified in the <code>threshold</code> statement must be greater than the value specified in the <code>minimum-interval</code> statement for the <code>transmit-interval</code> statement.</p> </div>	
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## trace (DHCP Relay Agent)

---

<b>Syntax</b>	trace;
<b>Hierarchy Level</b>	[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <b>interface</b> <i>interface-name</i> ], [edit forwarding-options dhcp-relay group <i>group-name</i> <b>interface</b> <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <b>interface</b> <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <b>interface</b> <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <b>interface</b> <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <b>interface</b> <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <b>interface</b> <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <b>interface</b> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4. Support at the [edit ... <b>dhcpv6</b> ] hierarchy levels introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	Enable trace operations for a group of interfaces or for a specific interface within a group. Use the statement at the [edit ... <b>dhcpv6</b> ] hierarchy levels to configure DHCPv6 support.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Tracing Extended DHCP Operations</i></li><li>• <i>Tracing Extended DHCP Operations for Specific Interfaces</i></li></ul>

## transmit-interval


<b>Syntax</b>	<pre>transmit-interval {     threshold milliseconds;     minimum-interval milliseconds; }</pre>
<b>Hierarchy Level</b>	<p>[edit system services dhcp-local-server liveness-detection method <a href="#">bfd</a>],          [edit system services dhcp-local-server dhcpv6 liveness-detection method <a href="#">bfd</a>],          [edit forwarding-options dhcp-relay liveness-detection method <a href="#">bfd</a>], [edit forwarding-options          dhcp-relay dhcpv6 liveness-detection method <a href="#">bfd</a>],          [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],          [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method  <a href="#">bfd</a>],          [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <a href="#">bfd</a>],          [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method  <a href="#">bfd</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.1.          Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Configure the Bidirectional Forwarding Detection (BFD) transmit interval.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.          routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li> <li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li> </ul>

## trust-option-82

---

<b>Syntax</b>	trust-option-82;
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay <a href="#">overrides</a>], [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <a href="#">overrides</a>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> <a href="#">overrides</a>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> <a href="#">overrides</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	Enable processing of DHCP client packets that have a gateway IP address (giaddr) of 0 (zero) and contain option 82 information. By default, the DHCP relay agent treats such packets as if they originated at an untrusted source, and drops them without further processing.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Trusting Option 82 Information on page 1451</a></li><li>• <a href="#">Overriding the Default DHCP Relay Configuration Settings on page 1447</a></li></ul>

## use-interface-description

<b>Syntax</b>	<code>use-interface-description (logical   device);</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay <b>dhcpv6</b> (<b>relay-agent-interface-id</b>   relay-agent-remote-id)],</p> <p>[edit forwarding-options dhcp-relay <b>dhcpv6</b> group <i>group-name</i> (<b>relay-agent-interface-id</b>   relay-agent-remote-id)],</p> <p>[edit forwarding-options dhcp-relay relay-option-82 (<b>circuit-id</b>   remote-id)],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> relay-option-82 (<b>circuit-id</b>   remote-id)],</p> <p>[edit logical-systems <i>logical-system-name</i> ... forwarding-options dhcp-relay dhcpv6 (<b>relay-agent-interface-id</b>   relay-agent-remote-id)],</p> <p>[edit logical-systems <i>logical-system-name</i> ... forwarding-options dhcp-relay ... relay-option-82 (<b>circuit-id</b>   remote-id)],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 (<b>relay-agent-interface-id</b>   relay-agent-remote-id)],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ... relay-option-82 (<b>circuit-id</b>   remote-id)]</p> <p>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security dhcpv6-options <b>option-18</b>],</p> <p>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security dhcpv6-options <b>option-37</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Support at the [edit ... <b>dhcpv6</b>] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p> <p>Support at the [edit ... <b>relay-agent-remote-id</b>] and [edit ... <b>remote-id</b>] hierarchy levels introduced in Junos OS Release 14.1.</p> <p>Support at the [edit vlans <i>vlan-name</i> dhcp-security dhcpv6-options <b>option-18</b>] and [edit vlans <i>vlan-name</i> dhcp-security dhcpv6-options <b>option-37</b>] hierarchy levels introduced in Junos OS Release 14.1X53-D10 for EX Series switches.</p>
<b>Description</b>	<p>Use the textual interface description instead of the interface identifier in the DHCP base option 82 Agent Circuit ID (suboption 1) or Agent Remote ID (suboption 2) information, or in the DHCPv6 option 18 (Relay Agent Interface ID) or option 37 (Relay Agent Remote ID) information in DHCP packets that the DHCP relay agent sends to a DHCP server.</p> <p>If you specify that the textual description be used and no description is configured for the interface, DHCP relay defaults to using the interface identifier. The textual description is configured using the <b>description</b> statement at the [edit interfaces <i>interface-name</i>] hierarchy level.</p>
<div>  <p><b>NOTE:</b> By default, DHCP relay accepts a maximum of 253 ASCII characters. If the textual interface description exceeds 253 characters, DHCP relay drops the packet, which results in the DHCP client failing to bind.</p> </div>	
<b>Options</b>	<p><b>logical</b>—Use the textual description that is configured for the logical interface.</p> <p><b>device</b>—Use the textual description that is configured for the device interface.</p>

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Including a Textual Description in DHCP Options on page 1468</a></li><li>• <a href="#">Using DHCP Relay Agent Option 82 Information on page 1464</a></li><li>• <i>Configuring DHCPv6 Relay Agent Options</i></li></ul>



## use-primary (DHCP Relay Agent)

<b>Syntax</b>	<code>use-primary <i>primary-profile-name</i>;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay dhcpv6 <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit forwarding-options dhcp-relay <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 <b>group</b> <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit forwarding-options dhcp-relay <b>group</b> <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 <b>group</b> <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay <b>group</b> <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <b>group</b> <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <b>group</b> <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <b>group</b> <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay <b>group</b> <i>group-name</i> <b>dynamic-profile</b> <i>profile-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.3.</p> <p>Support at the <b>[edit ... dhcpv6]</b> hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Specify the dynamic profile to configure as the primary dynamic profile. The primary dynamic profile is instantiated when the first subscriber logs in. Subsequent subscribers are not assigned the primary dynamic profile; instead, they are assigned the dynamic profile specified for the interface. When the first subscriber logs out, the next subscriber that logs in is assigned the primary dynamic profile.</p> <p>Use the statement at the <b>[edit ... dhcpv6]</b> hierarchy levels to configure DHCPv6 support.</p> <p>EX Series switches do not support DHCPv6.</p>
<b>Options</b>	<b><i>primary-profile-name</i></b> —Name of the dynamic profile to configure as the primary dynamic profile
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces](#)

## user-prefix (DHCP Relay Agent)

<b>Syntax</b>	<code>user-prefix user-prefix-string;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options dhcp-relay authentication <a href="#">username-include</a>],          [edit forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],          [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],          [edit forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],          [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],          [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],          [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],          [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>],          [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication <a href="#">username-include</a>],          [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication <a href="#">username-include</a>],          [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication <a href="#">username-include</a>],          [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication <a href="#">username-include</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels introduced in Junos OS Release 11.4.</p>
<b>Description</b>	Specify the user prefix that is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels to configure DHCPv6 support.
<b>Options</b>	<i>user-prefix-string</i> —User prefix string.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li> </ul>

## username-include (DHCP Relay Agent)

<b>Syntax</b>	<pre>username-include {   circuit-type;   client-id;   delimiter <i>delimiter-character</i>;   domain-name <i>domain-name-string</i>;   interface-name;   logical-system-name;   mac-address;   option-60;   option-82 &lt;circuit-id&gt; &lt;remote-id&gt;;   relay-agent-interface-id;   relay-agent-remote-id;   relay-agent-subscriber-id;   routing-instance-name;   user-prefix <i>user-prefix-string</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit forwarding-options dhcp-relay <a href="#">authentication</a>], [edit forwarding-options dhcp-relay dhcpv6 <a href="#">authentication</a>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <a href="#">authentication</a>], [edit forwarding-options dhcp-relay group <i>group-name</i> <a href="#">authentication</a>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the <a href="#">[edit ... dhcpv6]</a> hierarchy levels introduced in Junos OS Release 11.4.</p>
<b>Description</b>	<p>Configure the username that the router (or switch) passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router (or switch) accesses the local authentication service only and does not use external authentication services, such as RADIUS. Use the statement at the <a href="#">[edit...dhcpv6]</a> hierarchy levels to configure DHCPv6 support.</p> <p>The following statements are not supported in the DHCPv6 hierarchy levels:</p> <ul style="list-style-type: none"> <li>• <b>mac-address</b></li> <li>• <b>option-60</b></li> <li>• <b>option-82</b></li> </ul> <p>The following statements are supported in the DHCPv6 hierarchy levels only:</p> <ul style="list-style-type: none"> <li>• <b>relay-agent-interface-id</b></li> <li>• <b>relay-agent-remote-id</b></li> <li>• <b>relay-agent-subscriber-id</b></li> </ul>

The remaining statements are explained separately.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Creating Unique Usernames for DHCP Clients on page 1439</a></li><li>• <a href="#">Using External AAA Authentication Services with DHCP on page 1420</a></li></ul>

---

## version (BFD)

---

<b>Syntax</b>	version (0   1   automatic);
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec address bfd-liveness-detection], [edit system services dhcp-local-server liveness-detection method <b>bfd</b> ], [edit system services dhcp-local-server dhcpv6 liveness-detection method <b>bfd</b> ], [edit forwarding-options dhcp-relay liveness-detection method <b>bfd</b> ], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method <b>bfd</b> ], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method <b>bfd</b> ], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method <b>bfd</b> ], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method <b>bfd</b> ], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method <b>bfd</b> ], [edit protocols ldp oam bfd-liveness-detection], [edit protocols ldp oam fec address bfd-liveness-detection]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Configure the BFD protocol version to detect.
<b>Options</b>	<b>0</b> —Use BFD protocol version 0.  <b>1</b> —Use BFD protocol version 1.  <b>automatic</b> —Autodetect the BFD protocol version. <b>Default:</b> automatic
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 1395</a></li><li>• <a href="#">Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</a></li><li>• <a href="#">Configuring BFD for LDP LSPs</a></li></ul>

---

## Other Configuration Statements

---

- [cache-size](#) on page 1639
- [cache-timeout-negative](#) on page 1640
- [certificates](#) on page 1641
- [certification-authority](#) on page 1642
- [connection-limit](#) on page 1643
- [crl \(Encryption Interface\)](#) on page 1644
- [domain-search](#) on page 1644
- [encoding](#) on page 1645
- [enrollment-retry](#) on page 1645
- [enrollment-url](#) on page 1646
- [family \(for EX Series switches\)](#) on page 1647
- [file](#) on page 1650
- [ftp](#) on page 1651
- [http](#) on page 1652
- [https](#) on page 1653
- [ldap-url](#) on page 1654
- [lease-time](#) on page 1655
- [load-key-file](#) on page 1656
- [local](#) on page 1657
- [local-certificate](#) on page 1658
- [maximum-certificates](#) on page 1658
- [maximum-hop-count](#) on page 1659
- [maximum-lease-time \(DHCP\)](#) on page 1659
- [minimum-wait-time](#) on page 1660
- [name-server](#) on page 1660
- [no-listen](#) on page 1661
- [outbound-ssh](#) on page 1662
- [path-length](#) on page 1665
- [port \(HTTP/HTTPS\)](#) on page 1665
- [port \(SRC Server\)](#) on page 1666
- [process-inform](#) on page 1666
- [protocol-version](#) on page 1667
- [rate-limit](#) on page 1668
- [reconfigure](#) on page 1669
- [retransmission-attempt](#) on page 1670

- [retransmission-interval](#) on page 1671
- [server \(DNS and TFTP Service\)](#) on page 1671
- [server-address](#) on page 1672
- [server-identifier](#) on page 1673
- [servers](#) on page 1674
- [service-deployment](#) on page 1674
- [services \(System Services\)](#) on page 1675
- [session \(Time-out\)](#) on page 1677
- [sip-server](#) on page 1678
- [source-address \(SRC Software\)](#) on page 1678
- [source-address-giaddr](#) on page 1679
- [ssh](#) on page 1680
- [static-binding](#) on page 1681
- [system-generated-certificate](#) on page 1682
- [telnet](#) on page 1682
- [tftp](#) on page 1683
- [traceoptions \(DNS and TFTP Packet Forwarding\)](#) on page 1684
- [traceoptions](#) on page 1686
- [update-server](#) on page 1688
- [web-management](#) on page 1689
- [wins-server \(System\)](#) on page 1690

## cache-size

---

<b>Syntax</b>	cache-size <i>bytes</i> ;
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the cache size for digital certificates.
<b>Options</b>	<b>bytes</b> —Cache size for digital certificates. <b>Range:</b> 64 through 4,294,967,295 <b>Default:</b> 2 megabytes (MB)



**NOTE:** We recommend that you limit your cache size to 4 MB.

---

<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## cache-timeout-negative

---

<b>Syntax</b>	cache-timeout-negative <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Configure a negative cache for digital certificates.
<b>Options</b>	<b>seconds</b> —Negative time to cache digital certificates, in seconds. <b>Range:</b> 10 through 4,294,967,295 <b>Default:</b> 20



**CAUTION:** Configuring a large negative cache value can lead to a denial-of-service attack.

---

<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>



## certificates

<b>Syntax</b>	<pre> certificates {   cache-size bytes;   cache-timeout-negative seconds;   certification-authority ca-profile-name {     ca-name ca-identity;     crt file-name;     encoding (binary   pem);     enrollment-url url-name;     file certificate-filename;     ldap-url url-name;   }   enrollment-retry attempts;   local certificate-name {     certificate-key-string;     load-key-file URL filename;   }   maximum-certificates number;   path-length certificate-path-length; } </pre>
<b>Hierarchy Level</b>	[edit security]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>(Encryption interface on M Series and T Series routers and EX Series switches only)</p> <p>Configure the digital certificates for IPsec.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## certification-authority

---

<b>Syntax</b>	<code>certification-authority <i>ca-profile-name</i> {     <i>ca-name ca-identity</i>;     <i>crl file-name</i>;     <i>encoding</i> (binary   pem);     <i>enrollment-url url-name</i>;     <i>file certificate-filename</i>;     <i>ldap-url url-name</i>; }</code>
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Configure a certificate authority profile name.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li><i>Configuring Digital Certificates for an ES PIC</i></li></ul>

## connection-limit

<b>Syntax</b>	<code>connection-limit <i>limit</i>;</code>
<b>Hierarchy Level</b>	<code>[edit system services finger],</code> <code>[edit system services ftp],</code> <code>[edit system services netconf ssh],</code> <code>[edit system services ssh],</code> <code>[edit system services telnet],</code> <code>[edit system services xnm-clear-text],</code> <code>[edit system services xnm-ssl]</code>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure the maximum number of connections sessions for each type of system services (finger, ftp, ssh, telnet, xnm-clear-text, or xnm-ssl) per protocol (either IPv6 or IPv4).
<b>Options</b>	<p><b>limit</b>—(Optional) Maximum number of established connections per protocol (either IPv6 or IPv4).</p> <p><b>Range:</b> 1 through 250</p> <p><b>Default:</b> 75</p>



**NOTE:** The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</i></li> <li>• <i>Configuring DTCP-over-SSH Service for the Flow-Tap Application</i></li> <li>• <i>Configuring Finger Service for Remote Access to the Router</i></li> <li>• <i>Configuring FTP Service for Remote Access to the Router or Switch</i></li> <li>• <i>Configuring SSH Service for Remote Access to the Router or Switch</i></li> <li>• <i>Configuring Telnet Service for Remote Access to a Router or Switch</i></li> </ul>

## crl (Encryption Interface)

---

<b>Syntax</b>	<code>crl <i>file-name</i>;</code>
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the certificate revocation list (CRL). A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis.
<b>Options</b>	<i>file-name</i> —Specify the file from which to read the CRL.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>

## domain-search

---

<b>Syntax</b>	<code>domain-search [ <i>domain-list</i> ];</code>
<b>Hierarchy Level</b>	[edit system], [edit system services dhcp], [edit system services dhcp], [edit system services dhcp pool], [edit system services dhcp static-binding]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure a list of domains to be searched.
<b>Options</b>	<i>domain-list</i> —A list of domain names to search. The list can contain up to six domain names, with a total of up to 256 characters.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Reaching a Domain Name System Server</i></li><li>• <i>Configuring the Router or Interface to Act as a DHCP Server on J Series Services Routers</i></li><li>• <i>Configuring a DHCP Server on Switches (CLI Procedure)</i></li></ul>

## encoding

<b>Syntax</b>	encoding (binary   pem);
<b>Hierarchy Level</b>	[edit security ike policy <i>ike-peer-address</i> ], [edit security certificates <b>certification-authority</b> <i>ca-profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Specify the file format used for the <b>local-certificate</b> and <b>local-key-pair</b> statements.
<b>Options</b>	<b>binary</b> —Binary file format.  <b>pem</b> —Privacy-enhanced mail (PEM), an ASCII base 64 encoded format. <b>Default:</b> binary
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> <li>• <i>Configuring an IKE Policy for Digital Certificates for an ES PIC</i></li> </ul>

## enrollment-retry

<b>Syntax</b>	enrollment-retry <i>attempts</i> ;
<b>Hierarchy Level</b>	[edit security <b>certificates</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Specify how many times a router or switch can resend a digital certificate request.
<b>Options</b>	<b>attempts</b> —Number of enrollment retries. <b>Range:</b> 0 through 100 <b>Default:</b> 0
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## enrollment-url

---

<b>Syntax</b>	<code>enrollment-url <i>url-name</i>;</code>
<b>Hierarchy Level</b>	[edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Specify where your router or switch sends Simple Certificate Enrollment Protocol-based (SCEP-based) certificate enrollment requests (certificate authority URL).
<b>Options</b>	<i>url-name</i> —Certificate authority URL.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>

## family (for EX Series switches)

<b>Syntax</b>	<a href="#">family ccc on page 1647</a> <a href="#">family ethernet-switching on page 1647</a> <a href="#">family inet on page 1647</a> <a href="#">family inet6 on page 1648</a> <a href="#">family iso on page 1648</a> <a href="#">family mpls on page 1648</a>
<b>family ccc</b>	family ccc;
<b>family ethernet-switching</b>	<pre> family ethernet-switching {   filter [input   output] <i>filter-name</i>;   native-vlan-id <i>vlan-id</i>;   port-mode <i>mode</i>;   vlan (802.1Q Tagging) {     members [ (all   <i>names</i>   <i>vlan-ids</i>)];   } }</pre>
<b>family inet</b>	<pre> family inet {   address <i>address</i> {     arp <i>ip-address</i> (mac   multicast-mac) <i>mac-address</i> &lt;publish&gt;;     broadcast;     preferred;     primary;     vrrp-group <i>group-id</i> {       advertise-interval <i>milliseconds</i>;       preempt   no-preempt {         hold-time <i>seconds</i>;       }       priority <i>number</i>;       virtual-address [<i>addresses</i>];       virtual-link-local-address <i>ip-address</i>;     }   }   dhcp {     client-identifier (ascii <i>ascii</i>   hexadecimal <i>hexadecimal</i>);     lease-time (<i>seconds</i>   infinite);     retransmission-attempt <i>number</i>;     retransmission-interval <i>seconds</i>;     server-address <i>ip-address</i>;     update-server;     vendor-id <i>vendor-id</i>;   }   filter {     input <i>filter-name</i>;     output <i>filter-name</i>;   }   mtu <i>bytes</i>;   no-redirects;   no-neighbor-learn;   primary;   rpf-check; }</pre>

	<pre>        targeted-broadcast;     }  family inet6 {     address address {         eui-64;         nd6-stale-time seconds;         ndp ip-address (mac   multicast-mac) mac-address &lt;publish&gt;;         preferred;         primary;         vrrp-inet6-group group-id {             inet6-advertise-interval milliseconds;             preempt   preempt {                 hold-time seconds;             }             priority number;             virtual-inet6-address [addresses];             virtual-link-local-address ipv6-address;         }     }     (dad-disable   no-dad-disable);     filter {         input filter-name;         output filter-name;     }     mtu bytes;     no-neighbor-learn     rpf-check; }  family iso {     address interface-address;     mtu bytes; }  family mpls {     mtu bytes; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ], [edit interfaces interface-range <i>name</i> unit <i>logical-unit-number</i> ]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches, including options <b>ethernet-switching</b> , <b>inet</b> , and <b>iso</b> . Option <b>inet6</b> introduced in Junos OS Release 9.3 for EX Series switches. Options <b>ccc</b> and <b>mpls</b> introduced in Junos OS Release 9.5 for EX Series switches.
Description	Configure protocol family information for the logical interface on the switch.  You must configure a logical interface to be able to use the physical device.



**Default** Interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to **family ethernet-switching** by the default factory configuration. Before you can change the family setting for an interface to another family type, you must delete this default setting or any user-configured family setting. EX6200 and EX8200 switch interfaces do not have a default family setting.

**Options** See [Table 145 on page 1649](#) for protocol families available on the switch interfaces. Different protocol families support different subsets of the interface types on the switch. Interface types on the switch are:

- Aggregated Ethernet (**ae**)
- Gigabit Ethernet (**ge**)
- Interface-range configuration (**interface-range**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Routed VLAN interface (RVI) (**vlan**)
- Virtual management Ethernet (**vme**)
- 10-Gigabit Ethernet (**xe**)

If you are using an interface range, the supported protocol families are the ones supported by the interface types that compose the range.

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

**Table 145: Protocol Families and Supported Interface Types**

Family	Description	Supported Interface Types						
		ae	ge	lo0	me0	vlan	vme	xe
<b>ccc</b>	Circuit cross-connect protocol family	✓*	✓					✓
<b>ethernet-switching</b>	Ethernet switching protocol family	✓	✓		✓			✓
<b>inet</b>	IPv4 protocol family	✓	✓	✓	✓	✓	✓	✓
<b>inet6</b>	IPv6 protocol family	✓	✓	✓	✓	✓	✓	✓
<b>iso</b>	Junos OS protocol family for IS-IS traffic	✓	✓	✓	✓	✓	✓	✓
<b>mpls</b>	MPLS protocol family	✓	✓	✓	✓		✓	✓

\*Supported on EX8200 switches only

The remaining statements are explained separately.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring a DHCP Server on Switches (CLI Procedure)</i></li><li>• <i>Example: Configuring MPLS on EX Series Switches</i></li><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure) on page 2667</a></li><li>• <i>Configuring Routed VLAN Interfaces (CLI Procedure)</i></li></ul>

---

## file

---

<b>Syntax</b>	<code>file certificate-<i>filename</i>;</code>
<b>Hierarchy Level</b>	[edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Specify the file from which to read the digital certificate.
<b>Options</b>	<i>certificate-filename</i> —File from which to read the digital certificate.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>

---

## ftp

---

<b>Syntax</b>	<pre>ftp {     connection-limit limit;     rate-limit limit; }</pre>
<b>Hierarchy Level</b>	[edit system <a href="#">services</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Allow FTP requests from remote systems to the local router or switch.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring FTP Service for Remote Access to the Router or Switch</i></li></ul>

## http

---

<b>Syntax</b>	<pre>http {     interfaces [ <i>interface-names</i> ];     port <i>port</i>; }</pre>
<b>Hierarchy Level</b>	[edit system <a href="#">services web-management</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the port and interfaces for HTTP service, which is unencrypted.
<b>Options</b>	<p><b>interfaces [ <i>interface-names</i> ]</b>—Name of one or more interfaces on which to allow the HTTP service. By default, HTTP access is allowed through built-in Fast Ethernet or Gigabit Ethernet interfaces only.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Management Access for the EX Series Switch (J-Web Procedure) on page 529</a></li><li>• <a href="#">J-Web Interface User Guide</a></li><li>• <a href="#">https on page 1653</a></li><li>• <a href="#">port on page 1665</a></li><li>• <a href="#">web-management on page 1689</a></li></ul>

## https

<b>Syntax</b>	<pre>https {   interfaces [ <i>interface-names</i> ];   local-certificate <i>name</i>;   port <i>port</i>; }</pre>
<b>Hierarchy Level</b>	[edit system <a href="#">services web-management</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Configure the secure version of HTTP (HTTPS) service, which is encrypted.
<b>Options</b>	<p><b>interfaces [ <i>interface-names</i> ]</b>—Name of one or more interfaces on which to allow the HTTPS service. By default, HTTPS access is allowed through any ingress interface, but HTTP access is allowed through built-in Fast Ethernet or Gigabit Ethernet interfaces only.</p> <p><b>local-certificate <i>name</i></b>—Name of the X.509 certificate for a Secure Sockets Layer (SSL) connection. An SSL connection is configured at the <b>[edit security certificates local]</b> hierarchy.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Management Access for the EX Series Switch (J-Web Procedure) on page 529</a></li> <li>• <a href="#">J-Web Interface User Guide</a></li> <li>• <a href="#">http on page 1652</a></li> <li>• <a href="#">port on page 1665</a></li> <li>• <a href="#">web-management on page 1689</a></li> </ul>

## ldap-url

---

<b>Syntax</b>	<ldap-url <i>url-name</i> >;
<b>Hierarchy Level</b>	[edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) (Optional) Specify the Lightweight Directory Access Protocol (LDAP) URL for digital certificates.
<b>Options</b>	<i>url-name</i> —Name of the LDAP URL.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>


## lease-time

---

<b>Syntax</b>	lease-time ( <i>seconds</i>   infinite);
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices.
<b>Description</b>	Request a specific lease time for the IP address. The lease time is the length of time in seconds that a client holds the lease for an IP address assigned by a DHCP server.
<b>Default</b>	If no lease time is requested by client, then the server sends the lease time. The default lease time on a JUNOS OS DHCP server is one day.
<b>Options</b>	<b>seconds</b> —Request a lease time of a specific duration. <b>Range:</b> 60 through 2147483647 seconds  <b>infinite</b> —Request that the lease never expire.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a DHCP Client (CLI Procedure) on page 1413</a></li> <li>• <i>Example: Configuring the Device as a DHCP Client</i></li> <li>• <i>interfaces</i></li> <li>• <a href="#">unit on page 2831</a></li> <li>• <a href="#">family on page 1647</a></li> </ul>


## load-key-file

---

<b>Syntax</b>	load-key-file <i>URL filename</i> ;
<b>Hierarchy Level</b>	[edit system root-authentication], [edit system login user <i>username</i> authentication]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<div> <b>NOTE:</b> ECDSA is not supported on the QFabric system.</div> <p>Load RSA (SSH version 1 and SSH version 2) and DSA or ECDSA (SSH version 2) public keys from a previously-generated named file at a specified URL location or local path. The file contains one or more SSH keys that are copied into the configuration when the command is issued.</p>
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Root Password</i></li><li>• <i>Configuring the Root Password</i></li><li>• <i>Configuring Junos OS User Accounts</i></li><li>• <i>Configuring Junos OS User Accounts</i></li></ul>



## local

<b>Syntax</b>	<pre>local <i>certificate-name</i> {     <i>certificate-key-string</i>;     load-key-file <i>URL filename</i>; }</pre>
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Import a paired X.509 private key and authentication certificate, to enable Junos XML protocol client applications to establish Secure Sockets Layer (SSL) connections to the router or switch.
<div>  <p><b>NOTE:</b> For FIPS mode, the digital security certificates must be compliant with the National Institute of Standards and Technology (NIST) SP 800-131A standard.</p> </div>	
<b>Options</b>	<p><b><i>certificate-name</i></b><b><i>certificate-key-string</i></b>—String of alphanumeric characters that constitute the private key and certificate.</p> <p><b><i>certificate-name</i></b>—Name that uniquely identifies the certificate.</p> <p><b><i>load-key-file URL filename</i></b>—File that contains the private key and certificate. It can be one of two types of values:</p> <ul style="list-style-type: none"> <li>• Pathname of a file on the local disk (assuming you have already used another method to copy the certificate file to the router's or switch's local disk)</li> <li>• URL to the certificate file location (for instance, on the computer where the Junos XML protocol client application runs)</li> </ul>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Importing SSL Certificates for Junos XML Protocol Support</i></li> </ul>

## local-certificate

---

<b>Syntax</b>	local-certificate;
<b>Hierarchy Level</b>	[edit system services service-deployment], [edit system services web-management https], [edit system services xnm-ssl]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Import or reference an SSL certificate.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</i></li><li>• <a href="#">Generating SSL Certificates to Be Used for Secure Web Access on page 532</a></li><li>• <i>Importing SSL Certificates for Junos XML Protocol Support</i></li></ul>

## maximum-certificates

---

<b>Syntax</b>	maximum-certificates <i>number</i> ;
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the maximum number of peer digital certificates to be cached.
<b>Options</b>	<i>number</i> —Maximum number of peer digital certificates to be cached. <b>Range:</b> 64 through 4,294,967,295 peer certificates <b>Default:</b> 1024 peer certificates
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>

## maximum-hop-count

<b>Syntax</b>	<code>maximum-hop-count <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit forwarding-options helpers bootp], [edit forwarding-options helpers bootpinterface ( <i>interface-name</i>   <i>interface-group</i> )]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Specify the maximum number of hops allowed.
<b>Options</b>	<i>number</i> —Maximum number of hops. <b>Default:</b> 4 hops
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents</i></li> </ul>

## maximum-lease-time (DHCP)

<b>Syntax</b>	<code>maximum-lease-time <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit system services dhcp],
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	For J Series Services Routers and EX Series switches only. Specify the maximum length of time in seconds for which a client can request and hold a lease on a DHCP server.  An exception is that the dynamic BOOTP lease length can exceed the maximum lease length specified.
<b>Options</b>	<i>seconds</i> —The maximum number of seconds the lease can be held.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Router or Interface to Act as a DHCP Server on J Series Services Routers</i></li> <li>• <i>default-lease-time</i></li> </ul>

## minimum-wait-time

---

<b>Syntax</b>	<code>minimum-wait-time seconds;</code>
<b>Hierarchy Level</b>	[edit forwarding-options helpers bootp], [edit forwarding-options helpers bootpinterface ( <i>interface-name</i>   <i>interface-group</i> )]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Specify the minimum time allowed.
<b>Options</b>	<b>seconds</b> —Minimum time. <b>Default:</b> 0 seconds
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents</i></li></ul>

## name-server

---

<b>Syntax</b>	<code>name-server {     <i>address</i>; }</code>
<b>Hierarchy Level</b>	[edit system], [edit system services dhcp], [edit system services dhcp], [edit system services dhcp pool], [edit system services dhcp static-binding]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure one or more Domain Name System (DNS) name servers.
<b>Options</b>	<b>address</b> —Address of the name server. To configure multiple name servers, include a maximum of three <b>address</b> options.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Reaching a Domain Name System Server</i></li><li>• <i>Configuring the Router or Interface to Act as a DHCP Server on J Series Services Routers</i></li><li>• <i>Configuring a DHCP Server on Switches (CLI Procedure)</i></li></ul>

---

## no-listen

---

<b>Syntax</b>	no-listen;
<b>Hierarchy Level</b>	[edit forwarding-options helpers bootp interface ( <i>interface-name</i>   <i>interface-group</i> )], [edit forwarding-options helpers domain interface <i>interface-name</i> ], [edit forwarding-options helpers <b>tftp</b> interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Disable recognition of DNS requests or stop packets from being forwarded on a logical interface, a group of logical interfaces, a router, or a switch.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring DNS and TFTP Packet Forwarding</i></li><li>• <i>Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents</i></li></ul>

## outbound-ssh

---

**Syntax** [edit system services]  
outbound-ssh {  
  client *client-id* {  
    address {  
      port *port-number*;  
      retry *number*;  
      timeout *seconds*;  
    }  
    device-id *device-id*;  
    keep-alive {  
      retry *number*;  
      timeout *seconds*;  
    }  
    reconnect-strategy (in-order | sticky);  
    secret *password*;  
    services netconf;  
  }  
  traceoptions {  
    file filename <files *number*> <match *regex*> <size *size*> <world-readable |  
      no-world-readable>;  
    flag *flag*;  
    no-remote-trace;  
  }  
}

**Hierarchy Level** [edit system services]

**Release Information** Statement introduced in Junos OS Release 8.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure a router or switch running the Junos OS behind a firewall to communicate with client management applications on the other side of the firewall.

**Default** To configure transmission of the router's or switch's device ID to the application, include the **device-id** statement at the [edit system services] hierarchy level.

**Options** **client-id**—Identifies the **outbound-ssh** configuration stanza on the router or switch. Each **outbound-ssh** stanza represents a single outbound SSH connection. This attribute is not sent to the client.

**device-id**—Identifies the router or switch to the client during the initiation sequence.

**keep-alive**—(Optional) When configured, specifies that the router or switch send keepalive messages to the management server. To configure the keepalive message, you must set both the **timeout** and **retry** attributes.

**reconnect-strategy**—(Optional) Specify the method the router or switch uses to reestablish a disconnected outbound SSH connection. Two methods are available:

- **in-order**—Specify that the router or switch first attempt to establish an outbound SSH session based on the management server address list. The router or switch attempts to establish a session with the first server on the list. If this connection is not available, the router or switch attempts to establish a session with the next server, and so on down the list until a connection is established.
- **sticky**—Specify that the router or switch first attempt to reconnect to the management server that it was last connected to. If the connection is unavailable, it attempts to establish a connection with the next client on the list and so forth until a connection is made.

**retry**—Number of keepalive messages the router or switch sends without receiving a response from the client before the current SSH connection is disconnected. The default is three messages.

**secret**—(Optional) Router's or switch's public SSH host key. If added to the **outbound-ssh** statement, during the initialization of the outbound SSH service, the router or switch passes its public key to the management server. This is the recommended method of maintaining a current copy of the router's or switch's public key.

**timeout**—Length of time that the Junos server waits for data before sending a keep alive signal. The default is 15 seconds.

When reconnecting to a client, the router or switch attempts to reconnect to the client based on the **retry** and **timeout** values for each client listed.

**address**—Hostname or the IPv4 address of the NSM application server. You can list multiple clients by adding each client's IP address or hostname along with the following connection parameters:

- **port**—Outbound SSH port for the client. The default is port 22.
- **retry**—Number of times the router or switch attempts to establish an outbound SSH connection before giving up. The default is three tries.
- **timeout**—Length of time that the router or switch attempts to establish an outbound SSH connection before giving up. The default is fifteen seconds.

**filename**—(Optional) By default, the filename of the log file used to record the trace options is the name of the traced process (for example, **mib2d** or **snmpd**). Use this option to override the default value.

**files**—(Optional) Maximum number of trace files generated. By default, the maximum number of trace files is 10. Use this option to override the default value.

When a trace file reaches its maximum size, the system archives the file and starts a new file. The system archives trace files by appending a number to the filename in sequential order from 1 to the maximum value (specified by the default value or the options value set here). Once the maximum value is reached, the numbering sequence is restarted at 1, overwriting the older file.

**size**—(Optional) Maximum size of the trace file in kilobytes (KB). Once the maximum file size is reached, the system archives the file. The default value is 1000 KB. Use this option to override the default value.

**match**—(Optional) When used, the system only adds lines to the trace file that match the regular expression specified. For example, if the match value is set to **=error**, the system only records lines to the trace file that include the string **error**.

**services**—Services available for the session. Currently, NETCONF is the only service available.

**world-readable | no-world-readable**—(Optional) Whether the files are accessible by the originator of the trace operation only or by any user. By default, log files are only accessible by the user that started the trace operation (**no-world-readable**).

**all | configuration | connectivity**—(Optional) Type of tracing operation to perform.

**all**—Log all events.

**configuration**—Log all events pertaining to the configuration of the router or switch.

**connectivity**—Log all events pertaining to the establishment of a connection between the client server and the router or switch.

**no-remote-trace**—(Optional) Disable remote tracing.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
---------------------------------	---

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Outbound SSH Service</i></li><li>• <i>System Management Configuration Statements</i></li></ul>
------------------------------	---



## path-length

<b>Syntax</b>	<code>path-length <i>certificate-path-length</i>;</code>
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the digital certificate path length.
<b>Options</b>	<b><i>certificate-path-length</i></b> —Digital certificate path length. <b>Range:</b> 2 through 15 certificates <b>Default:</b> 15 certificates
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## port (HTTP/HTTPS)

<b>Syntax</b>	<code>port <i>port-number</i>;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">services web-management</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the port on which the HTTP or HTTPS service is connected.
<b>Options</b>	<b><i>port-number</i></b> —The TCP port number on which the specified service listens.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Table 47 on page 530</a></li> <li>• <i>J-Web Interface User Guide</i></li> <li>• <a href="#">http on page 1652</a></li> <li>• <a href="#">https on page 1653</a></li> <li>• <a href="#">web-management on page 1689</a></li> </ul>

## port (SRC Server)

---

<b>Syntax</b>	<code>port <i>port-number</i>;</code>
<b>Hierarchy Level</b>	[edit system services service-deployment servers <i>server-address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the port number on which to contact the SRC server.
<b>Options</b>	<i>port-number</i> —(Optional) The TCP port number for the SRC server. <b>Default:</b> 3333
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Junos OS to Work with SRC Software</i></li></ul>

## process-inform

---

<b>Syntax</b>	<pre>process-inform {   pool <i>pool-name</i> network <i>address-range</i> }</pre>
<b>Hierarchy Level</b>	[edit system services dhcp-local-server overrides]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	For extended Dynamic Host Configuration Protocol (DHCP) servers, enable the processing of DHCP information request messages sent from the client to the server to request DHCP options. The messages are also passed to the configured server list.
<b>Default</b>	Information request messages are not processed.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring a DHCP Server on Switches (CLI Procedure)</i></li></ul>

---

## protocol-version

---

<b>Syntax</b>	<code>protocol-version <i>version</i>;</code>
<b>Hierarchy Level</b>	[edit system services ssh]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the secure shell (SSH) protocol version.
<b>Default</b>	<b>v2</b> —SSH protocol version 2 is the default, introduced in Junos OS Release 11.4.
<b>Options</b>	<b><i>version</i></b> —SSH protocol version: <b>v1</b> , <b>v2</b> , or both.
<b>Required Privilege Level</b>	<b>admin</b> —To view this statement in the configuration. <b>admin-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring SSH Service for Remote Access to the Router or Switch</i></li></ul>

## rate-limit

---

<b>Syntax</b>	<code>rate-limit <i>limit</i>;</code>
<b>Hierarchy Level</b>	[edit system services finger], [edit system services ftp], [edit system services netconf ssh], [edit system services ssh], [edit system services telnet], [edit system services xnm-clear-text], [edit system services xnm-ssl]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the maximum number of connections attempts per protocol (either IPv6 or IPv4) on an access service.
<b>Default</b>	150 connections
<b>Options</b>	<b>rate-limit <i>limit</i></b> —(Optional) Maximum number of connection attempts allowed per minute, per IP protocol (either IPv4 or IPv6). <b>Range:</b> 1 through 250 <b>Default:</b> 150
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</i></li></ul>

## reconfigure

<b>Syntax</b>	<pre>reconfigure {     attempts <i>attempts</i>;     clear-on-abort;     timeout interval;     token token; }</pre>
<b>Hierarchy Level</b>	[edit system services <a href="#">dhcp-local-server</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	For extended Dynamic Host Configuration Protocol (DHCP) servers, enable dynamic reconfiguration triggered by the server of all DHCP clients.
<b>Options</b>	<p><b>attempts</b>—Number of attempts made to reconfigure all DHCP clients.</p> <p><b>clear-on-abort</b>—Delete all DHCP clients when reconfiguration fails; that is, when the maximum number of retry attempts have been made without success.</p> <p><b>timeout interval</b>—Initial value (in seconds) between attempts to reconfigure all DHCP clients. Each successive attempt doubles the interval between attempts. For example, if the first value is 2, the first retry is attempted 2 seconds after the first attempt fails. The second retry is attempted 4 seconds after the first retry fails. The third retry is attempted 8 seconds after the second retry fails, and so on. A group configuration takes precedence over a DHCP local server configuration.</p> <p><b>token</b>—Configure a plain-text token for all DHCP clients. The token enables rudimentary entity authentication to protect against inadvertently instantiated DHCP servers. A null token (empty string) indicates that the configuration token functionality is not enabled.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring a DHCP Server on Switches (CLI Procedure)</i></li> </ul>

## retransmission-attempt

---

<b>Syntax</b>	<code>retransmission-attempt <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices.
<b>Description</b>	Specify the number of times the device retransmits a Dynamic Host Control Protocol (DHCP) packet if a DHCP server fails to respond. After the specified number of attempts, no further attempts at reaching a server are made.
<b>Options</b>	<b><i>number</i></b> —Number of retransmit attempts.. <b>Range:</b> 0 through 6 <b>Default:</b> 4
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a DHCP Client (CLI Procedure) on page 1413</a></li><li>• <a href="#">Example: Configuring the Device as a DHCP Client</a></li><li>• <a href="#">interfaces</a></li><li>• <a href="#">unit on page 2831</a></li><li>• <a href="#">family on page 1647</a></li></ul>

## retransmission-interval

<b>Syntax</b>	<code>retransmission-interval seconds;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>family</b> inet dhcp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the time between successive retransmissions of the client DHCP request if a DHCP server fails to respond.
<b>Options</b>	<p><b>seconds</b>—Number of seconds between successive retransmissions.</p> <p><b>Range:</b> 4 through 64 seconds</p> <p><b>Default:</b> 4 seconds</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a DHCP Client (CLI Procedure) on page 1413</a></li> </ul>

## server (DNS and TFTP Service)

<b>Syntax</b>	<code>server address &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt;;</code>
<b>Hierarchy Level</b>	<p>[edit forwarding-options helpers domain],</p> <p>[edit forwarding-options helpers domain interface <i>interface-name</i>],</p> <p>[edit forwarding-options helpers <b>tftp</b>],</p> <p>[edit forwarding-options helpers <b>tftp</b> interface <i>interface-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Specify the DNS or TFTP server for forwarding DNS or TFTP requests. Only one server can be specified for each interface.
<b>Options</b>	<p><b>address</b>—Address of the server.</p> <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Logical system of the server.</p> <p><b>routing-instance [ <i>routing-instance-names</i> ]</b>—(Optional) Set the routing instance name or names that belong to the DNS server or TFTP server.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring DNS and TFTP Packet Forwarding</a></li> </ul>

## server-address

---

<b>Syntax</b>	<code>server-address <i>ip-address</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices.
<b>Description</b>	Specify the address of the DHCP server that the client should accept DHCP offers from. If this option is included in the DHCP configuration, the client accepts offers only from this server and ignores all other offers.
<b>Default</b>	The client accepts the first offer it receives from any DHCP server.
<b>Options</b>	<i>ip-address</i> —DHCP server address.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a DHCP Client (CLI Procedure) on page 1413</a></li><li>• <i>Example: Configuring the Device as a DHCP Client</i></li><li>• <i>interfaces</i></li><li>• <a href="#">unit on page 2831</a></li><li>• <a href="#">family on page 1647</a></li></ul>



## server-identifier

<b>Syntax</b>	<code>server-identifier address;</code>
<b>Hierarchy Level</b>	[edit system services dhcp], [edit system services dhcp pool], [edit system services dhcp static-binding]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	<p>For J Series Services Routers and EX Series switches only. Configure a server identifier. The identifier can be used to identify a DHCP server in a DHCP message. It can also be used as a destination address from clients to servers (for example, when the boot file is set, but not the boot server).</p> <p>Servers include the server identifier in <b>DHCPOFFER</b> messages so that clients can distinguish between multiple lease offers. Clients include the server identifier in <b>DHCPREQUEST</b> messages to select a lease and indicate which offer is accepted from multiple lease offers. Also, clients can use the server identifier to send unicast request messages to specific DHCP servers to renew a current lease.</p> <p>This address must be a manually assigned, static IP address. The server cannot send a request and receive an IP address from itself or another DHCP server.</p>
<b>Default</b>	If no server identifier is set, the DHCP server sets the server identifier based on the primary interface address used by the server to receive a client request. For example, if the client sends a DHCP request and the server receives it on <b>fe-0/0/0</b> and the primary interface address is <b>1.1.1.1</b> , then the server identifier is set to <b>1.1.1.1</b> .
<b>Options</b>	<b>address</b> —IPv4 address of the server. This address must be accessible by all clients served within a specified range of addresses (based on an address pool or static binding).
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring the Router or Interface to Act as a DHCP Server on J Series Services Routers</i></li> </ul>

## servers

---

<b>Syntax</b>	<code>servers server-address {     port port-number; }</code>
<b>Hierarchy Level</b>	[edit system services <a href="#">service-deployment</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure an IPv4 address for the Session and Resource Control (SRC) server.
<b>Options</b>	<b>server-address</b> —The TCP port number. <b>Default:</b> 3333  The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Junos OS to Work with SRC Software</i></li></ul>

## service-deployment

---

<b>Syntax</b>	<code>service-deployment {     <a href="#">servers</a> server-address {         port port-number;     }     <a href="#">source-address</a> source-address; }</code>
<b>Hierarchy Level</b>	[edit system services]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Enable Junos OS to work with the Session and Resource Control (SRC) software.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Junos OS to Work with SRC Software</i></li></ul>

## services (System Services)

```
Syntax  services {
        dhcp { \* DHCP not supported on a DCF
            dhcp_services;
        }
        finger {
            connection-limit limit;
            rate-limit limit;
        }
        ftp {
            connection-limit limit;
            rate-limit limit;
        }
        service-deployment {
            servers address {
                port-number port-number;
            }
            source-address address;
        }
        ssh {
            connection-limit limit;
            protocol-version [v1 v2];
            rate-limit limit;
            root-login (allow | deny | deny-password);
        }
        telnet {
            connection-limit limit;
            rate-limit limit;
        }
        web-management {
            http {
                interfaces [ names ];
                port port;
            }
            https {
                interfaces [ names ];
                local-certificate name;
                port port;
            }
            session {
                idle-timeout [ minutes ];
                session-limit [ limit ];
            }
        }
        xnm-clear-text {
            connection-limit limit;
            rate-limit limit;
        }
        xnm-ssl {
            connection-limit limit;
            local-certificate name;
            rate-limit limit;
            ssl-renegotiation;
        }
    }
```

```
}  
}
```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure the router or switch so that users on remote systems can access the local router or switch through the DHCP server, finger, rlogin, SSH, telnet, Web management, Junos XML protocol clear-text, Junos XML protocol SSL, and network utilities or enable Junos OS to work with the Session and Resource Control (SRC) software.

The remaining statements are explained separately.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring clear-text or SSL Service for Junos XML Protocol Client Applications*
- *Configuring the Router or Interface to Act as a DHCP Server on J Series Services Routers*
- *Configuring the Junos OS to Work with SRC Software*

## session (Time-out)

---

<b>Syntax</b>	<pre>session {     idle-timeout <i>minutes</i>;     session-limit <i>session-limit</i>; }</pre>
<b>Hierarchy Level</b>	[edit system <a href="#">services web-management</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Configure limits for the number of minutes a session can be idle before it times out, and configure the number of simultaneous J-Web user login sessions.
<b>Options</b>	<p><b>idle-timeout <i>minutes</i></b>—Configure the number of minutes a session can be idle before it times out.</p> <p><b>Range:</b> 1 through 1440</p> <p><b>Default:</b> 1440</p> <p><b>session-limit <i>session-limit</i></b>—Configure the maximum number of simultaneous J-Web user login sessions.</p> <p><b>Range:</b> 1 through 1024</p> <p><b>Default:</b> Unlimited</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>J-Web Interface User Guide</i></li> </ul>

## sip-server

---

<b>Syntax</b>	<code>sip-server [address   name];</code>
<b>Hierarchy Level</b>	[edit system services dhcp], [edit system <a href="#">services</a> dhcp], [edit system <a href="#">services</a> dhcp pool], [edit system <a href="#">services</a> dhcp <a href="#">static-binding</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for EX Series switches.
<b>Description</b>	Configure Session Initiation Protocol (SIP) server addresses or names for DHCP servers.
<b>Options</b>	<p><b>address</b>—IPv4 address of the SIP server. To configure multiple SIP servers, include multiple <b>address</b> options. This address must be accessible by all clients served within a specified range of addresses (based on an address pool or static binding).</p> <p><b>name</b>—Fully qualified domain name of the SIP server. To configure multiple SIP servers, include multiple <b>name</b> options. This domain name must be accessible by all clients served within a specified range of addresses (based on an address pool or static binding).</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a DHCP SIP Server (CLI Procedure) on page 1412</a></li><li>• <a href="#">Configuring a DHCP Server on Switches (CLI Procedure)</a></li></ul>

## source-address (SRC Software)

---

<b>Syntax</b>	<code>source-address source-address;</code>
<b>Hierarchy Level</b>	[edit system services service-deployment]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Enable Junos OS to work with the Session and Resource Control (SRC) software.
<b>Options</b>	<b>source-address</b> — Local IPv4 address to be used as source address for traffic to the SRC server. The source address restricts traffic within the out-of-band network.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Junos OS to Work with SRC Software</a></li></ul>

## source-address-giaddr

<b>Syntax</b>	source-address-giaddr;
<b>Hierarchy Level</b>	[edit forwarding-options helpers bootp], [edit forwarding-options helpers bootp interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for EX Series switches.
<b>Description</b>	<p>Configure the gateway IP address (giaddr) as the source IP address of the switch for relayed DHCP packets when the switch is used as the DHCP relay agent.</p> <p>When this statement is entered in the [edit forwarding-options helpers bootp] hierarchy, the gateway IP address is configured as the source IP address of the switch for relayed DHCP packets exiting all interfaces on the switch.</p> <p>When this statement is entered in the [edit forwarding-options helpers bootp interface <i>interface-name</i>] hierarchy, the gateway IP address is configured as the source IP address of the switch for relayed DHCP packets exiting the specified interface of the switch.</p> <p>In Junos OS Release 10.1 for EX Series switches and later releases, the IP address of the interface that the DHCP packet exits on the switch acting as a DHCP relay agent is used as the source IP address for relayed DHCP packets by default.</p> <p>In Junos OS Releases 9.6 and 10.0 for EX Series switches, the gateway IP address of the switch is always used as the source IP address for relayed DHCP packets when the switch is used as the DHCP relay agent.</p> <p>In Junos OS Releases 9.3 through 9.5 for EX Series switches, the IP address of the interface that the DHCP packet exits on the switch acting as a DHCP relay agent is always used as the source IP address for relayed DHCP packets.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>DHCP/BOOTP Relay for Switches Overview</i></li> </ul>

## ssh

---

**Syntax**    ssh {  
              ciphers [ *cipher-1 cipher-2 cipher-3 ...*];  
              client-alive-count-max *seconds*;  
              client-alive-interval *seconds*;  
              connection-limit *limit*;  
              hostkey-algorithm <*algorithm*|no-*algorithm*>;  
              key-exchange <*algorithm*>;  
              macs <*algorithm*>;  
              max-sessions-per-connection <*number*>;  
              no-passwords;  
              no-tcp-forwarding;  
              protocol-version [*v1 v2*];  
              rate-limit *limit*;  
              root-login (*allow* | *deny* | *deny-password*);  
              }

**Hierarchy Level**    [edit system services]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Statement introduced in Junos OS Release 11.1 for the QFX Series.  
                              **client-alive-interval** and **client-alive-max-count** statements introduced in Junos OS Release 12.2.  
                              **no-passwords** statement introduced in Junos OS Release 13.3.

**Description**    Allow SSH requests from remote systems to the local router or switch.  
  
                      The remaining statements are explained separately.

**Required Privilege Level**    system—To view this statement in the configuration.  
                                  system-control—To add this statement to the configuration.

**Related Documentation**    • *Configuring SSH Service for Remote Access to the Router or Switch*



## static-binding

<b>Syntax</b>	<pre>static-binding <i>mac-address</i> {     client-identifier (ascii <i>client-id</i>   hexadecimal <i>client-id</i>);     fixed-address {         <i>address</i>;     }     host-name <i>client-hostname</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit system services dhcp], [edit system services dhcp]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>For J Series Services routers and EX Series switches only. Set static bindings for DHCP clients. A static binding is a mapping between a fixed IP address and the client's MAC address or client identifier.</p>
<b>Options</b>	<p><b><i>mac-address</i></b>—The MAC address of the client. This is a hardware address that uniquely identifies a client on the network.</p> <p><b><i>fixed-address address</i></b>—Fixed IP address assigned to the client. Typically a client has one address assigned, but you can assign more.</p> <p><b><i>host-name client-hostname</i></b>—Hostname of the client requesting the DHCP server. The name can include the local domain name. Otherwise, the name is resolved based on the <b><i>domain-name</i></b> statement.</p> <p><b><i>client-identifier (ascii client-id   hexadecimal client-id)</i></b>—Used by the DHCP server to index the database of address bindings. The client identifier is an ASCII string or hexadecimal number and can include a type-value pair as specified in RFC 1700, <i>Assigned Numbers</i>. Either a client identifier or the client's MAC address must be configured to uniquely identify the client on the network.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Router or Interface to Act as a DHCP Server on J Series Services Routers</i></li> <li>• <i>Configuring a DHCP Server on Switches (CLI Procedure)</i></li> </ul>

## system-generated-certificate

---

<b>Syntax</b>	system-generated-certificate;
<b>Hierarchy Level</b>	[edit system <a href="#">services web-management https</a> ]
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for EX Series switches.
<b>Description</b>	Configure the automatically generated self-signed certificate for enabling HTTPS services..
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling HTTPS and XNM-SSL Services on Switches Using Self-Signed Certificates (CLI Procedure)</a> on page 1417</li></ul>

## telnet

---

<b>Syntax</b>	telnet { <a href="#">connection-limit</a> <i>limit</i> ; <a href="#">rate-limit</a> <i>limit</i> ; }
<b>Hierarchy Level</b>	[edit system services]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Provide Telnet connections from remote systems to the local router or switch.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Telnet Service for Remote Access to a Router or Switch</i></li></ul>

## tftp

<b>Syntax</b>	<pre> tftp {   description <i>text-description</i>;   interface <i>interface-name</i> {     broadcast;     description <i>text-description</i>;     no-listen;     server address &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance       <i>routing-instance-name</i>&gt;;   }   server address &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance     <i>routing-instance-name</i>&gt;; } </pre>
<b>Hierarchy Level</b>	[edit forwarding-options helpers]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Enable TFTP request packet forwarding.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring DNS and TFTP Packet Forwarding</i></li> </ul>

## traceoptions (DNS and TFTP Packet Forwarding)

---

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;size <i>bytes</i>&gt; &lt;world-readable           no-world-readable&gt;;     flag <i>flag</i>;     level <i>level</i>;     &lt;no-remote-trace&gt;; }</pre>
<b>Hierarchy Level</b>	[edit forwarding-options helpers]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement standardized and <b>match</b> option introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure tracing operations for BOOTP, DNS and TFTP packet forwarding.
<b>Default</b>	If you do not include this statement, no tracing operations are performed.
<b>Options</b>	<p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" "). All files are placed in a file named <b>fud</b> in the directory <b>/var/log</b>. If you include the <b>file</b> statement, you must specify a filename.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option and a filename.</p> <p><b>Range:</b> 2 through 1000 <b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>address</b>—Trace address management events</li><li>• <b>all</b>—Trace all events</li><li>• <b>bootp</b>—Trace BOOTP or DHCP services events</li><li>• <b>config</b>—Trace configuration events</li><li>• <b>domain</b>—Trace DNS service events</li><li>• <b>ifdb</b>—Trace interface database operations</li><li>• <b>io</b>—Trace I/O operations</li><li>• <b>main</b>—Trace main loop events</li><li>• <b>port</b>—Trace arbitrary protocol events</li></ul>

- **rtsock**—Trace routing socket operations
- **tftp**—Trace TFTP service events
- **trace**—Trace tracing operations
- **ui**—Trace user interface operations
- **util**—Trace miscellaneous utility operations

**match *regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

**no-remote-trace**—(Optional) Disable remote tracing globally or for a specific tracing operation.

**no-world-readable**—(Optional) Restrict file access to the owner.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** file again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 0 bytes through 4,294,967,295 KB

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

<b>Required Privilege</b>	interface—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Tracing BOOTP, DNS, and TFTP Forwarding Operations</i></li> </ul>

## traceoptions

---

**Syntax**    traceoptions {  
              file *filename* <files *number*> <size *size*>;  
              flag all;  
              flag certificates;  
              flag database;  
              flag general;  
              flag ike;  
              flag parse;  
              flag policy-manager;  
              flag routing-socket;  
              flag timer;  
              level  
              no-remote-trace  
          }

**Hierarchy Level**    [edit security],  
                          [edit services ipsec-vpn]

Trace options can be configured at either the [edit security] or the [edit services ipsec-vpn] hierarchy level, but not at both levels.

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure security trace options.

To specify more than one trace option, include multiple **flag** statements. Trace option output is recorded in the `/var/log/kmd` file.



**NOTE:** The `traceoptions` statement is not supported on QFabric systems.

---

**Options**    **files *number***—(Optional) Maximum number of trace files. When a trace file (for example, **kmd**) reaches its maximum size, it is renamed **kmd.0**, then **kmd.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 0 files

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB). When a trace file (for example, **kmd**) reaches this size, it is renamed, **kmd.0**, then **kmd.1** and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

**Default:** 1024 KB

**flag *flag***—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.

- **all**—Trace all security events.
- **certificates**—Trace certificate events.
- **database**—Trace database events.
- **general**—Trace general events.
- **ike**—Trace IKE module processing.
- **parse**—Trace configuration processing.
- **policy-manager**—Trace policy manager processing.
- **routing-socket**—Trace routing socket messages.
- **timer**—Trace internal timer events.

**level *level***—(Optional) Set traceoptions level.

- **all**—match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match conditions that should be handled specially.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

**no-remote-trace**—(Optional) Disable remote tracing

<b>Required Privilege</b>	admin—To view the configuration.
<b>Level</b>	admin-control—To add this statement to the configuration.

<b>Related Documentation</b>	• <i>Configuring Tracing Operations for Security Services</i>
------------------------------	---

## update-server

---

<b>Syntax</b>	update-server;
<b>Hierarchy Level</b>	[edit Interfaces <i>interface-name</i> unit <i>logical-unit-number</i> inet dhcp]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices.
<b>Description</b>	Propagate TCP/IP settings learned from an external DHCP server to the DHCP server running on the switch, router, or device.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a DHCP Client (CLI Procedure) on page 1413</a></li><li>• <i>Example: Configuring the Device as a DHCP Client</i></li><li>• <i>interfaces</i></li><li>• <a href="#">unit on page 2831</a></li><li>• <a href="#">family on page 1647</a></li></ul>



## web-management

<b>Syntax</b>	<pre> web-management {   http {     interfaces [ <i>interface-names</i> ];     port <i>port</i>;   }   https {     interfaces [ <i>interface-names</i> ];     local-certificate <i>name</i>;     port <i>port</i>;   } } </pre>
<b>Hierarchy Level</b>	[edit system <a href="#">services</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Configure settings for HTTP or HTTPS access. HTTP access allows management of the router or switch using the browser-based J-Web graphical user interface. HTTPS access allows secure management of the router or switch using the J-Web interface. With HTTPS access, communication between the router or switch Web server and your browser is encrypted.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Table 47 on page 530</a></li> <li>• <a href="#">J-Web Interface User Guide</a></li> <li>• <a href="#">http on page 1652</a></li> <li>• <a href="#">https on page 1653</a></li> <li>• <a href="#">port on page 1665</a></li> </ul>

## wins-server (System)

---

<b>Syntax</b>	<code>wins-server {     address; }</code>
<b>Hierarchy Level</b>	[edit system services dhcp], [edit system services dhcp], [edit system services dhcp pool], [edit system services dhcp static-binding]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	For J Series Services Routers and EX Series switches only. Specify one or more NetBIOS Name Servers. When a DHCP client is added to the network and assigned an IP address, the NetBIOS Name Server manages the Windows Internet Name Service (WINS) database that matches IP addresses (such as <b>192.168.1.3</b> ) to Windows NetBIOS names (such as <b>\\Marketing</b> ). List servers in order of preference.
<b>Options</b>	<b>address</b> —IPv4 address of the NetBIOS Name Server running WINS. To configure multiple servers, include multiple <b>address</b> options.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Router or Interface to Act as a DHCP Server on J Series Services Routers</i></li><li>• <i>Configuring a DHCP Server on Switches (CLI Procedure)</i></li></ul>

## CHAPTER 26

# Administration

- [Routine Monitoring on page 1691](#)
- [Verifying and Managing DHCP Local Server Configurations on page 1695](#)
- [Verifying and Managing DHCP Relay Agent Configurations on page 1697](#)
- [DHCP Local Server Monitoring Commands on page 1697](#)
- [DHCP Relay Agent Monitoring Commands on page 1732](#)
- [Other Operational Commands on page 1788](#)

### Routine Monitoring

---

- [Monitoring DHCP Services on page 1691](#)

### Monitoring DHCP Services

#### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

A switch or router can operate as a DHCP server. Use the monitoring functionality to view information about dynamic and static DHCP leases, conflicts, pools, and statistics.

**Action** To monitor the DHCP server in the J-Web interface, select **Monitor > Services > DHCP**.

To monitor the DHCP server in the CLI, enter the following CLI commands:

- `show system services dhcp binding`
- `show system services dhcp conflict`
- `show system services dhcp pool`
- `show system services dhcp statistics`
- `show system services dhcp relay-statistics`
- `show system services dhcp global`
- `show system services dhcp client`
- `clear system services dhcp binding`

- `clear system services dhcp conflict`
- `clear system services dhcp statistics`
- `clear dhcp relay-statistics`

On EX4300 switches, to monitor the DHCP server in the CLI, enter the following CLI commands:

- `show dhcp server binding`
- `show dhcp server statistics`
- `show dhcp relay binding`
- `show dhcp relay statistics`
- `clear dhcp server binding`
- `clear dhcp server statistics`
- `clear dhcp relay binding`
- `clear dhcp relay statistics`

**Meaning** [Table 146 on page 1692](#) summarizes the output fields in DHCP displays in the J-Web interface.

**Table 146: Summary of DHCP Output Fields**

Field	Values	Additional Information
Global tab		
Name	This column displays the following information: <ul style="list-style-type: none"><li>• Boot lease length</li><li>• Domain Name</li><li>• Name servers</li><li>• Server identifier</li><li>• Domain search</li><li>• Gateway routers</li><li>• WINS server</li><li>• Boot file</li><li>• Boot server</li><li>• Default lease time</li><li>• Minimum lease time</li><li>• Maximum lease time</li></ul>	
Value	Displays the value for each of the parameters in the Name column.	
Bindings tab		
Allocated Address	List of IP addresses the DHCP server has assigned to clients.	

Table 146: Summary of DHCP Output Fields (*continued*)

Field	Values	Additional Information
MAC Address	Corresponding media access control (MAC) address of the client.	
Binding Type	Type of binding assigned to the client: <b>dynamic</b> or <b>static</b> .	DHCP servers can assign a dynamic binding from a pool of IP addresses or a static binding to one or more specific IP addresses.
Lease Expires	Date and time the lease expires, or <b>never</b> for leases that do not expire.	
<b>Pools tab</b>		
Pool Name	Subnet on which the IP address pool is defined.	
Low Address	Lowest address in the IP address pool.	
High Address	Highest address in the IP address pool.	
Excluded Addresses	Addresses excluded from the address pool.	
<b>Clients tab</b>		
Interface Name	Name of the logical interface.	
Hardware Address	Vendor identification.	
Status	State of the client binding.	
Address Obtained	IP address obtained from the DHCP server.	
Update Server	Indicates whether server update is enabled.	
Lease Obtained	Date and time the lease was obtained.	
Lease Expires	Date and time the lease expires.	
Renew	Reacquires an IP address from the server for the interface. When you click this option, the command sends a discover message if the client state is INIT and a renew request message if the client state is BOUND. For all other states it performs no action.	

Table 146: Summary of DHCP Output Fields (*continued*)

Field	Values	Additional Information
Release	Clears other resources received earlier from the server, and reinitializes the client state to INIT for the particular interface.	
<b>Conflicts tab</b>		
Detection Time	Date and time the client detected the conflict.	
Detection Method	How the conflict was detected.	Only client-detected conflicts are displayed.
Address	IP address where the conflict occurs.	The addresses in the conflicts list remain excluded until you use the <b>clear system services dhcp conflict</b> command to manually clear the list.
<b>DHCP Statistics</b>		
<b>Relay Statistics tab</b>		
Packet Counters	Displays the number of packet counters.	
Dropped Packet Counters	Graphically displays the number of dropped packet counters.	
<b>Statistics tab</b>		
Packets dropped	Total number of packets dropped and the number of packets dropped due to a particular condition.	
Messages received	Number of BOOTREQUEST, DHCPDECLINE, DHCPDISCOVER, DHCPINFORM, DHCPRELEASE, and DHCPREQUEST messages sent from DHCP clients and received by the DHCP server.	
Messages sent	Number of BOOTREPLY, DHCPACK, DHCPOFFER, DHCPNAK, and DHCPFORCERENEW messages sent from the DHCP server to DHCP clients.	

[Table 147 on page 1694](#) summarizes the output fields in DHCP displays in EX4300 switches in the J-Web interface.

Table 147: Summary of DHCP Output Fields for EX4300 Switches

Field	Values	Additional Information
<b>Binding Information tab</b>		
IP Address	IP address of the DHCP client..	
Session ID	Session ID of the subscriber session.	

Table 147: Summary of DHCP Output Fields for EX4300 Switches (*continued*)

Field	Values	Additional Information
Hardware Address	Hardware address of the DHCP client.	
Expires	Number of seconds in which the lease expires.	
State	State of the address binding table on the extended DHCP local server: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has an active IP address lease.</li> <li>• <b>FORCERENEW</b>—Client has received the <b>FORCERENEW</b> message from the server.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>RELEASE</b>—Client is releasing the IP address lease.</li> <li>• <b>RENEWING</b>—Client is sending a request to renew the IP address lease.</li> <li>• <b>REQUESTING</b>—Client is requesting a DHCP server.</li> <li>• <b>SELECTING</b>—Client is receiving offers from DHCP servers.</li> </ul>	
Interface	Interface on which the request was received.	

[Table 148 on page 1695](#) summarizes the output fields in DHCP Statistics Information for EX4300 switches in the J-Web interface.

Table 148: Summary of the DHCP Statistics Information Output for EX4300 switches

Field	Values	Additional Information
Message Counters		
Message Counters	Graphically displays the number of messages sent and received.	
Dropped packet Counters		
MAC Limit	Graphically displays the number of dropped packet counters.	

- Related Documentation**
- [Configuring DHCP Services \(J-Web Procedure\) on page 1405](#)
  - [Understanding DHCP Services for Switches](#)

## Verifying and Managing DHCP Local Server Configurations

- [Verifying and Managing DHCP Local Server Configuration on page 1696](#)
- [Verifying and Managing DHCPv6 Local Server Configuration on page 1696](#)

## Verifying and Managing DHCP Local Server Configuration

**Purpose** View or clear information about client address bindings and statistics for the extended DHCP local server.



**NOTE:** If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

**Action**

- To display the address bindings in the client table on the extended DHCP local server:

```
user@host> show dhcp server binding routing-instance customer routing instance
```

- To display extended DHCP local server statistics:

```
user@host> show dhcp server statistics routing-instance customer routing instance
```

- To clear the binding state of a DHCP client from the client table on the extended DHCP local server:

```
user@host> clear dhcp server binding routing-instance customer routing instance
```

- To clear all extended DHCP local server statistics:

```
user@host> clear dhcp server statistics routing-instance customer routing instance
```

**Related Documentation**

- [CLI Explorer](#)

## Verifying and Managing DHCPv6 Local Server Configuration

**Purpose** View or clear information about client address bindings and statistics for the DHCPv6 local server.

**Action**

- To display the address bindings in the client table on the DHCPv6 local server:

```
user@host> show dhcpv6 server binding
```

- To display DHCPv6 local server statistics:

```
user@host> show dhcpv6 server statistics
```

- To clear all DHCPv6 local server statistics:

```
user@host> clear dhcpv6 server binding
```

- To clear all DHCPv6 local server statistics:

```
user@host> clear dhcpv6 server statistics
```

**Related Documentation**

- [CLI Explorer](#)



## Verifying and Managing DHCP Relay Agent Configurations

---

- [Verifying and Managing DHCP Relay Configuration on page 1697](#)
- [Verifying and Managing DHCPv6 Relay Configuration on page 1697](#)

### Verifying and Managing DHCP Relay Configuration

**Purpose** View or clear address bindings or statistics for extended DHCP relay agent clients:

- Action**
- To display the address bindings for extended DHCP relay agent clients:  
`user@host> show dhcp relay binding routing-instance customer routing instance`
  - To display extended DHCP relay agent statistics:  
`user@host> show dhcp relay statistics routing-instance customer routing instance`
  - To clear the binding state of DHCP relay agent clients:  
`user@host> clear dhcp relay binding routing-instance customer routing instance`
  - To clear all extended DHCP relay agent statistics:  
`user@host> clear dhcp relay statistics routing-instance customer routing instance`

**Related Documentation**

- [CLI Explorer](#)

### Verifying and Managing DHCPv6 Relay Configuration

**Purpose** View or clear address bindings or statistics for extended DHCPv6 relay agent clients:

- Action**
- To display the address bindings for extended DHCPv6 relay agent clients:  
`user@host> show dhcpv6 relay binding`
  - To display extended DHCPv6 relay agent statistics:  
`user@host> show dhcpv6 relay statistics`
  - To clear the binding state of DHCPv6 relay agent clients:  
`user@host> clear dhcpv6 relay binding`
  - To clear all extended DHCPv6 relay agent statistics:  
`user@host> clear dhcpv6 relay statistics`

**Related Documentation**

- [CLI Explorer](#)

## DHCP Local Server Monitoring Commands

---

- `clear dhcp server binding`
- `clear dhcp server statistics`

- `clear dhcpv6 server binding`
- `clear dhcpv6 server statistics`
- `request dhcp server reconfigure`
- `request dhcpv6 server reconfigure`
- `request system reboot`
- `show dhcp server binding`
- `show dhcp server statistics`
- `show dhcpv6 server binding`
- `show dhcpv6 server statistics`

## clear dhcp server binding

**Syntax**    `clear dhcp server binding`  
               `<address>`  
               `<all>`  
               `<interface interface-name>`  
               `<interfaces-vlan>`  
               `<interfaces-wildcard>`  
               `<logical-system logical-system-name>`  
               `<routing-instance routing-instance-name>`

**Release Information**    Command introduced in Junos OS Release 9.0.  
                                  Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

**Description**    Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table on the extended DHCP local server.



**NOTE:** If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

**Options**    **address**—(Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

**all**—(Optional) Clear the binding state for all DHCP clients.

**interface interface-name**—(Optional) Clear the binding state for DHCP clients on the specified interface.



**NOTE:** This option clears all bindings whose initial login requests were received over the specified interface. Dynamic demux login requests are not received over the dynamic demux interface, but rather the underlying interface of the dynamic demux interface. To clear a specific dynamic demux interface, use the *ip-address* or *mac-address* options.

**interfaces-vlan**—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

**interfaces-wildcard**—(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (\*).

**logical-system** *logical-system-name*—(Optional) Clear the binding state for DHCP clients on the specified logical system.

**routing-instance** *routing-instance-name*—(Optional) Clear the binding state for DHCP clients on the specified routing instance.

**Required Privilege Level** view

**Related Documentation**

- [Clearing DHCP Bindings for Subscriber Access](#)
- [show dhcp server binding on page 1716](#)

**List of Sample Output**

- [clear dhcp server binding <ip-address> on page 1700](#)
- [clear dhcp server binding all on page 1700](#)
- [clear dhcp server binding interface on page 1701](#)
- [clear dhcp server binding <interfaces-vlan> on page 1701](#)
- [clear dhcp server binding <interfaces-wildcard> on page 1701](#)

**Output Fields** See [show dhcp server binding](#) for an explanation of output fields.

## Sample Output

### [clear dhcp server binding <ip-address>](#)

The following sample output displays the address bindings in the DHCP client table on the extended DHCP local server before and after the **clear dhcp server binding** command is issued.

```
user@host> show dhcp server binding
```

```
2 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)
```

IP address	Hardware address	Type	Lease expires at
100.20.32.1	90:00:00:01:00:01	active	2007-01-17 11:38:47 PST
100.20.32.3	90:00:00:02:00:01	active	2007-01-17 11:38:41 PST

```
user@host> clear dhcp server binding 10.20.32.1
```

```
user@host> show dhcp server binding
```

```
1 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)
```

IP address	Hardware address	Type	Lease expires at
100.20.32.3	90:00:00:02:00:01	active	2007-01-17 11:38:41 PST

### [clear dhcp server binding all](#)

The following command clears all DHCP local server bindings:

```
user@host> clear dhcp server binding all
```

### clear dhcp server binding interface

The following command clears DHCP local server bindings on a specific interface:

```
user@host> clear dhcp server binding interface fe-0/0/2
```

### clear dhcp server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP local server bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp server binding ae0
```

### clear dhcp server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP local server bindings over a specific interface:

```
user@host> clear dhcp server binding ge-1/0/0.*
```

## clear dhcp server statistics

---

<b>Syntax</b>	<b>clear dhcp server statistics</b> <b>&lt;interface <i>interface-name</i>&gt;</b> <b>&lt;logical-system <i>logical-system-name</i>&gt;</b> <b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0.
<b>Description</b>	Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.
<b>Options</b>	<b>logical-system <i>logical-system-name</i></b> —(Optional) Clear the statistics for DHCP clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.  <b>routing-instance <i>routing-instance-name</i></b> —(Optional) Clear the statistics for DHCP clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">clear dhcp server statistics on page 1702</a>
<b>Output Fields</b>	See <a href="#">show dhcp server statistics</a> for an explanation of output fields.

## Sample Output

### clear dhcp server statistics

The following sample output displays the extended DHCP local server statistics before and after the **clear dhcp server statistics** command is issued.

```
user@host> show dhcp server statistics
Packets dropped:
  Total                1
  Lease Time Violation 1

Messages received:
  BOOTREQUEST          89163
  DHCPDECLINE           0
  DHCPDISCOVER          8110
  DHCPINFORM            0
  DHCPRELEASE           0
  DHCPREQUEST          81053

Messages sent:
  BOOTREPLY             32420
  DHCPOFFER             8110
  DHCPACK               8110
  DHCPNAK               8100

user@host> clear dhcp server statistics
user@host> show dhcp server statistics
```

Packets dropped:	
Total	0
Messages received:	
BOOTREQUEST	0
DHCPDECLINE	0
DHCPDISCOVER	0
DHCPINFORM	0
DHCPRELEASE	0
DHCPREQUEST	0
Messages sent:	
BOOTREPLY	0
DHCPOFFER	0
DHCPACK	0
DHCPNAK	0

## clear dhcpv6 server binding

---

<b>Syntax</b>	<b>clear dhcpv6 server binding</b> <b>&lt;address&gt;</b> <b>&lt;all&gt;</b> <b>&lt;interface <i>interface-name</i>&gt;</b> <b>&lt;interfaces-vlan&gt;</b> <b>&lt;interfaces-wildcard&gt;</b> <b>&lt;logical-system <i>logical-system-name</i>&gt;</b> <b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.6. Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.
<b>Description</b>	Clear the binding state of a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) client from the client table on the extended DHCPv6 local server.
<b>Options</b>	<p><b><i>address</i></b>—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"><li>• <i>CID</i>—The specified Client ID (CID).</li><li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li><li>• <i>session-id</i>—The specified session ID.</li></ul> <p><b><i>all</i></b>—(Optional) Clear the binding state for all DHCPv6 clients.</p> <p><b><i>interface interface-name</i></b>—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.</p> <p><b><i>interfaces-vlan</i></b>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><b><i>interfaces-wildcard</i></b>—(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).</p> <p><b><i>logical-system logical-system-name</i></b>—(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.</p> <p><b><i>routing-instance routing-instance-name</i></b>—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Clearing DHCP Bindings for Subscriber Access</i></li><li>• <a href="#">show dhcpv6 server binding on page 1724</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear dhcpv6 server binding all on page 1705</a> <a href="#">clear dhcpv6 server binding &lt;ipv6-prefix&gt; on page 1705</a>



[clear dhcpv6 server binding interface on page 1705](#)  
[clear dhcpv6 server binding <interfaces-vlan> on page 1705](#)  
[clear dhcpv6 server binding <interfaces-wildcard> on page 1705](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear dhcpv6 server binding all

The following command clears all DHCPv6 local server bindings:

```
user@host> clear dhcpv6 server binding all
```

### clear dhcpv6 server binding <ipv6-prefix>

The following command clears DHCPv6 local server bindings for a specific IPv6 prefix:

```
user@host> clear dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
```

### clear dhcpv6 server binding interface

The following command clears DHCPv6 local server bindings on a specific interface:

```
user@host> clear dhcpv6 server binding interface fe-0/0/2
```

### clear dhcpv6 server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 local server bindings on top of the underlying interface **ae0**, which clears DHCPv6 bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcpv6 server binding interface ae0
```

### clear dhcpv6 server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 local server bindings over a specific interface:

```
user@host> clear dhcpv6 server binding ge-1/0/0.*
```

## clear dhcpv6 server statistics

---

<b>Syntax</b>	<code>clear dhcpv6 server statistics</code> <code>&lt;interface <i>interface-name</i>&gt;</code> <code>&lt;logical-system <i>logical-system-name</i>&gt;</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.6.
<b>Description</b>	Clear all extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(Optional) Clear the statistics for DHCPv6 clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Clear the statistics for DHCPv6 clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show dhcpv6 server statistics on page 1730</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear dhcpv6 server statistics on page 1706</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### clear dhcpv6 server statistics

```
user@host> clear dhcpv6 server statistics
```

## request dhcp server reconfigure

**Syntax** `request dhcp server reconfigure (all | address | interface interface-name | logical-system logical-system-name | routing-instance routing-instance-name)`

**Release Information** Command introduced in Junos OS Release 10.0.  
Command introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** Initiate reconfiguration processing for the specified DHCP clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the **clear dhcp server binding** command.

When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfiguring state and the local server sends a **forcerenew** message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the **forcerenew** message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the **clear-on-abort** statement to configure the client to be cleared when reconfiguration fails.

**Options** **all**—Initiate reconfiguration for all DHCP clients.

***address***—Initiate reconfiguration for DHCP client with the specified IP address or MAC address.

**interface *interface-name***—Initiate reconfiguration for all DHCP clients on this logical interface (clients whose initial login requests were received over the specified interface).



**NOTE:** You cannot use the **interface *interface-name*** option with the **request dhcp server reconfigure** command for DHCP passive clients (clients that are added as a result of DHCP snooped packets). For passive clients, the interface is not guaranteed to be the next-hop interface to the client, as is the case for active clients.

**logical-system *logical-system-name***—Initiate reconfiguration for all DHCP clients on the specified logical system.

**routing-instance *routing-instance-name***—Initiate reconfiguration reconfigured for all DHCP clients in the specified routing instance.

**Required Privilege Level** view

**Related Documentation** • [Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429](#)

**List of Sample Output** [request dhcp server reconfigure on page 1708](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[request dhcp server reconfigure](#)

```
user@host> request dhcp server reconfigure interface fe-0/0/0.100
```

## request dhcpv6 server reconfigure

<b>Syntax</b>	<code>request dhcpv6 server reconfigure (all   <i>address</i>   <i>client-id</i>   interface <i>interface-name</i>   logical-system <i>logical-system-name</i>   routing-instance <i>routing-instance-name</i>   session-id)</code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	<p>Initiate reconfiguration processing for the specified DHCPv6 clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the <b>clear dhcpv6 server binding</b> command.</p> <p>When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfigure state and the local server sends a reconfigure message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the reconfigure message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the <b>clear-on-abort</b> statement to configure the client to be cleared when reconfiguration fails.</p>
<b>Options</b>	<p><b>all</b>—Initiate reconfiguration for all DHCPv6 clients.</p> <p><b><i>address</i></b>—Initiate reconfiguration for DHCPv6 client with the specified IPv6 address.</p> <p><b><i>client-id</i></b>—Initiate reconfiguration for DHCPv6 client with the specified client ID.</p> <p><b>interface <i>interface-name</i></b>—Initiate reconfiguration for all DHCPv6 clients on this logical interface (clients whose initial login requests were received over the specified interface).</p> <p><b>logical-system <i>logical-system-name</i></b>—Initiate reconfiguration for all DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—Initiate reconfiguration reconfigured for all DHCPv6 clients in the specified routing instance.</p> <p><b><i>session-id</i></b>—Initiate reconfiguration for DHCPv6 client with the specified session ID.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Dynamic Client Reconfiguration of Extended Local Server Clients on page 1429</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request dhcpv6 server reconfigure on page 1710</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

request dhcpv6 server reconfigure

```
user@host> request dhcpv6 server reconfigure 2001::2/16
```

## request system reboot

<b>List of Syntax</b>	<a href="#">Syntax on page 1711</a> <a href="#">Syntax (EX Series Switches) on page 1711</a> <a href="#">Syntax (TX Matrix Router) on page 1711</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1711</a> <a href="#">Syntax (MX Series Router) on page 1711</a>
<b>Syntax</b>	<pre>request system reboot &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;media (compact-flash   disk   removable-compact-flash   usb)&gt; &lt;message "<i>text</i>"&gt; &lt;other-routing-engine&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>request system reboot &lt;all-members&gt; &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;local&gt; &lt;media (external   internal)&gt; &lt;member <i>member-id</i>&gt; &lt;message "<i>text</i>"&gt; &lt;other-routing-engine&gt; &lt;slice <i>slice</i>&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>request system reboot &lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt; &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;media (compact-flash   disk)&gt; &lt;message "<i>text</i>"&gt; &lt;other-routing-engine&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>request system reboot &lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt; &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;media (compact-flash   disk)&gt; &lt;message "<i>text</i>"&gt; &lt;other-routing-engine&gt; &lt;partition (1   2   alternate)&gt;</pre>
<b>Syntax (MX Series Router)</b>	<pre>request system reboot &lt;all-members&gt; &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;local&gt;</pre>

```
<media (external | internal)>  
<member member-id>  
<message "text">  
<other-routing-engine>
```

<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Option <b>other-routing-engine</b> introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Option <b>sfc</b> introduced for the TX Matrix Plus router in Junos OS Release 9.6. Option <b>both-routing-engines</b> introduced in Junos OS Release 12.1.
<b>Description</b>	Reboot the software.
<b>Options</b>	<p><b>none</b>—Reboot the software immediately.</p> <p><b>all-chassis</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all routers connected to the TX Matrix or TX Matrix Plus router, respectively.</p> <p><b>all-lcc</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively.</p> <p><b>all-members</b>—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on all members of the Virtual Chassis configuration.</p> <p><b>at <i>time</i></b>—(Optional) Time at which to reboot the software, specified in one of the following ways:</p> <ul style="list-style-type: none"><li>• <b>now</b>—Stop or reboot the software immediately. This is the default.</li><li>• <b>+<i>minutes</i></b>—Number of minutes from now to reboot the software.</li><li>• <b><i>yymmddhhmm</i></b>—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.</li><li>• <b><i>hh:mm</i></b>—Absolute time on the current day at which to stop the software, specified in 24-hour time.</li></ul> <p><b>both-routing-engines</b>—(Optional) Reboot both Routing Engines at the same time.</p> <p><b>in <i>minutes</i></b>—(Optional) Number of minutes from now to reboot the software. This option is an alias for the <b>at +<i>minutes</i></b> option.</p> <p><b>lcc <i>number</i></b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis number. Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"><li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li><li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li></ul>



- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the local Virtual Chassis member.

**media (compact-flash | disk | removable-compact-flash | usb)**—(Optional) Boot medium for next boot. (The options **removable-compact-flash** and **usb** pertain to the J Series routers only.)

**media (external | internal)**—(EX Series switches and MX Series routers only) (Optional) Reboot the boot media:

- **external**—Reboot the external mass storage device.
- **internal**—Reboot the internal flash device.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

**message "*text*"**—(Optional) Message to display to all system users before stopping or rebooting the software.

**other-routing-engine**—(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is rebooted.

**partition**—(TX Matrix Plus routers only) (Optional) Reboot using the specified partition on the boot media. This option has the following suboptions:

- 1—Reboot from partition 1.
- 2—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition.

**scc**—(TX Matrix routers only) (Optional) Reboot the Routing Engine on the TX Matrix switch-card chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Reboot the Routing Engine on the TX Matrix Plus switch-fabric chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted. Replace ***number*** with 0.

**slice *slice***—(EX Series switches only) (Optional) Reboot a partition on the boot media. This option has the following suboptions:

- 1—Power off partition 1.
- 2—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

**Additional Information** Reboot requests are recorded in the system log files, which you can view with the **show log** command (see [show log](#)). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the **show system processes** command (see [show system processes](#)).

On a TX Matrix or TX Matrix Plus router, if you issue the **request system reboot** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are rebooted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are rebooted.



**NOTE:** Before issuing the **request system reboot** command on a TX Matrix Plus router with no options or the **all-chassis**, **all-lcc**, **lcc number**, or **sfc** options, verify that master Routing Engine for all routers in the routing matrix are in the same slot number. If the master Routing Engine for a line-card chassis is in a different slot number than the master Routing Engine for a TX Matrix Plus router, the line-card chassis might become logically disconnected from the routing matrix after the **request system reboot** command.



**NOTE:** To reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the master Routing Engine.

**Required Privilege Level** maintenance

**Related Documentation**

- [clear system reboot on page 84](#)
- [request system halt on page 101](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

- [request system reboot on page 1715](#)
- [request system reboot \(at 2300\) on page 1715](#)
- [request system reboot \(in 2 Hours\) on page 1715](#)
- [request system reboot \(Immediately\) on page 1715](#)
- [request system reboot \(at 1:20 AM\) on page 1715](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system reboot

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

### request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes
```

```
shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

### request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
user@host> request system reboot in 120
user@host> request system reboot at 19:00
```

### request system reboot (Immediately)

```
user@host> request system reboot at now
```

### request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

## show dhcp server binding

---

**Syntax**    `show dhcp server binding`  
              `<address>`  
              `<interfaces-vlan><brief | detail | summary>`  
              `<interface interface-name>`  
              `<interfaces-vlan>`  
              `<interfaces-wildcard>`  
              `<logical-system logical-system-name>`  
              `<routing-instance routing-instance-name>`

**Release Information**    Command introduced in Junos OS Release 9.0.  
                              Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

**Description**    Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol (DHCP) local server.



**NOTE:** If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

**Options**    **address**—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

**brief | detail | summary**—(Optional) Display the specified level of output about active client bindings. The default is **brief**, which produces the same output as **show dhcp server binding**.

**interface interface-name**—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

**interfaces-vlan**—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.

**interfaces-wildcard**—(Optional) The set of interfaces on which to show the binding state information. This option supports the use of the wildcard character (\*).

**logical-system logical-system-name**—(Optional) Display information about active client bindings for DHCP clients on the specified logical system.

**routing-instance routing-instance-name**—(Optional) Display information about active client bindings for DHCP clients on the specified routing instance.

**Required Privilege Level** view

**Related Documentation**

- *Clearing DHCP Bindings for Subscriber Access*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- [clear dhcp server binding on page 1699](#)

**List of Sample Output**

[show dhcp server binding on page 1719](#)  
[show dhcp server binding detail on page 1719](#)  
[show dhcp server binding detail \(ACI Interface Set Configured\) on page 1719](#)  
[show dhcp server binding interface <vlan-id> on page 1720](#)  
[show dhcp server binding interface <svlan-id> on page 1720](#)  
[show dhcp server binding <ip-address> on page 1720](#)  
[show dhcp server binding <session-id> on page 1720](#)  
[show dhcp server binding summary on page 1720](#)  
[show dhcp server binding <interfaces-vlan> on page 1720](#)  
[show dhcp server binding <interfaces-wildcard> on page 1720](#)

**Output Fields** [Table 149 on page 1717](#) lists the output fields for the **show dhcp server binding** command. Output fields are listed in the approximate order in which they appear.

**Table 149: show dhcp server binding Output Fields**

Field Name	Field Description	Level of Output
<i>number</i> clients, ( <i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	<b>summary</b>
IP address	IP address of the DHCP client.	<b>brief detail</b>
Session Id	Session ID of the subscriber session.	<b>brief detail</b>
Hardware address	Hardware address of the DHCP client.	<b>brief detail</b>
Expires	Number of seconds in which lease expires.	<b>brief detail</b>

Table 149: show dhcp server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	State of the address binding table on the extended DHCP local server: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has active IP address lease.</li> <li>• <b>FORCERENEW</b>—Client has received forcerenew message from server.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>RELEASE</b>—Client is releasing IP address lease.</li> <li>• <b>RENEWING</b>—Client sending request to renew IP address lease.</li> <li>• <b>REQUESTING</b>—Client requesting a DHCP server.</li> <li>• <b>SELECTING</b>—Client receiving offers from DHCP servers.</li> </ul>	<b>brief</b> <b>detail</b>
<b>Interface</b>	Interface on which the request was received.	<b>brief</b>
<b>Lease Expires</b>	Date and time at which the client's IP address lease expires.	<b>detail</b>
<b>Lease Expires in</b>	Number of seconds in which lease expires.	<b>detail</b>
<b>Lease Start</b>	Date and time at which the client's IP address lease started.	<b>detail</b>
<b>Lease time violated</b>	Lease time violation has occurred.	<b>detail</b>
<b>Last Packet Received</b>	Date and time at which the router received the last packet.	<b>detail</b>
<b>Incoming Client Interface</b>	Client's incoming interface.	<b>detail</b>
<b>Client Interface Svlan Id</b>	S-VLAN ID of the client's incoming interface.	<b>detail</b>
<b>Client Interface Vlan Id</b>	VLAN ID of the client's incoming interface.	<b>detail</b>
<b>Demux Interface</b>	Name of the IP demultiplexing (demux) interface.	<b>detail</b>
<b>Server IP Address or Server Identifier</b>	IP address of DHCP server.	<b>detail</b>
<b>Server Interface</b>	Interface of DHCP server.	<b>detail</b>
<b>Client Pool Name</b>	Name of address pool used to assign client IP address lease.	<b>detail</b>
<b>ACI Interface Set Name</b>	Internally generated name of the dynamic agent circuit identifier (ACI) interface set.	<b>detail</b>
<b>ACI Interface Set Index</b>	Index number of the dynamic ACI interface set.	<b>detail</b>
<b>ACI Interface Set Session ID</b>	Identifier of the dynamic ACI interface set entry in the session database.	<b>detail</b>

## Sample Output

### show dhcp server binding

```
user@host> show dhcp server binding
```

IP address	Session Id	Hardware address	Expires	State	Interface
100.20.20.15	6	00:10:94:00:00:01	86180	BOUND	ge-1/0/0.0
100.20.20.16	7	00:10:94:00:00:02	86180	BOUND	ge-1/0/0.0
100.20.20.17	8	00:10:94:00:00:03	86180	BOUND	ge-1/0/0.0
100.20.20.18	9	00:10:94:00:00:04	86180	BOUND	ge-1/0/0.0
100.20.20.19	10	00:10:94:00:00:05	86180	BOUND	ge-1/0/0.0

### show dhcp server binding detail

```
user@host> show dhcp server binding detail
Client IP Address: 100.20.20.15
  Hardware Address: 00:10:94:00:00:01
  State: BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

  Lease Expires: 2009-07-21 10:10:25 PDT
  Lease Expires in: 86151 seconds
  Lease Start: 2009-07-20 10:10:25 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address: 100.20.20.9
  Server Interface: none
  Session Id: 6
  Client Pool Name: 6
  Client IP Address: 100.20.20.16
  Hardware Address: 00:10:94:00:00:02
  State: BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

  Lease Expires: 2009-07-21 10:10:25 PDT
  Lease Expires in: 86151 seconds
  Lease Start: 2009-07-20 10:10:25 PDT
  Lease time violated: yes
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address: 100.20.20.9
  Server Interface: none
  Session Id: 7
  Client Pool Name: 7
```

### show dhcp server binding detail (ACI Interface Set Configured)

```
user@host> show dhcp server binding detail
Client IP Address: 100.20.22.14
  Hardware Address: 00:00:64:34:01:02
  State: BOUND(LOCAL_SERVER_STATE_BOUND)
  Lease Expires: 2012-03-13 09:53:32 PDT
  Lease Expires in: 82660 seconds
  Lease Start: 2012-03-12 10:23:32 PDT
  Last Packet Received: 2012-03-12 10:23:32 PDT
  Incoming Client Interface: demux0.1073741827
  Client Interface Svlan Id: 1802
  Client Interface Vlan Id: 302
  Demux Interface: demux0.1073741832
  Server Identifier: 100.20.200.202
  Session Id: 11
```

```
Client Pool Name:          poolA
Client Profile Name:       DEMUXprofile
ACI Interface Set Name:    aci-1002-demux0.1073741827
ACI Interface Set Index:   2
ACI Interface Set Session ID: 6
```

#### show dhcp server binding interface <vlan-id>

```
user@host> show dhcp server binding interface ge-1/1/0:100
IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.15    6          00:10:94:00:00:01 86124    BOUND  ge-1/1/0:100
```

#### show dhcp server binding interface <svlan-id>

```
user@host> show dhcp server binding interface ge-1/1/0:10-100
IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.16    7          00:10:94:00:00:02 86124    BOUND  ge-1/1/0:10-100
```

#### show dhcp server binding <ip-address>

```
user@host> show dhcp server binding 100.20.20.19
IP address      Session Id  Hardware address  Expires  State  Interface
100.20.20.19    10         00:10:94:00:00:05 86081    BOUND  ge-1/0/0.0
```

#### show dhcp server binding <session-id>

```
user@host> show dhcp server binding 6
IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.15    6          00:10:94:00:00:01 86124    BOUND  ge-1/0/0.0
```

#### show dhcp server binding summary

```
user@host> show dhcp server binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

#### show dhcp server binding <interfaces-vlan>

```
user@host> show dhcp server binding ge-1/0/0:100-200
IP address      Session Id  Hardware address  Expires  State  Interface
192.168.0.17    42         00:10:94:00:00:02 86346    BOUND  ge-1/0/0.1073741827
192.168.0.16    41         00:10:94:00:00:01 86346    BOUND  ge-1/0/0.1073741827
```

#### show dhcp server binding <interfaces-wildcard>

```
user@host> show dhcp server binding ge-1/3/*
IP address      Session Id  Hardware address  Expires  State  Interface
192.168.0.9     24         00:10:94:00:00:04 86361    BOUND  ge-1/3/0.110
192.168.0.8     23         00:10:94:00:00:03 86361    BOUND  ge-1/3/0.110
192.168.0.7     22         00:10:94:00:00:02 86361    BOUND  ge-1/3/0.110
```



## show dhcp server statistics

---

<b>Syntax</b>	<b>show dhcp server statistics</b> <b>&lt;logical-system <i>logical-system-name</i>&gt;</b> <b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0.
<b>Description</b>	Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(Optional) Display information about extended DHCP local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Display information about extended DHCP local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear dhcp server statistics on page 1702</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show dhcp server statistics on page 1722</a>
<b>Output Fields</b>	<a href="#">Table 150 on page 1722</a> lists the output fields for the <b>show dhcp server statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 150: show dhcp server statistics Output Fields

Field Name	Field Description
<b>Packets dropped</b>	<p>Number of packets discarded by the extended DHCP local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCP local server</li> <li>• <b>Authentication</b>—Number of packets discarded because they could not be authenticated</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified</li> <li>• <b>Dynamic profile</b>—Number of packets discarded due to dynamic profile information</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error</li> <li>• <b>Send error</b>—Number of packets that the extended DHCP local server could not send</li> </ul>
<b>Messages received</b>	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTP protocol data units (PDUs) received</li> <li>• <b>DHCPDECLINE</b>—Number of DHCP PDUs of type DECLINE received</li> <li>• <b>DHCPDISCOVER</b>—Number of DHCP PDUs of type DISCOVER received</li> <li>• <b>DHCPINFORM</b>—Number of DHCP PDUs of type INFORM received</li> <li>• <b>DHCPRELEASE</b>—Number of DHCP PDUs of type RELEASE received</li> <li>• <b>DHCPREQUEST</b>—Number of DHCP PDUs of type REQUEST received</li> </ul>
<b>Messages sent</b>	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREPLY</b>—Number of BOOTP PDUs transmitted</li> <li>• <b>DHCPOFFER</b>—Number of DHCP OFFER PDUs transmitted</li> <li>• <b>DHCPACK</b>—Number of DHCP ACK PDUs transmitted</li> <li>• <b>DHCPNACK</b>—Number of DHCP NACK PDUs transmitted</li> <li>• <b>DHCPFORCERENEW</b>—Number of DHCP FORCERENEW PDUs transmitted</li> </ul>

## Sample Output

### show dhcp server statistics

```

user@host> show dhcp server statistics
Packets dropped:
    Total                  1

```

Lease Time Violation	1
Messages received:	
BOOTREQUEST	25
DHCPDECLINE	0
DHCPDISCOVER	10
DHCPINFORM	0
DHCPRELEASE	4
DHCPREQUEST	10
Messages sent:	
BOOTREPLY	20
DHCPOFFER	10
DHCPACK	10
DHCPNAK	0
DHCPFORCERENEW	0

## show dhcpv6 server binding

---

Syntax	<pre>show dhcpv6 server binding &lt;address&gt; &lt;brief   detail   summary&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</pre>
Release Information	Command introduced in Junos OS Release 9.6. Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.
Description	Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server.
Options	<p><b>address</b>—(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:</p> <ul style="list-style-type: none"><li>• <i>CID</i>—The specified Client ID (CID).</li><li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li><li>• <i>session-id</i>—The specified session ID.</li></ul> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output about active client bindings. The default is <b>brief</b>, which produces the same output as <b>show dhcpv6 server binding</b>.</p> <p><b>interface interface-name</b>—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.</p> <p><b>interfaces-vlan</b>—(Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.</p> <p><b>interfaces-wildcard</b>—(Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Display information about active client bindings for DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Display information about active client bindings for DHCPv6 clients on the specified routing instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><li>• <i>Clearing DHCP Bindings for Subscriber Access</i></li><li>• <a href="#">clear dhcpv6 server binding on page 1704</a></li></ul>

**List of Sample Output**

- [show dhcpv6 server binding on page 1726](#)
- [show dhcpv6 server binding detail on page 1726](#)
- [show dhcpv6 server binding interface on page 1727](#)
- [show dhcpv6 server binding interface detail on page 1727](#)
- [show dhcpv6 server binding \(IPv6 Prefix\) on page 1728](#)
- [show dhcpv6 server binding \(Session ID\) on page 1728](#)
- [show dhcpv6 server binding \(Interfaces VLAN\) on page 1728](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 1728](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 1728](#)
- [show dhcpv6 server binding summary on page 1729](#)

**Output Fields** [Table 151 on page 1725](#) lists the output fields for the **show dhcpv6 server binding** command. Output fields are listed in the approximate order in which they appear.

**Table 151: show dhcpv6 server binding Output Fields**

Field Name	Field Description	Level of Output
<i>number clients</i> , ( <i>number init</i> , <i>number bound</i> , <i>number selecting</i> , <i>number requesting</i> , <i>number renewing</i> , <i>number releasing</i> )	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	<b>summary</b>
<b>Prefix</b>	Client's DHCPv6 prefix, or prefix used to support multiple address assignment.	<b>brief detail</b>
<b>Session Id</b>	Session ID of the subscriber session.	<b>brief detail</b>
<b>Expires</b>	Number of seconds in which lease expires.	<b>brief detail</b>
<b>State</b>	State of the address binding table on the extended DHCPv6 local server: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has active IP address lease.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>RECONFIGURE</b>—Server has sent reconfigure message to client.</li> <li>• <b>RELEASE</b>—Client is releasing IP address lease.</li> <li>• <b>RENEWING</b>—Client sending request to renew IP address lease.</li> <li>• <b>REQUESTING</b>—Client requesting a DHCPv6 server.</li> <li>• <b>SELECTING</b>—Client receiving offers from DHCPv6 servers.</li> </ul>	<b>brief detail</b>
<b>Interface</b>	Interface on which the DHCPv6 request was received.	<b>brief</b>
<b>Client IPv6 Address</b>	Client's IPv6 address.	<b>detail</b>
<b>Client IPv6 Prefix</b>	Client's IPv6 prefix.	<b>detail</b>
<b>Client DUID</b>	Client's DHCP Unique Identifier (DUID).	<b>brief detail</b>
<b>Lease expires</b>	Date and time at which the client's IP address lease expires.	<b>detail</b>

Table 151: show dhcpv6 server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Lease expires in	Number of seconds in which lease expires.	detail
Preferred Lease Expires	Date and UTC time at which the client's IPv6 prefix expires.	detail
Preferred Lease Expires in	Number of seconds at which client's IPv6 prefix expires.	detail
Lease Start	Date and time at which the client's address lease was obtained.	detail
Lease time violated	Lease time violation has occurred.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server IP Address	IP address of DHCPv6 server.	detail
Server Interface	Interface of DHCPv6 server.	detail
Client Pool Name	Address pool used to assign IPv6 address.	detail
Client Prefix Pool Name	Address pool used to assign IPv6 prefix.	detail
Client Id length	Length of the DHCPv6 client ID, in bytes.	detail
Client Id	ID of the DHCPv6 client.	detail

## Sample Output

### show dhcpv6 server binding

```

user@host> show dhcpv6 server binding
Prefix          Session Id Expires State Interface Client DUID
2001:bd8:1111:2222::/64 6 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:01
2001:bd8:1111:2222::/64 7 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
2001:bd8:1111:2222::/64 8 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
2001:bd8:1111:2222::/64 9 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:04
2001:bd8:1111:2222::/64 10 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:05
2002::1/74 11 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:06

```

### show dhcpv6 server binding detail

```

user@host> show dhcpv6 server binding detail

```

```

Session Id: 6
  Client IPv6 Prefix:      2001:bd8:1111:2222::/64
  Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:01

  State:
  BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
    Lease Expires:         2009-07-21 10:41:15 PDT
    Lease Expires in:      86308 seconds
    Preferred Lease Expires: 2012-07-24 00:18:14 UTC
    Preferred Lease Expires in: 600 seconds
    Lease Start:           2009-07-20 10:41:15 PDT
    Lease time violated:    yes
    Incoming Client Interface: ge-1/0/0.0
    Server Ip Address:      0.0.0.0
    Server Interface:       none
    Client Id Length:       14
    Client Id:
    /0x00010001/0x02e159c0/0x00109400/0x0001

```

```

Session Id: 7
  Client IPv6 Address:     2002::1/128
  Client IPv6 Prefix:      2001:bd8:1111:2222::/64
  Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

  State:
  BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
    Lease Expires:         2009-07-21 10:41:15 PDT
    Lease Expires in:      86308 seconds
    Preferred Lease Expires: 2012-07-24 00:18:14 UTC
    Preferred Lease Expires in: 600 seconds
    Lease Start:           2009-07-20 10:41:15 PDT
    Incoming Client Interface: ge-1/0/0.0
    Server Ip Address:      0.0.0.0
    Client Pool Name:       bos-v6-pool
    Client Prefix Pool Name: bos-v6-prefix-pool
    Client Id Length:       14
    Client Id:
    /0x00010001/0x02e159c0/0x00109400/0x0002

```

### show dhcpv6 server binding interface

```

user@host> show dhcpv6 server binding interface ge-1/0/0:10-101
Prefix          Session Id Expires State Interface Client DUID
2001:bd8:1111:2222::/64 1      86055   BOUND   ge-1/0/0.100
LL_TIME0x1-0x4b0a53b9-00:10:94:00:00:01

```

### show dhcpv6 server binding interface detail

```

user@host> show dhcpv6 server binding interface ge-1/0/0:10-101 detail
Session Id: 7
  Client IPv6 Prefix:      2001:bd8:1111:2222::/64
  Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

  State:                   BOUND(bound)
  Lease Expires:           2009-07-21 10:41:15 PDT
  Lease Expires in:        86136 seconds
  Preferred Lease Expires: 2012-07-24 00:18:14 UTC
  Preferred Lease Expires in: 600 seconds
  Lease Start:             2009-07-20 10:41:15 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address:       0.0.0.0

```

```

Server Interface:          none
Client Id Length:         14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

### show dhcpv6 server binding (IPv6 Prefix)

```

user@host> show dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
detail
Session Id: 7
Client IPv6 Prefix:      2001:bd8:1111:2222::/64
Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

State:                   BOUND(bound)
Lease Expires:           2009-07-21 10:41:15 PDT
Lease Expires in:        86136 seconds
Preferred Lease Expires: 2012-07-24 00:18:14 UTC
Preferred Lease Expires in: 600 seconds
Lease Start:             2009-07-20 10:41:15 PDT
Incoming Client Interface: ge-1/0/0.0
Server Ip Address:        0.0.0.0
Server Interface:        none
Client Id Length:        14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

### show dhcpv6 server binding (Session ID)

```

user@host> show dhcpv6 server binding 8
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 8      86235 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03

```

### show dhcpv6 server binding (Interfaces VLAN)

```

user@host> show dhcpv6 server binding ge-1/0/0:100-200
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 11      87583 BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 12      87583 BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

### show dhcpv6 server binding (Interfaces Wildcard)

```

user@host> show dhcpv6 server binding demux0
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 30      79681 BOUND demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 31      79681 BOUND demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32 32      79681 BOUND demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

### show dhcpv6 server binding (Interfaces Wildcard)

```

user@host> show dhcpv6 server binding ge-1/3/*
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 22      79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 33      79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```



```
2001:CB9::/32      24      79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

#### show dhcpv6 server binding summary

```
user@host> show dhcpv6 server binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

## show dhcpv6 server statistics

---

Syntax	<b>show dhcpv6 server statistics</b> <b>&lt;logical-system <i>logical-system-name</i>&gt;</b> <b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>
Release Information	Command introduced in Junos OS Release 9.6.
Description	Display extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.
Options	<b>logical-system <i>logical-system-name</i></b> —(Optional) Display information about extended DHCPv6 local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.  <b>routing-instance <i>routing-instance-name</i></b> —(Optional) Display information about extended DHCPv6 local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">clear dhcpv6 server statistics on page 1706</a></li></ul>
List of Sample Output	<a href="#">show dhcpv6 server statistics on page 1731</a>
Output Fields	<a href="#">Table 152 on page 1731</a> lists the output fields for the <b>show dhcpv6 server statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 152: show dhcpv6 server statistics Output Fields

Field Name	Field Description
<b>Packets dropped</b>	<p>Number of packets discarded by the extended DHCPv6 local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCPv6 local server</li> <li>• <b>Strict Reconfigure</b>—Number of solicit messages discarded because the client does not support reconfiguration</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error</li> <li>• <b>Send error</b>—Number of packets that the extended DHCPv6 local server could not send</li> </ul>
<b>Messages received</b>	<p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> <li>• <b>DHCPV6_CONFIRM</b>—Number of DHCPv6 CONFIRM PDUs received.</li> <li>• <b>DHCPV6_DECLINE</b>—Number of DHCPv6 DECLINE PDUs received.</li> <li>• <b>DHCPV6_INFORMATION_REQUEST</b>—Number of DHCPv6 INFORMATION-REQUEST PDUs received.</li> <li>• <b>DHCPV6_REBIND</b>—Number of DHCPv6 REBIND PDUs received.</li> <li>• <b>DHCPV6_RELAY_FORW</b>—Number of DHCPv6 RELAY-FORW PDUs received.</li> <li>• <b>DHCPV6_RELAY_REPL</b>—Number of DHCPv6 RELAY-REPL PDUs received.</li> <li>• <b>DHCPV6_RELEASE</b>—Number of DHCPv6 RELEASE PDUs received.</li> <li>• <b>DHCPV6_RENEW</b>—Number of DHCPv6 RENEW PDUs received.</li> <li>• <b>DHCPV6_REQUEST</b>—Number of DHCPv6 REQUEST PDUs received.</li> <li>• <b>DHCPV6_SOLICIT</b>—Number of DHCPv6 SOLICIT PDUs received.</li> </ul>
<b>Messages sent</b>	<p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> <li>• <b>DHCPV6_ADVERTISE</b>—Number of DHCPv6 ADVERTISE PDUs transmitted.</li> <li>• <b>DHCPV6_REPLY</b>—Number of DHCPv6 ADVERTISE PDUs transmitted.</li> <li>• <b>DHC6_RECONFIGURE</b>—Number of DHCPv6 RECONFIGURE PDUs transmitted.</li> </ul>

## Sample Output

### show dhcpv6 server statistics

```
user@host> show dhcpv6 server statistics
```

```
Dhcpv6 Packets dropped:
  Total          1
  Lease Time Violation 1

Messages received:
  DHCPV6_DECLINE      0
  DHCPV6_SOLICIT      9
  DHCPV6_INFORMATION_REQUEST 0
  DHCPV6_RELEASE      0
  DHCPV6_REQUEST      5
  DHCPV6_CONFIRM      0
  DHCPV6_RENEW        0
  DHCPV6_REBIND       0
  DHCPV6_RELAY_FORW   0
  DHCPV6_RELAY_REPL   0

Messages sent:
  DHCPV6_ADVERTISE    9
  DHCPV6_REPLY        5
  DHCPV6_RECONFIGURE  0
```

---

## DHCP Relay Agent Monitoring Commands

- `clear dhcp relay binding`
- `clear dhcp relay statistics`
- `clear dhcpv6 relay binding`
- `clear dhcpv6 relay statistics`
- `show dhcp relay binding`
- `show dhcp relay statistics`
- `show dhcpv6 relay binding`
- `show dhcpv6 relay statistics`
- `show route extensive`
- `show route protocol`

## clear dhcp relay binding

<b>Syntax</b>	<pre>clear dhcp relay binding &lt;address&gt; &lt;all&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.3.</p> <p>Options <b>all</b> and <b>interface</b> added in Junos OS Release 8.4.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.</p>
<b>Description</b>	Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table.
<b>Options</b>	<p><b>address</b>—(Optional) Clear the binding state for the DHCP client, using one of the following entries:</p> <ul style="list-style-type: none"> <li><i>ip-address</i>—The specified IP address.</li> <li><i>mac-address</i>—The specified MAC address.</li> <li><i>session-id</i>—The specified session ID.</li> </ul> <p><b>all</b>—(Optional) Clear the binding state for all DHCP clients.</p> <p><b>interface interface-name</b>—(Optional) Clear the binding state for DHCP clients on the specified interface.</p> <p><b>interfaces-vlan</b>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><b>interfaces-wildcard</b>—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Clear the binding state for DHCP clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Clear the binding state for DHCP clients on the specified routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Clearing DHCP Bindings for Subscriber Access</i></li> <li><a href="#">show dhcp relay binding on page 1743</a></li> </ul>

**List of Sample Output** [clear dhcp relay binding on page 1734](#)  
[clear dhcp relay binding all on page 1734](#)  
[clear dhcp relay binding interface on page 1734](#)  
[clear dhcp relay binding <interfaces-vlan> on page 1734](#)  
[clear dhcp relay binding <interfaces-wildcard> on page 1734](#)

**Output Fields** See [show dhcp relay binding](#) for an explanation of output fields.

## Sample Output

### clear dhcp relay binding

The following sample output displays the address bindings in the DHCP client table before and after the **clear dhcp relay binding** command is issued.

```
user@host> show dhcp relay binding
IP address      Hardware address  Type    Lease expires at
100.20.32.1     90:00:00:01:00:01 active    2007-02-08 16:41:17 EST
192.168.14.8    90:00:01:01:02:01 active    2007-02-10 10:01:06 EST
```

```
user@host> clear dhcp relay binding 100.20.32.1
```

```
user@host> show dhcp relay binding
IP address      Hardware address  Type    Lease expires at
192.168.14.8    90:00:01:01:02:01 active    2007-02-10 10:01:06 EST
```

### clear dhcp relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcp relay binding all
```

### clear dhcp relay binding interface

The following command clears DHCP relay agent bindings on a specific interface:

```
user@host> clear dhcp relay binding interface fe-0/0/3
```

### clear dhcp relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP relay agent bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp relay binding interface ae0
```

### clear dhcp relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP relay agent bindings over a specific interface:

```
user@host> clear dhcp relay binding ge-1/0/0.*
```

## clear dhcp relay statistics

---

<b>Syntax</b>	<pre>clear dhcp relay statistics &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>
<b>Syntax</b>	<p>Syntax for EX Series switches:</p> <pre>show dhcp relay statistics &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.</p>
<b>Description</b>	Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show dhcp relay statistics on page 1748</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear dhcp relay statistics on page 1736</a>
<b>Output Fields</b>	<a href="#">Table 153 on page 1736</a> lists the output fields for the <b>clear dhcp relay statistics</b> command.

Table 153: clear dhcp relay statistics Output Fields

Field Name	Field Description
<b>Packets dropped</b>	<p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the <b>Packets dropped</b> output. When all of the Packets dropped statistics are 0 (zero), only the <b>Total</b> field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCP relay agent application.</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified.</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified.</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified.</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified.</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment.</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface.</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance.</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address.</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short.</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error.</li> <li>• <b>Send error</b>—Number of packets that the extended DHCP relay application could not send.</li> <li>• <b>Option 60</b>—Number of packets discarded containing DHCP option 60 vendor-specific information.</li> <li>• <b>Option 82</b>—Number of packets discarded because DHCP option 82 information could not be added.</li> </ul>
<b>Messages received</b>	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTP protocol data units (PDUs) received</li> <li>• <b>DHCPDECLINE</b>—Number of DHCP PDUs of type DECLINE received</li> <li>• <b>DHCPDISCOVER</b>—Number of DHCP PDUs of type DISCOVER received</li> <li>• <b>DHCPINFORM</b>—Number of DHCP PDUs of type INFORM received</li> <li>• <b>DHCPRELEASE</b>—Number of DHCP PDUs of type RELEASE received</li> <li>• <b>DHCPREQUEST</b>—Number of DHCP PDUs of type REQUEST received</li> </ul>
<b>Messages sent</b>	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREPLY</b>—Number of BOOTP PDUs transmitted</li> <li>• <b>DHCPOFFER</b>—Number of DHCP OFFER PDUs transmitted</li> <li>• <b>DHCPACK</b>—Number of DHCP ACK PDUs transmitted</li> <li>• <b>DHC PNACK</b>—Number of DHCP NACK PDUs transmitted</li> </ul>

## Sample Output

### clear dhcp relay statistics

The following sample output displays the DHCP relay statistics before and after the **clear dhcp relay statistics** command is issued.

```
user@host> show dhcp relay statistics
```



```
Packets dropped:
  Total          1
  Lease Time Violated 1

Messages received:
  BOOTREQUEST    116
  DHCPDECLINE    0
  DHCPDISCOVER   11
  DHCPINFORM     0
  DHCPRELEASE    0
  DHCPREQUEST    105

Messages sent:
  BOOTREPLY      44
  DHCPOFFER      11
  DHCPACK        11
  DHCPNAK        11
```

```
user@host> clear dhcp relay statistics
```

```
user@host> show dhcp relay statistics
```

```
Packets dropped:
  Total          0

Messages received:
  BOOTREQUEST    0
  DHCPDECLINE    0
  DHCPDISCOVER   0
  DHCPINFORM     0
  DHCPRELEASE    0
  DHCPREQUEST    0

Messages sent:
  BOOTREPLY      0
  DHCPOFFER      0
  DHCPACK        0
  DHCPNAK        0
```

## clear dhcpv6 relay binding

---

<b>Syntax</b>	<code>clear dhcpv6 relay binding</code> <code>&lt;address&gt;</code> <code>&lt;all&gt;</code> <code>&lt;interface <i>interface-name</i>&gt;</code> <code>&lt;interfaces-vlan&gt;</code> <code>&lt;interfaces-wildcard&gt;</code> <code>&lt;logical-system <i>logical-system-name</i>&gt;</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 12.3R2 for EX Series switches. Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.
<b>Description</b>	Clear the binding state of Dynamic Host Configuration Protocol for IPv6 (DHCPv6) clients from the client table.
<b>Options</b>	<p><b><i>address</i></b>—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"><li>• <i>CID</i>—The specified Client ID (CID).</li><li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li><li>• <i>session-id</i>—The specified session ID.</li></ul> <p><b><i>all</i></b>—(Optional) Clear the binding state for all DHCPv6 clients.</p> <p><b><i>interfaces-vlan</i></b>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><b><i>interfaces-wildcard</i></b>—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p><b><i>interface interface-name</i></b>—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.</p> <p><b><i>logical-system logical-system-name</i></b>—(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.</p> <p><b><i>routing-instance routing-instance-name</i></b>—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Clearing DHCP Bindings for Subscriber Access</i></li><li>• <a href="#">show dhcpv6 relay binding on page 1751</a></li></ul>

- List of Sample Output** [clear dhcpv6 relay binding on page 1739](#)  
[clear dhcpv6 relay binding <prefix> on page 1739](#)  
[clear dhcpv6 relay binding all on page 1739](#)  
[clear dhcv6p relay binding interface on page 1739](#)  
[clear dhcpv6 relay binding <interfaces-vlan> on page 1740](#)  
[clear dhcpv6 relay binding <interfaces-wildcard> on page 1740](#)
- Output Fields** See [show dhcpv6 relay binding](#) for an explanation of output fields.

## Sample Output

### clear dhcpv6 relay binding

The following sample output displays the DHCPv6 bindings before and after the **clear dhcpv6 relay binding** command is issued.

```
user@host> show dhcpv6 relay binding
```

Prefix	Session Id	Expires	State	Interface	Client DUID
2001:bd8:3c4d:15::/64	1	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01					
2001:bd8:3c4d:16::/64	2	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02					
2001:bd8:3c4d:17::/64	3	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03					
2001:bd8:3c4d:18::/64	4	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04					
2001:bd8:3c4d:19::/64	5	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05					
2001:bd8:3c4d:20::/64	6	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06					

### clear dhcpv6 relay binding <prefix>

```
user@host> clear dhcpv6 relay binding 2001:bd8:3c4d:15::/64
```

```
user@host> show dhcpv6 relay binding
```

Prefix	Session Id	Expires	State	Interface	Client DUID
2001:bd8:3c4d:16::/64	2	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02					
2001:bd8:3c4d:17::/64	3	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03					
2001:bd8:3c4d:18::/64	4	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04					
2001:bd8:3c4d:19::/64	5	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05					
2001:bd8:3c4d:20::/64	6	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06					

### clear dhcpv6 relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcpv6 relay binding all
```

### clear dhcv6p relay binding interface

The following command clears DHCPv6 relay agent bindings on a specific interface:

```
user@host> clear dhcpv6 relay binding interface fe-0/0/2
```

#### clear dhcpv6 relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 relay agent bindings on top of the underlying interface **ae0**, which clears DHCPv6 bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcpv6 relay binding interface ae0
```

#### clear dhcpv6 relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 relay agent bindings over a specific interface:

```
user@host> clear dhcpv6 relay binding ge-1/0/0.*
```

## clear dhcpv6 relay statistics

<b>Syntax</b>	<b>clear dhcpv6 relay statistics</b> <b>&lt;logical-system <i>logical-system-name</i>&gt;</b> <b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.
<b>Description</b>	Clear all Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.
<b>Options</b>	<b>logical-system <i>logical-system-name</i></b> —(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.  <b>routing-instance <i>routing-instance-name</i></b> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">clear dhcpv6 relay statistics on page 1741</a>
<b>Output Fields</b>	See <a href="#">show dhcpv6 relay statistics</a> for an explanation of output fields.

## Sample Output

### clear dhcpv6 relay statistics

The following sample output displays the DHCPv6 relay statistics before and after the **clear dhcpv6 relay statistics** command is issued.

```

user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                0
    Lease Time Violated  1

Messages received:
    DHCPV6_DECLINE        0
    DHCPV6_SOLICIT        10
    DHCPV6_INFORMATION_REQUEST  0
    DHCPV6_RELEASE        0
    DHCPV6_REQUEST        10
    DHCPV6_CONFIRM        0
    DHCPV6_RENEW          0
    DHCPV6_REBIND         0
    DHCPV6_RELAY_REPL     0

Messages sent:
    DHCPV6_ADVERTISE      0
    DHCPV6_REPLY           0
    DHCPV6_RECONFIGURE    0
    DHCPV6_RELAY_FORW     0

```

```
user@host> clear dhcpv6 relay statistics
```

```
user@host> show dhcpv6 relay statistics
```

```
DHCPv6 Packets dropped:
```

```
    Total                                0
```

```
Messages received:
```

```
    DHCPV6_DECLINE                        0
```

```
    DHCPV6_SOLICIT                        0
```

```
    DHCPV6_INFORMATION_REQUEST           0
```

```
    DHCPV6_RELEASE                        0
```

```
    DHCPV6_REQUEST                        0
```

```
    DHCPV6_CONFIRM                        0
```

```
    DHCPV6_RENEW                          0
```

```
    DHCPV6_REBIND                         0
```

```
    DHCPV6_RELAY_REPL                     0
```

```
Messages sent:
```

```
    DHCPV6_ADVERTISE                      0
```

```
    DHCPV6_REPLY                          0
```

```
    DHCPV6_RECONFIGURE                    0
```

```
    DHCPV6_RELAY_FORW                     0
```

## show dhcp relay binding

<b>Syntax</b>	<pre> show dhcp relay binding &lt;address&gt; &lt;brief&gt; &lt;detail&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;ip-address   mac-address&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt; &lt;summary&gt; </pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.3.</p> <p>Options <b>interface</b> and <b>mac-address</b> added in Junos OS Release 8.4.</p> <p>Options <b>interfaces-vlan</b> and <b>interfaces-wildcard</b> added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.</p>
<b>Description</b>	Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.
<b>Options</b>	<p><b>address</b>—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:</p> <ul style="list-style-type: none"> <li>• <b>ip-address</b>—The specified IP address.</li> <li>• <b>mac-address</b>—The specified MAC address.</li> <li>• <b>session-id</b>—The specified session ID.</li> </ul> <p><b>brief</b>—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as <b>show dhcp relay binding</b>.</p> <p><b>detail</b>—(Optional) Display detailed client binding information.</p> <p><b>interface interface-name</b>—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.</p> <p><b>interfaces-vlan</b>—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.</p> <p><b>interfaces-wildcard</b>—(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Perform this operation on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Perform this operation on the specified routing instance.</p> <p><b>summary</b>—(Optional) Display a summary of DHCP client information.</p>

**Required Privilege Level** view

**Related Documentation**

- *Clearing DHCP Bindings for Subscriber Access*
- [clear dhcp relay binding on page 1733](#)

**List of Sample Output**

- [show dhcp relay binding on page 1745](#)
- [show dhcp relay binding detail on page 1746](#)
- [show dhcp relay binding interface on page 1746](#)
- [show dhcp relay binding interface vlan-id on page 1746](#)
- [show dhcp relay binding interface svlan-id on page 1746](#)
- [show dhcp relay binding ip-address on page 1747](#)
- [show dhcp relay binding mac-address on page 1747](#)
- [show dhcp relay binding session-id on page 1747](#)
- [show dhcp relay binding <interfaces-vlan> on page 1747](#)
- [show dhcp relay binding <interfaces-wildcard> on page 1747](#)
- [show dhcp relay binding summary on page 1747](#)

**Output Fields** Table 154 on page 1744 lists the output fields for the **show dhcp relay binding** command. Output fields are listed in the approximate order in which they appear.

**Table 154: show dhcp relay binding Output Fields**

Field Name	Field Description	Level of Output
<i>number</i> clients, ( <i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> rebinding, <i>number</i> releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	<b>summary</b>
<b>IP address</b>	IP address of the DHCP client.	<b>briefdetail</b>
<b>Session Id</b>	Session ID of the subscriber session.	<b>briefdetail</b>
<b>Generated Remote ID</b>	Remote ID generated by the Option 82 Agent Remote ID (suboption 1)	<b>detail</b>
<b>Hardware address</b>	Hardware address of the DHCP client.	<b>briefdetail</b>
<b>Expires</b>	Number of seconds in which the lease expires.	<b>briefdetail</b>



Table 154: show dhcp relay binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	State of the DHCP relay address binding table on the DHCP client: <ul style="list-style-type: none"> <li><b>BOUND</b>—Client has an active IP address lease.</li> <li><b>INIT</b>—Initial state.</li> <li><b>REBINDING</b>—Client is broadcasting a request to renew the IP address lease.</li> <li><b>RELEASE</b>—Client is releasing the IP address lease.</li> <li><b>RENEWING</b>—Client is sending a request to renew the IP address lease.</li> <li><b>REQUESTING</b>—Client is requesting a DHCP server.</li> <li><b>SELECTING</b>—Client is receiving offers from DHCP servers.</li> </ul>	<b>briefdetail</b>
<b>Interface</b>	Incoming client interface.	brief
<b>Lease Expires</b>	Date and time at which the client's IP address lease expires.	detail
<b>Lease Expires in</b>	Number of seconds in which the lease expires.	detail
<b>Lease Start</b>	Date and time at which the client's IP address lease started.	detail
<b>Lease time violated</b>	Lease time violation has occurred.	detail
<b>Incoming Client Interface</b>	Client's incoming interface.	detail
<b>Server IP Address</b>	IP address of the DHCP server.	detail
<b>Server Interface</b>	Interface of the DHCP server.	detail
<b>Bootp Relay Address</b>	IP address of BOOTP relay.	detail
<b>Type</b>	Type of DHCP packet processing performed on the router: <ul style="list-style-type: none"> <li><b>active</b>—Router actively processes and relays DHCP packets.</li> <li><b>passive</b>—Router passively snoops DHCP packets passing through the router.</li> </ul>	All levels
<b>Lease expires at</b>	Date and time at which the client's IP address lease expires.	All levels

## Sample Output

### show dhcp relay binding

```

user@host> show dhcp relay binding
IP address      Session Id  Hardware address  Expires   State   Interface
100.20.32.11    41         00:10:94:00:00:01 86371     BOUND   ge-1/0/0.0
100.20.32.12    42         00:10:94:00:00:02 86371     BOUND   ge-1/0/0.0

```

100.20.32.13	43	00:10:94:00:00:03	86371	BOUND	ge-1/0/0.0
100.20.32.14	44	00:10:94:00:00:04	86371	BOUND	ge-1/0/0.0
100.20.32.15	45	00:10:94:00:00:05	86371	BOUND	ge-1/0/0.0

**show dhcp relay binding detail**

```
user@host> show dhcp relay binding detail
```

```
Client IP Address: 100.20.32.11
  Hardware Address: 00:10:94:00:00:01
  State: BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
  Lease Expires: 2009-07-21 11:00:06 PDT
  Lease Expires in: 86361 seconds
  Lease Start: 2009-07-20 11:00:06 PDT
  Lease time violated: yes
  Last Packet Received: 2009-07-20 11:00:06 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address: 100.20.22.2
  Server Interface: none
  Bootp Relay Address: 100.20.32.2
  Session Id: 41

Client IP Address: 100.20.32.12
  Hardware Address: 00:10:94:00:00:02
  State: BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
  Lease Expires: 2009-07-21 11:00:06 PDT
  Lease Expires in: 86361 seconds
  Lease Start: 2009-07-20 11:00:06 PDT
  Last Packet Received: 2009-07-20 11:00:06 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address: 100.20.22.2
  Server Interface: none
  Bootp Relay Address: 100.20.32.2
  Session Id: 42
  Generated Remote ID: host:ge-1/0/0:100
```

**show dhcp relay binding interface**

```
user@host> show dhcp relay binding interface fe-0/0/2
```

IP address	Hardware address	Type	Lease expires at
100.20.32.1	90:00:00:01:00:01	active	2007-03-27 15:06:20 EDT

**show dhcp relay binding interface vlan-id**

```
user@host> show dhcp relay binding interface ge-1/1/0:100
```

IP address	Session Id	Hardware address	Expires	State	Interface
200.20.20.15	6	00:10:94:00:00:01	86124	BOUND	ge-1/1/0:100

**show dhcp relay binding interface svlan-id**

```
user@host> show dhcp relay binding interface ge-1/1/0:10-100
```

IP address	Session Id	Hardware address	Expires	State	Interface
------------	------------	------------------	---------	-------	-----------

```

200.20.20.16      7          00:10:94:00:00:02  86124      BOUND
ge-1/1/0:10-100

```

#### show dhcp relay binding ip-address

```

user@host> show dhcp relay binding 100.20.32.13
IP address      Session Id  Hardware address  Expires    State      Interface
100.20.32.13    43         00:10:94:00:00:03  86293     BOUND     ge-1/0/0.0

```

#### show dhcp relay binding mac-address

```

user@host> show dhcp relay binding 00:10:94:00:00:05
IP address      Session Id  Hardware address  Expires    State      Interface
100.20.32.15    45         00:10:94:00:00:05  86279     BOUND     ge-1/0/0.0

```

#### show dhcp relay binding session-id

```

user@host> show dhcp relay binding 41
IP address      Session Id  Hardware address  Expires    State      Interface
100.20.32.11    41         00:10:94:00:00:01  86305     BOUND     ge-1/0/0.0

```

#### show dhcp relay binding <interfaces-vlan>

```

user@host> show dhcp relay binding ge-1/0/0:100-200
IP address      Session Id  Hardware address  Expires    State      Interface
192.168.0.17    42         00:10:94:00:00:02  86346     BOUND     ge-1/0/0.1073741827
192.168.0.16    41         00:10:94:00:00:01  86346     BOUND     ge-1/0/0.1073741827

```

#### show dhcp relay binding <interfaces-wildcard>

```

user@host> show dhcp relay binding ge-1/3/*
IP address      Session Id  Hardware address  Expires    State      Interface
192.168.0.9     24         00:10:94:00:00:04  86361     BOUND     ge-1/3/0.110
192.168.0.8     23         00:10:94:00:00:03  86361     BOUND     ge-1/3/0.110
192.168.0.7     22         00:10:94:00:00:02  86361     BOUND     ge-1/3/0.110

```

#### show dhcp relay binding summary

```

user@host> show dhcp relay binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 rebinding,
0 releasing)

```

## show dhcp relay statistics

---

Syntax	<code>show dhcp relay statistics</code> <code>&lt;logical-system <i>logical-system-name</i>&gt;</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
Syntax	Syntax for EX Series switches:  <code>show dhcp relay statistics</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
Release Information	Command introduced in Junos OS Release 8.3. Command introduced in Junos OS Release 12.1 for EX Series switches. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.
Description	Display Dynamic Host Configuration Protocol (DHCP) relay statistics.
Options	<code>logical-system <i>logical-system-name</i></code> —(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.  <code>routing-instance <i>routing-instance-name</i></code> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">clear dhcp relay statistics on page 1735</a></li></ul>
List of Sample Output	<a href="#">show dhcp relay statistics on page 1750</a>
Output Fields	<a href="#">Table 155 on page 1749</a> lists the output fields for the <code>show dhcp relay statistics</code> command. Output fields are listed in the approximate order in which they appear.

Table 155: show dhcp relay statistics Output Fields

Field Name	Field Description
<b>Packets dropped</b>	<p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the <b>Packets dropped</b> output. When all of the Packets dropped statistics are 0 (zero), only the <b>Total</b> field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCP relay agent application.</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified.</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified.</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified.</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified.</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment.</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface.</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance.</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address.</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short.</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error.</li> <li>• <b>Send error</b>—Number of packets that the extended DHCP relay application could not send.</li> <li>• <b>Option 60</b>—Number of packets discarded containing DHCP option 60 vendor-specific information.</li> <li>• <b>Option 82</b>—Number of packets discarded because DHCP option 82 information could not be added.</li> </ul>
<b>Messages received</b>	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTP protocol data units (PDUs) received</li> <li>• <b>DHCPDECLINE</b>—Number of DHCP PDUs of type DECLINE received</li> <li>• <b>DHCPDISCOVER</b>—Number of DHCP PDUs of type DISCOVER received</li> <li>• <b>DHCPINFORM</b>—Number of DHCP PDUs of type INFORM received</li> <li>• <b>DHCPRELEASE</b>—Number of DHCP PDUs of type RELEASE received</li> <li>• <b>DHCPREQUEST</b>—Number of DHCP PDUs of type REQUEST received</li> </ul>
<b>Messages sent</b>	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREPLY</b>—Number of BOOTP PDUs transmitted</li> <li>• <b>DHCPOFFER</b>—Number of DHCP OFFER PDUs transmitted</li> <li>• <b>DHCPACK</b>—Number of DHCP ACK PDUs transmitted</li> <li>• <b>DHCPNACK</b>—Number of DHCP NACK PDUs transmitted</li> <li>• <b>DHCPFORCERENEW</b>—Number of DHCP FORCERENEW PDUs transmitted</li> </ul>
<b>External Server Response</b>	State of the external DHCP server responsiveness.
<b>Packets forwarded</b>	<p>Number of packets forwarded.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTREQUEST protocol data units (PDUs) forwarded</li> <li>• <b>BOOTREPLY</b>—Number of BOOTREPLY protocol data units (PDUs) forwarded</li> </ul>

Table 155: show dhcp relay statistics Output Fields (*continued*)

Field Name	Field Description
External Server Response	State of the external DHCP server responsiveness.

## Sample Output

### show dhcp relay statistics

```

user@host> show dhcp relay statistics
Packets dropped:
    Total                34
    Bad hardware address  1
    Bad opcode            1
    Bad options           3
    Invalid server address 5
    Lease Time Violation  1
    No available addresses 1
    No interface match    2
    No routing instance match 9
    No valid local address 4
    Packet too short      2
    Read error            1
    Send error            1
    Option 60             1
    Option 82             2

Messages received:
    BOOTREQUEST          116
    DHCPDECLINE           0
    DHCPDISCOVER          11
    DHCPINFORM            0
    DHCPRELEASE           0
    DHCPREQUEST          105

Messages sent:
    BOOTREPLY             0
    DHCPOFFER             2
    DHCPACK               1
    DHCPNAK               0
    DHCPFORCERENEW        0

Packets forwarded:
    Total                4
    BOOTREQUEST          2
    BOOTREPLY            2

External Server Response:
    State                Responding

```

## show dhcpv6 relay binding

<b>Syntax</b>	<pre>show dhcpv6 relay binding &lt;address&gt; &lt;brief&gt; &lt;detail&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt; &lt;summary&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 11.4.</p> <p><i>interfaces-vlan</i> and <i>interfaces-wildcard</i> options introduced in Junos OS Release 12.1.</p>
<b>Description</b>	Display the DHCPv6 address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.
<b>Options</b>	<p><b>address</b>—(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:</p> <ul style="list-style-type: none"> <li>• <i>CID</i>—The specified Client ID (CID).</li> <li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li> <li>• <i>session-id</i>—The specified session ID.</li> </ul> <p><b>brief</b>—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as <b>show dhcpv6 relay binding</b>.</p> <p><b>detail</b>—(Optional) Display detailed client binding information.</p> <p><b>interface interface-name</b>—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and S-VLAN ID.</p> <p><b>interfaces-vlan</b>—(Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.</p> <p><b>interfaces-wildcard</b>—(Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Perform this operation on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Perform this operation on the specified routing instance.</p> <p><b>summary</b>—(Optional) Display a summary of DHCPv6 client information.</p>
<b>Required Privilege Level</b>	view

- Related Documentation**
- [Clearing DHCP Bindings for Subscriber Access](#)
  - [clear dhcpv6 relay binding on page 1738](#)

- List of Sample Output**
- [show dhcpv6 relay binding on page 1753](#)
  - [show dhcpv6 relay binding \(Address\) on page 1754](#)
  - [show dhcpv6 relay binding detail \(Client ID\) on page 1754](#)
  - [show dhcpv6 relay binding detail on page 1754](#)
  - [show dhcpv6 relay binding detail \(Multi-Relay Topology\) on page 1755](#)
  - [show dhcpv6 relay binding \(Session ID\) on page 1755](#)
  - [show dhcpv6 relay binding \(Interfaces VLAN\) on page 1755](#)
  - [show dhcpv6 relay binding \(Interfaces Wildcard\) on page 1755](#)
  - [show dhcpv6 relay binding \(Interfaces Wildcard\) on page 1756](#)
  - [show dhcpv6 relay binding summary on page 1756](#)

- Output Fields** [Table 156 on page 1752](#) lists the output fields for the **show dhcpv6 relay binding** command. Output fields are listed in the approximate order in which they appear.

Table 156: show dhcpv6 relay binding Output Fields

Field Name	Field Description	Level of Output
<i>number clients, (number init, number bound, number selecting, number requesting, number renewing, number rebinding, number releasing)</i>	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	<b>summary</b>
<b>Client IPv6 Prefix</b>	Prefix of the DHCPv6 client.	<b>brief detail</b>
<b>Client DUID</b>	DHCP for IPv6 Unique Identifier (DUID) of the client.	<b>brief detail</b>
<b>Session Id</b>	Session ID of the subscriber session.	<b>brief detail</b>
<b>Expires</b>	Number of seconds in which the lease expires.	<b>brief detail</b>
<b>State</b>	State of the DHCPv6 relay address binding table on the DHCPv6 client: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has an active IP address lease.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>REBINDING</b>—Client is broadcasting a request to renew the IP address lease.</li> <li>• <b>RELEASE</b>—Client is releasing the IP address lease.</li> <li>• <b>RENEWING</b>—Client is sending a request to renew the IP address lease.</li> <li>• <b>REQUESTING</b>—Client is requesting a DHCPv6 server.</li> <li>• <b>SELECTING</b>—Client is receiving offers from DHCPv6 servers.</li> </ul>	<b>brief detail</b>
<b>Interface</b>	Incoming client interface.	<b>brief</b>
<b>Lease Expires</b>	Date and time at which the client's IP address lease expires.	<b>detail</b>



Table 156: show dhcpv6 relay binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Lease Expires in	Number of seconds in which the lease expires.	detail
Preferred Lease Expires	Date and UTC time at which the client's IPv6 prefix expires.	detail
Preferred Lease Expires in	Number of seconds at which the client's IPv6 prefix expires.	detail
Lease Start	Date and time at which the client's IP address lease started.	detail
Lease time violated	Lease time violation has occurred.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server Address	IP address of the DHCPv6 server.  Displays <b>unknown</b> for a DHCPv6 relay agent in a multi-relay topology that is not directly adjacent to the DHCPv6 server and does not detect the IP address of the server. In that case, the output instead displays the <b>Next Hop Server Facing Relay</b> field.	detail
Next Hop Server Facing Relay	Next-hop address in the direction of the DHCPv6 server.	detail
Server Interface	Interface of the DHCPv6 server.	detail
Relay Address	IP address of the relay.	detail
Client Pool Name	Address pool that granted the client lease.	detail
Client ID Length	Length of client ID.	All levels
Client Id	Client ID.	All levels
Generated Circuit ID	Circuit ID generated by the DHCPv6 Interface-ID option (option 18)	detail
Generated Remote ID Enterprise Number	The Juniper Networks IANA private enterprise number	detail
Generated Remote ID	Remote ID generated by the DHCPv6 Remote-ID option (option 37)	detail

## Sample Output

### show dhcpv6 relay binding

```

user@host> show dhcpv6 relay binding
Prefix          Session Id  Expires  State  Interface  Client DUID
2001:bd8:3c4d:15::/64  1          83720    BOUND  ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
2001:bd8:3c4d:16::/64  2          83720    BOUND  ge-1/0/0.0

```

```

LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:bd8:3c4d:17::/64      3      83720      BOUND      ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:bd8:3c4d:18::/64      4      83720      BOUND      ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:bd8:3c4d:19::/64      5      83720      BOUND      ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:bd8:3c4d:20::/64      6      83720      BOUND      ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06

```

### show dhcpv6 relay binding (Address)

```

user@host> show dhcpv6 relay binding 2001:bd8:1111:2222::/64 detail
Session Id: 1
  Client IPv6 Prefix:      2001:bd8:3c4d:15::/64
  Client DUID:              LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

  State:                    BOUND(RELAY_STATE_BOUND)
  Lease Expires:            2011-05-25 07:12:09 PDT
  Lease Expires in:         77115 seconds
  Preferred Lease Expires:  2012-07-24 00:18:14 UTC
  Preferred Lease Expires in: 600 seconds
  Lease Start:              2011-05-24 07:12:09 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Address:           2008:aaaa:bbbb::1
  Server Interface:         none
  Relay Address:            2001:bd8:1111:2222::
  Client Pool Name:         pool-25
  Client Id Length:         14
  Client Id:
    /0x00010001/0x4bfa26af/0x00109400/0x0001

```

### show dhcpv6 relay binding detail (Client ID)

```

user@host> show dhcpv6 relay binding 14/0x00010001/0x4bfa26af/0x00109400/0x0001
detail
Session Id: 1
  Client IPv6 Prefix:      2001:bd8:3c4d:15::/64
  Client DUID:              LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

  State:                    BOUND(RELAY_STATE_BOUND)
  Lease Expires:            2011-05-25 07:12:09 PDT
  Lease Expires in:         77115 seconds
  Preferred Lease Expires:  2012-07-24 00:18:14 UTC
  Preferred Lease Expires in: 600 seconds
  Lease Start:              2011-05-24 07:12:09 PDT
  Lease time violated:      yes
  Incoming Client Interface: ge-1/0/0.0
  Server Address:           2008:aaaa:bbbb::1
  Server Interface:         none
  Relay Address:            2001:bd8:1111:2222::
  Client Pool Name:         pool-25
  Client Id Length:         14
  Client Id:
    /0x00010001/0x4bfa26af/0x00109400/0x0001

```

### show dhcpv6 relay binding detail

```

user@host> show dhcpv6 relay binding detail
Session Id: 1
  Client IPv6 Prefix:      2001:bd8:3c4d:15::/64

```

```

Client DUID:                               LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

State:                                     BOUND(RELAY_STATE_BOUND)
Lease Expires:                            2011-05-25 07:12:09 PDT
Lease Expires in:                          77115 seconds
Preferred Lease Expires:                   2012-07-24 00:18:14 UTC
Preferred Lease Expires in:                600 seconds
Lease Start:                              2011-05-24 07:12:09 PDT
Lease time violated:                       yes
Incoming Client Interface:                 ge-1/0/0.0
Server Address:                           2008:aaaa:bbbb::1
Server Interface:                         none
Relay Address:                             2001:bd8:1111:2222::
Client Pool Name:                          pool-25
Client Id Length:                          14
Client Id:                                /0x00010001/0x4bfa26af/0x00109400/0x0001
Generated Remote ID Enterprise Number:     1411
Generated Remote ID:                       host:ge-1/0/0:100

```

### show dhcpv6 relay binding detail (Multi-Relay Topology)

```

user@host > show dhcpv6 relay binding detail
Session Id: 13
Client IPv6 Prefix:                       3000:0:0:8001::5/128
Client DUID:                              LL0x1-00:00:65:03:01:02
State:                                    BOUND(DHCPV6_RELAY_STATE_BOUND)
Lease Expires:                            2011-11-21 06:14:50 PST
Lease Expires in:                          293 seconds
Preferred Lease Expires:                   2012-07-24 00:18:14 UTC
Preferred Lease Expires in:                600 seconds
Lease Start:                              2011-11-21 06:09:50 PST
Incoming Client Interface:                 ge-1/0/0.0
Server Address:                           unknown
Next Hop Server Facing Relay:              4000::2
Server Interface:                         none
Client Id Length:                          10
Client Id:                                /0x00030001/0x00006503/0x0102

```

### show dhcpv6 relay binding (Session ID)

```

user@host> show dhcpv6 relay binding 41
Prefix          Session Id Expires   State   Interface   Client DUID
2001:bd8:3c4d:15::/64  41      78837   BOUND   ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

```

### show dhcpv6 relay binding (Interfaces VLAN)

```

user@host> show dhcpv6 relay binding ge-1/0/0:100-200
Prefix          Session Id Expires   State   Interface   Client DUID
2001:DB8::/32   11        87583   BOUND   ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   12        87583   BOUND   ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

### show dhcpv6 relay binding (Interfaces Wildcard)

```

user@host> show dhcpv6 relay binding demux0
Prefix          Session Id Expires   State   Interface   Client DUID
2001:DB8::/32   30        79681   BOUND   demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   31        79681   BOUND   demux0.1073741825

```

```
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32      32      79681    BOUND    demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

#### show dhcpv6 relay binding (Interfaces Wildcard)

```
user@host> show dhcpv6 relay binding ge-1/3/*
Prefix          Session Id Expires State Interface Client DUID
2001:DB8::/32   22      79681    BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   33      79681    BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32   24      79681    BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

#### show dhcpv6 relay binding summary

```
user@host> show dhcpv6 relay binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

## show dhcpv6 relay statistics

<b>Syntax</b>	<b>show dhcpv6 relay statistics</b> <code>&lt;logical-system <i>logical-system-name</i>&gt;</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Display Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear dhcpv6 relay statistics on page 1741</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show dhcpv6 relay statistics on page 1758</a>
<b>Output Fields</b>	<a href="#">Table 157 on page 1757</a> lists the output fields for the <b>show dhcpv6 relay statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 157: show dhcpv6 relay statistics Output Fields**

Field Name	Field Description
DHCPv6 Packets dropped	<p>Number of packets discarded by the extended DHCPv6 relay agent application due to errors. Only nonzero statistics appear in the <b>Packets dropped</b> output. When all of the Packets dropped statistics are 0 (zero), only the <b>Total</b> field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the DHCPv6 relay agent application.</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified.</li> <li>• <b>Bad send</b>—Number of packets that the extended DHCP relay application could not send.</li> <li>• <b>Bad src address</b>—Number of packets discarded because the family type was not AF_INET6.</li> <li>• <b>No client id</b>—Number of packets discarded because they could not be matched to a client.</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No safd</b>—Number of packets discarded because they arrived on an unconfigured interface.</li> <li>• <b>Short packet</b>—Number of packets discarded because they were too short.</li> <li>• <b>Relay hop count</b>—Number of packets discarded because the hop count in the packet exceeded 32.</li> </ul>

Table 157: show dhcpv6 relay statistics Output Fields (*continued*)

Field Name	Field Description
<b>Messages received</b>	<p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> <li><b>DHCPV6_DECLINE</b>—Number of DHCPv6 PDUs of type DECLINE received</li> <li><b>DHCPV6_SOLICIT</b>—Number of DHCPv6 PDUs of type SOLICIT received</li> <li><b>DHCPV6_INFORMATION_REQUEST</b>—Number of DHCPv6 PDUs of type INFORMATION-REQUEST received</li> <li><b>DHCPV6_RELEASE</b>—Number of DHCPv6 PDUs of type RELEASE received</li> <li><b>DHCPV6_REQUEST</b>—Number of DHCPv6 PDUs of type REQUEST received</li> <li><b>DHCPV6_CONFIRM</b>—Number of DHCPv6 PDUs of type CONFIRM received</li> <li><b>DHCPV6_RENEW</b>—Number of DHCPv6 PDUs of type RENEW received</li> <li><b>DHCPV6_REBIND</b>—Number of DHCPv6 PDUs of type REBIND received</li> <li><b>DHCPV6_RELAY_REPL</b>—Number of DHCPv6 PDUs of type RELAY-REPL received</li> </ul>
<b>Messages sent</b>	<p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> <li><b>DHCPV6_ADVERTISE</b>—Number of DHCPv6 ADVERTISE PDUs transmitted</li> <li><b>DHCP_REPLY</b>—Number of DHCPv6 REPLY PDUs transmitted</li> <li><b>DHCP_RECONFIGURE</b>—Number of DHCPv6 RECONFIGURE PDUs transmitted</li> <li><b>DHCP_RELAY_FORW</b>—Number of DHCPv6 RELAY-FORW PDUs transmitted</li> </ul>
<b>Packets forwarded</b>	<p>Number of packets forwarded by the extended DHCPv6 relay agent application.</p> <ul style="list-style-type: none"> <li><b>FWD REQUEST</b>—Number of DHCPv6 REQUEST packets forwarded</li> <li><b>FWD REPLY</b>—Number of DHCPv6 REPLY packets forwarded</li> </ul>
<b>External Server Response</b>	State of the external DHCP server responsiveness.

## Sample Output

### show dhcpv6 relay statistics

```

user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
  Total                  1
  Lease Time Violation   1

Messages received:
  DHCPV6_DECLINE         0
  DHCPV6_SOLICIT         10
  DHCPV6_INFORMATION_REQUEST 0
  DHCPV6_RELEASE         0
  DHCPV6_REQUEST         10
  DHCPV6_CONFIRM         0
  DHCPV6_RENEW           0
  DHCPV6_REBIND          0
  DHCPV6_RELAY_REPL      0

Messages sent:
  DHCPV6_ADVERTISE       0
  DHCPV6_REPLY            0

```

DHCPV6_RECONFIGURE	0
DHCPV6_RELAY_FORW	0
Packets forwarded:	
Total	4
FWD REQUEST	2
FWD REPLY	2
External Server Response:	
State	Responding

## show route extensive

<b>List of Syntax</b>	<a href="#">Syntax on page 1760</a> <a href="#">Syntax (EX Series Switches) on page 1760</a>
<b>Syntax</b>	show route extensive <destination-prefix> <logical-system (all   logical-system-name)>
<b>Syntax (EX Series Switches)</b>	show route extensive <destination-prefix>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Display extensive information about the active entries in the routing tables.
<b>Options</b>	<b>none</b> —Display all active entries in the routing table.  <b>destination-prefix</b> —(Optional) Display active entries for the specified address or range of addresses.  <b>logical-system (all   logical-system-name)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show route extensive on page 1767</a> <a href="#">show route extensive (Access Route) on page 1773</a> <a href="#">show route extensive (BGP PIC Edge) on page 1774</a> <a href="#">show route extensive (FRR and LFA) on page 1774</a> <a href="#">show route extensive (Route Reflector) on page 1775</a> <a href="#">show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 1775</a> <a href="#">show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 1776</a>
<b>Output Fields</b>	<a href="#">Table 158 on page 1760</a> describes the output fields for the <b>show route extensive</b> command. Output fields are listed in the approximate order in which they appear.

**Table 158: show route extensive Output Fields**

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.



Table 158: show route extensive Output Fields (*continued*)

Field Name	Field Description
<i>number routes</i>	<p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active).</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive).</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy).</li> </ul>
<i>route-destination</i> (entry, announced)	<p>Route destination (for example: 10.0.0.1/24). The <b>entry</b> value is the number of route for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul>
<b>TSI</b>	Protocol header information.
<b>label stacking</b>	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>
<b>[protocol, preference]</b>	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>

Table 158: show route extensive Output Fields (*continued*)

Field Name	Field Description
<b>Level</b>	(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
<b>Route Distinguisher</b>	IP subnet augmented with a 64-bit prefix.
<b>PMSI</b>	Provider multicast service interface (MVPN routing table).
<b>Next-hop type</b>	Type of next hop. For a description of possible values for this field, see the Output Field table in the <a href="#">show route detail</a> command.
<b>Next-hop reference count</b>	Number of references made to the next hop.
<b>Flood nexthop branches exceed maximum message</b>	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
<b>Source</b>	IP address of the route source.
<b>Next hop</b>	Network layer address of the directly reachable neighboring system.
<b>via</b>	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b> . This field can also contain the following information: <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> </ul>
<b>Label-switched-path lsp-path-name</b>	Name of the LSP used to reach the next hop.
<b>Label operation</b>	MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).
<b>Offset</b>	Whether the metric has been increased or decreased by an offset value.
<b>Interface</b>	(Local only) Local interface name.
<b>Protocol next hop</b>	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.

Table 158: show route extensive Output Fields (*continued*)

Field Name	Field Description
<b><i>label-operation</i></b>	MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).
<b>Indirect next hops</b>	<p>When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.</p> <p>When BGP PIC Edge is enabled, the output lines that contain <b>Indirect next hop: weight</b> follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values:</p> <ul style="list-style-type: none"> <li>• 0x1 indicates active next hops.</li> <li>• 0x4000 indicates passive next hops.</li> </ul>
<b>State</b>	State of the route (a route can be in more than one state). See the Output Field table in the <a href="#">show route detail</a> command.
<b>Session ID</b>	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
<b>Weight</b>	<p>Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown.</p> <p>For sample output, see <a href="#">show route table</a>.</p>

Table 158: show route extensive Output Fields (*continued*)

Field Name	Field Description
Inactive reason	<p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> <li>• <b>Active preferred</b>—Currently active route was selected over this route.</li> <li>• <b>Always compare MED</b>—Path with a lower multiple exit discriminator (MED) is available.</li> <li>• <b>AS path</b>—Shorter AS path is available.</li> <li>• <b>Cisco Non-deterministic MED selection</b>—Cisco nondeterministic MED is enabled and a path with a lower MED is available.</li> <li>• <b>Cluster list length</b>—Path with a shorter cluster list length is available.</li> <li>• <b>Forwarding use only</b>—Path is only available for forwarding purposes.</li> <li>• <b>IGP metric</b>—Path through the next hop with a lower IGP metric is available.</li> <li>• <b>IGP metric type</b>—Path with a lower OSPF link-state advertisement type is available.</li> <li>• <b>Interior &gt; Exterior &gt; Exterior via Interior</b>—Direct, static, IGP, or EBGP path is available.</li> <li>• <b>Local preference</b>—Path with a higher local preference value is available.</li> <li>• <b>Next hop address</b>—Path with a lower metric next hop is available.</li> <li>• <b>No difference</b>—Path from a neighbor with a lower IP address is available.</li> <li>• <b>Not Best in its group</b>—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed).</li> <li>• <b>Number of gateways</b>—Path with a higher number of next hops is available.</li> <li>• <b>Origin</b>—Path with a lower origin code is available.</li> <li>• <b>OSPF version</b>—Path does not support the indicated OSPF version.</li> <li>• <b>RIB preference</b>—Route from a higher-numbered routing table is available.</li> <li>• <b>Route distinguisher</b>—64-bit prefix added to IP subnets to make them unique.</li> <li>• <b>Route metric or MED comparison</b>—Route with a lower metric or MED is available.</li> <li>• <b>Route preference</b>—Route with a lower preference value is available.</li> <li>• <b>Router ID</b>—Path through a neighbor with a lower ID is available.</li> <li>• <b>Unusable path</b>—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved.</li> <li>• <b>Update source</b>—Last tiebreaker is the lowest IP address value.</li> </ul>
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	<p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see <a href="#">show route table</a>.</p>

Table 158: show route extensive Output Fields (*continued*)

Field Name	Field Description
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that announce this route. <b>n-Resolve inet</b> indicates that the route is used for route resolution for next hops found in the routing table. <b>n</b> is an index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Unverified</b>—Indicates that origin validation is not enabled for the BGP peers.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.

Table 158: show route extensive Output Fields (*continued*)

Field Name	Field Description
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the <a href="#">show route detail</a> command for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: <b>none</b> or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.

Table 158: show route extensive Output Fields (*continued*)

Field Name	Field Description
<b>Forwarding nexthops</b>	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

## Sample Output

### show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 29
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 69
        Age: 1:34:06
        Task: RT
        Announcement bits (2): 0-KRT 3-Resolve tree 2
        AS path: I

10.31.1.0/30 (2 entries, 1 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 2
        Next hop: via so-0/3/0.0, selected
        State: <Active Int>
        Local AS: 69
        Age: 1:32:40
        Task: IF
        Announcement bits (1): 3-Resolve tree 2
        AS path: I
    OSPF Preference: 10
        Next-hop reference count: 1
        Next hop: via so-0/3/0.0, selected
        State: <Int>
        Inactive reason: Route Preference
        Local AS: 69
        Age: 1:32:40 Metric: 1
        Area: 0.0.0.0
        Task: OSPF
        AS path: I

10.31.1.1/32 (1 entry, 1 announced)
    *Local Preference: 0
        Next hop type: Local
        Next-hop reference count: 7
        Interface: so-0/3/0.0
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:32:43
        Task: IF
        Announcement bits (1): 3-Resolve tree 2
        AS path: I

```

```
...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.31.2.0/30 -> {10.31.1.6}
    *OSPF   Preference: 10
            Next-hop reference count: 9
            Next hop: via so-0/3/0.0
            Next hop: 10.31.1.6 via ge-3/1/0.0, selected
            State: <Active Int>
            Local AS:    69
            Age: 1:32:19   Metric: 2
            Area: 0.0.0.0
            Task: OSPF
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kerne1 224.0.0.2/32 -> {}
    *PIM    Preference: 0
            Next-hop reference count: 18
            State: <Active NoReadvrt Int>
            Local AS:    69
            Age: 1:34:08
            Task: PIM Recv
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kerne1 224.0.0.22/32 -> {}
    *IGMP   Preference: 0
            Next-hop reference count: 18
            State: <Active NoReadvrt Int>
            Local AS:    69
            Age: 1:34:06
            Task: IGMP
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
    State: <FlashAll>
    *RSVP   Preference: 7
            Next-hop reference count: 6
            Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
            Label-switched-path green-r1-r3
            Label operation: Push 100096
            State: <Active Int>
            Local AS:    69
            Age: 1:28:12   Metric: 2
            Task: RSVP
            Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
            AS path: I
```



```

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP Preference: 7
    Next-hop reference count: 6
    Next hop: via so-0/3/0.0 weight 0x1, selected
    Label-switched-path green-r1-r2
    State: <Active Int>
    Local AS: 69
    Age: 1:28:12 Metric: 1
    Task: RSVP
    Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
    AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:34:07
    Task: IF
    AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
  *MPLS Preference: 0
    Next hop type: Receive
    Next-hop reference count: 6
    State: <Active Int>
    Local AS: 69
    Age: 1:34:08 Metric: 1
    Task: MPLS
    Announcement bits (1): 0-KRT
    AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299840 (1 entry, 1 announced)
TSI:
KRT in-kernel 299840 /52 -> {indirect(1048575)}
  *RSVP Preference: 7/2
    Next hop type: Flood
    Address: 0x9174a30
    Next-hop reference count: 4
    Next hop type: Router, Next hop index: 798
    Address: 0x9174c28
    Next-hop reference count: 2
    Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
    Label-switched-path R2-to-R4-2p2mp

```

```

Label operation: Pop
Next hop type: Router, Next hop index: 1048574
Address: 0x92544f0
Next-hop reference count: 2
Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
Label-switched-path R2-to-R200-p2mp
Label operation: Pop
Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
Label operation: Pop
State: <Active Int>
Age: 1:29      Metric: 1
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I...

```

800010 (1 entry, 1 announced)

TSI:

```

KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: via vt-3/2/0.32769, selected
    Label operation: Pop
    State: <Active Int>
    Age: 1:31:53
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

```

vt-3/2/0.32769 (1 entry, 1 announced)

TSI:

```

KRT in-kernel vt-3/2/0.32769.0 /16 -> {indirect(1048574)}
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
    Label-switched-path green-r1-r3
    Label operation: Push 800012, Push 100096(top)
    Protocol next hop: 10.255.70.103
    Push 800012
    Indirect next hop: 87272e4 1048574
    State: <Active Int>
    Age: 1:31:53      Metric2: 2
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Indirect next hops: 1
      Protocol next hop: 10.255.70.103 Metric: 2
      Push 800012
      Indirect next hop: 87272e4 1048574
      Indirect path forwarding next hops: 1
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
        10.255.70.103/32 Originating RIB: inet.3
        Metric: 2      Node path count: 1
        Forwarding nexthops: 1
        Nexthop: 10.31.1.6 via ge-3/1/0.0

```

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)

```

*Direct Preference: 0
  Next hop type: Interface
  Next-hop reference count: 1
  Next hop: via lo0.0, selected
  State: <Active Int>
  Local AS: 69
  Age: 1:34:07
  Task: IF
  AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0
  Next hop type: Interface
  Next-hop reference count: 1
  Next hop: via lo0.0, selected
  State: <Active NoReadvrt Int>
  Local AS: 69
  Age: 1:34:07
  Task: IF
  AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
  *MLD Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:06
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0

```

Next hop type: Interface  
Next-hop reference count: 1  
Next hop: via lo0.16385, selected  
State: <Active NoReadvrt Int>  
Age: 1:34:07  
Task: IF  
AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)

\*BGP Preference: 170/-101  
Route Distinguisher: 10.255.70.103:1  
Next-hop reference count: 7  
Source: 10.255.70.103  
Protocol next hop: 10.255.70.103  
Indirect next hop: 2 no-forward  
State: <Secondary Active Int Ext>  
Local AS: 69 Peer AS: 69  
Age: 1:28:12 Metric2: 1  
Task: BGP\_69.10.255.70.103+179  
Announcement bits (1): 0-green-l2vpn  
AS path: I  
Communities: target:11111:1 Layer2-info: encaps:VPLS,  
control flags:, mtu: 0  
Label-base: 800008, range: 8  
Localpref: 100  
Router ID: 10.255.70.103  
Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)

TSI:

Page 0 idx 0 Type 1 val 8699540

\*L2VPN Preference: 170/-1  
Next-hop reference count: 5  
Protocol next hop: 10.255.71.52  
Indirect next hop: 0 -  
State: <Active Int Ext>  
Age: 1:34:03 Metric2: 1  
Task: green-l2vpn  
Announcement bits (1): 1-BGP.0.0.0.0+179  
AS path: I  
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,  
mtu: 0  
Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)

TSI:

Page 0 idx 0 Type 1 val 8699528

\*L2VPN Preference: 170/-101  
Next-hop reference count: 5  
Protocol next hop: 10.255.71.52  
Indirect next hop: 0 -  
State: <Active Int Ext>  
Age: 1:34:03 Metric2: 1  
Task: green-l2vpn  
Announcement bits (1): 1-BGP.0.0.0.0+179  
AS path: I  
Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0  
Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)

```
*L2CKT Preference: 7
  Next hop: via so-1/1/2.0 weight 1, selected
  Label-switched-path my-lsp
  Label operation: Push 100000[0]
  Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
  State: <Active Int>
  Local AS: 99
  Age: 10:21
  Task: l2 circuit
  Announcement bits (1): 0-LDP
  AS path: I
  VC Label 100000, MTU 1500, VLAN ID 512
```

55.0.0.0/24 (1 entry, 1 announced)

TSI:

KRT queued (pending) add

55.0.0.0/24 -> {Push 300112}

```
*BGP Preference: 170/-101
  Next hop type: Router
  Address: 0x925c208
  Next-hop reference count: 2
  Source: 10.0.0.9
  Next hop: 10.0.0.9 via ge-1/2/0.15, selected
  Label operation: Push 300112
  Label TTL action: prop-ttl
  State: <Active Ext>
  Local AS: 7019 Peer AS: 13979
  Age: 1w0d 23:06:56
  AIGP: 25
  Task: BGP_13979.10.0.0.9+56732
  Announcement bits (1): 0-KRT
  AS path: 13979 7018 I
  Accepted
  Route Label: 300112
  Localpref: 100
  Router ID: 10.9.9.1
```

### show route extensive (Access Route)

user@host> show route 13.160.0.102 extensive

inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)

13.160.0.102/32 (1 entry, 1 announced)

TSI:

KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}

OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern

```
*Access Preference: 13
  Next-hop reference count: 78472
  Next hop: 13.160.0.2 via fe-0/0/0.0, selected
  State: <Active Int>
```

Age: 12

```
Task: RPD Unix Domain Server./var/run/rpd_serv.local
Announcement bits (2): 0-KRT 1-OSPFv2
AS path: I
```

### show route extensive (BGP PIC Edge)

```
user@host> show route 1.1.1.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  1.1.1.6/32 (3 entries, 2 announced)
    State: <CalcForwarding>
    TSI:
    KRT in-kerne1 1.1.1.6/32 -> {indirect(1048574), indirect(1048577)}
    Page 0 idx 0 Type 1 val 9219e30
    Nexthop: Self
    AS path: [2] 3 I
    Communities: target:2:1
    Path 1.1.1.6 from 1.1.1.4 Vector len 4. Val: 0
  ..
    #Multipath Preference: 255
    Next hop type: Indirect
    Address: 0x93f4010
    Next-hop reference count: 2
  ..
    Protocol next hop: 1.1.1.4
    Push 299824
    Indirect next hop: 944c000 1048574 INH Session ID: 0x3
    Indirect next hop: weight 0x1
    Protocol next hop: 1.1.1.5
    Push 299824
    Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
    Indirect next hop: weight 0x4000
    State: <ForwardingOnly Int Ext>
    Inactive reason: Forwarding use only
    Age: 25 Metric2: 15
    Validation State: unverified
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: 3 I
    Communities: target:2:1
```

### show route extensive (FRR and LFA)

```
user@host> show route 20.31.2.0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
  20.31.2.0/24 (2 entries, 1 announced)
    State: FlashAll
    TSI:
    KRT in-kerne1 20.31.2.0/24 -> {Push 299776, Push 299792}
    *RSVP Preference: 7/1
    Next hop type: Router, Next hop index: 1048574
    Address: 0xbbbc010
    Next-hop reference count: 5
    Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
    Label-switched-path europa-d-to-europa-e
    Label operation: Push 299776
    Label TTL action: prop-ttl
    Session Id: 0x201
    Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
    Label-switched-path europa-d-to-europa-e
    Label operation: Push 299792
    Label TTL action: prop-ttl
    Session Id: 0x202
    State: Active Int
    Local AS: 100
    Age: 5:31 Metric: 2
```

```

Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I

```

### show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
*BGP Preference: 170/-101
Source: 192.168.4.214
Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
State: <Active Int Ext>
Local AS: 10458 Peer AS: 10458
Age: 3:09 Metric: 0 Metric2: 0
Task: BGP_10458.192.168.4.214+1033
Announcement bits (2): 0-KRT 4-Resolve inet.0
AS path: 3944 7777 I <Originator>
Cluster list: 1.1.1.1
Originator ID: 10.255.245.88
Communities: 7777:7777
Localpref: 100
Router ID: 4.4.4.4
Indirect next hops: 1
    Protocol next hop: 207.17.136.192 Metric: 0
    Indirect next hop: 84ac908 40
    Indirect path forwarding next hops: 0
    Next hop type: Discard

```

### show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
*LDP Preference: 9
Next hop type: Flood
Next-hop reference count: 3
Address: 0x9097d90
Next hop: via vt-0/1/0.1
Next-hop index: 661
Label operation: Pop
Address: 0x9172130
Next hop: via so-0/0/3.0
Next-hop index: 654
Label operation: Swap 299872
State: **Active Int>

```

```

Local AS: 1001
Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

### show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
    Next hop type: Flood
    Address: 0x2735208
    Next-hop reference count: 3
    Next hop type: Router, Next hop index: 1397
    Address: 0x2735d2c
    Next-hop reference count: 3
    Next hop: 1.3.8.2 via ge-1/2/22.0
    Label operation: Pop
    Load balance label: None;
    Next hop type: Router, Next hop index: 1395
    Address: 0x2736290
    Next-hop reference count: 3
    Next hop: 1.3.4.2 via ge-1/2/18.0
    Label operation: Pop
    Load balance label: None;
    State: <Active Int AckRequest MulticastRPF>
    Local AS: 10
    Age: 54:05      Metric: 1
    Validation State: unverified
    Task: LDP
    Announcement bits (1): 0-KRT
    AS path: I
    FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
      Primary Upstream : 1.1.1.3:0--1.1.1.2:0
        RPF Nexthops :
          ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
          ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
      Backup Upstream : 1.1.1.3:0--1.1.1.6:0
        RPF Nexthops :
          ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
          ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```



## show route protocol

<b>List of Syntax</b>	<a href="#">Syntax on page 1777</a> <a href="#">Syntax (EX Series Switches) on page 1777</a>
<b>Syntax</b>	<pre>show route protocol <i>protocol</i> &lt;brief   detail   extensive   terse&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show route protocol <i>protocol</i> &lt;brief   detail   extensive   terse&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>ospf2</b> and <b>ospf3</b> options introduced in Junos OS Release 9.2.</p> <p><b>ospf2</b> and <b>ospf3</b> options introduced in Junos OS Release 9.2 for EX Series switches.</p> <p><b>flow</b> option introduced in Junos OS Release 10.0.</p> <p><b>flow</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p>
<b>Description</b>	Display the route entries in the routing table that were learned from a particular protocol.
<b>Options</b>	<p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>protocol</i></b>—Protocol from which the route was learned:</p> <ul style="list-style-type: none"> <li>• <b>access</b>—Access route for use by DHCP application</li> <li>• <b>access-internal</b>—Access-internal route for use by DHCP application</li> <li>• <b>aggregate</b>—Locally generated aggregate route</li> <li>• <b>arp</b>—Route learned through the Address Resolution Protocol</li> <li>• <b>atmvpn</b>—Asynchronous Transfer Mode virtual private network</li> <li>• <b>bgp</b>—Border Gateway Protocol</li> <li>• <b>ccc</b>—Circuit cross-connect</li> <li>• <b>direct</b>—Directly connected route</li> <li>• <b>dvmrp</b>—Distance Vector Multicast Routing Protocol</li> <li>• <b>esis</b>—End System-to-Intermediate System</li> <li>• <b>flow</b>—Locally defined flow-specification route</li> <li>• <b>frr</b>—Precomputed protection route or backup route used when a link goes down</li> <li>• <b>isis</b>—Intermediate System-to-Intermediate System</li> <li>• <b>ldp</b>—Label Distribution Protocol</li> <li>• <b>l2circuit</b>—Layer 2 circuit</li> </ul>

- **l2vpn**—Layer 2 virtual private network
- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



**NOTE:** EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><li>• <i>MPLS Feature Support on Juniper Switches</i></li></ul>
List of Sample Output	<a href="#">show route protocol access on page 1779</a> <a href="#">show route protocol access-internal extensive on page 1779</a> <a href="#">show route protocol arp on page 1779</a> <a href="#">show route protocol bgp on page 1780</a> <a href="#">show route protocol bgp detail on page 1780</a> <a href="#">show route protocol bgp extensive on page 1780</a> <a href="#">show route protocol bgp terse on page 1781</a> <a href="#">show route protocol direct on page 1781</a> <a href="#">show route protocol frr on page 1782</a> <a href="#">show route protocol l2circuit detail on page 1782</a> <a href="#">show route protocol l2vpn extensive on page 1783</a> <a href="#">show route protocol ldp on page 1784</a> <a href="#">show route protocol ldp extensive on page 1784</a> <a href="#">show route protocol ospf (Layer 3 VPN) on page 1785</a> <a href="#">show route protocol ospf detail on page 1786</a>

[show route protocol rip on page 1786](#)  
[show route protocol rip detail on page 1786](#)  
[show route protocol ripng table inet6 on page 1787](#)  
[show route protocol static detail on page 1787](#)

**Output Fields** For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

## Sample Output

### show route protocol access

```

user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0

```

### show route protocol access-internal extensive

```

user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
    *Access-internal Preference: 12
        Next-hop reference count: 200000
        Next hop: 13.160.0.2 via fe-0/0/0.0, selected
        State: <Active Int>
    Age: 36
        Task: RPD Unix Domain Server./var/run/rpd_serv.local
        Announcement bits (1): 0-KRT
        AS path: I

```

### show route protocol arp

```

user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.4/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.5/32      [ARP/4294967293] 00:04:32, from 20.20.1.1
                  Unusable
20.20.1.6/32      [ARP/4294967293] 00:04:34, from 20.20.1.1
                  Unusable
20.20.1.7/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable

```

```

20.20.1.8/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.9/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.10/32     [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.11/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
20.20.1.12/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
20.20.1.13/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
...

```

### show route protocol bgp

```

user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21    *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                  AS path: 10458 14203 2914 4788 4788 I
                  > to 192.168.167.254 via fxp0.0

```

### show route protocol bgp detail

```

user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24    (1 entry, 1 announced)
   *BGP           Preference: 170/-101
                   Next hop type: Indirect
                   Next-hop reference count: 1006436
                   Source: 192.168.69.71
                   Next hop type: Router, Next hop index: 324
                   Next hop: 192.168.167.254 via fxp0.0, selected
                   Protocol next hop: 192.168.69.71
                   Indirect next hop: 8e166c0 342
                   State: <Active Ext>
                   Local AS: 69 Peer AS: 10458
                   Age: 6d 10:42:42 Metric2: 0
                   Task: BGP_10458.192.168.69.71+179
                   Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1

   AS path: 10458 14203 2914 4788 4788 I
   Communities: 2914:410 2914:2403 2914:3400
   Accepted
   Localpref: 100
   Router ID: 207.17.136.192

```

### show route protocol bgp extensive

```

user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
  Nexthop: Self
  AS path: [69] 10458 14203 2914 4788 4788 I
  Communities: 2914:410 2914:2403 2914:3400
  Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1

```

```

*BGP      Preference: 170/-101
          Next hop type: Indirect
          Next-hop reference count: 1006502
          Source: 192.168.69.71
          Next hop type: Router, Next hop index: 324
          Next hop: 192.168.167.254 via fxp0.0, selected
          Protocol next hop: 192.168.69.71
          Indirect next hop: 8e166c0 342
          State: <Active Ext>
          Local AS: 69 Peer AS: 10458
          Age: 6d 10:44:45 Metric2: 0
          Task: BGP_10458.192.168.69.71+179
          Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1
          AS path: 10458 14203 2914 4788 4788 I
          Communities: 2914:410 2914:2403 2914:3400
          Accepted
          Localpref: 100
          Router ID: 207.17.136.192
          Indirect next hops: 1
            Protocol next hop: 192.168.69.71
            Indirect next hop: 8e166c0 342
            Indirect path forwarding next hops: 1
              Next hop type: Router
              Next hop: 192.168.167.254 via fxp0.0
            192.168.0.0/16 Originating RIB: inet.0
              Node path count: 1
              Forwarding nexthops: 1
                Nexthop: 192.168.167.254 via fxp0.0

```

### show route protocol bgp terse

```

user@host> show route protocol bgp 192.168.64.0/21 terse

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
192.168.64.0/21   B 170      100          >100.1.3.2    10023 21 I

```

### show route protocol direct

```

user@host> show route protocol direct

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

8.8.8.0/24         *[Direct/0] 17w0d 10:31:49
> via fe-1/3/1.0
10.255.165.1/32    *[Direct/0] 25w4d 04:13:18
> via lo0.0
30.30.30.0/24      *[Direct/0] 17w0d 23:06:26
> via fe-1/3/2.0
192.168.164.0/22   *[Direct/0] 25w4d 04:13:20
> via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
*[Direct/0] 25w4d 04:13:21

```

```
> via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

abcd::10:255:165:1/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0
```

### show route protocol frr

```
user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.3 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.4 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32      *[FRR/200] 00:05:35, from 20.20.1.1
                  > to 20.20.1.5 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32      *[FRR/200] 00:05:37, from 20.20.1.1
                  > to 20.20.1.6 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.7 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.8 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.9 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32     *[FRR/200] 00:05:38, from 20.20.1.1
...

```

### show route protocol l2circuit detail

```
user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop: via ge-2/0/0.0, selected
        Label operation: Pop      Offset: 4
        State: <Active Int>
        Local AS: 99
        Age: 9:52
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

```

```

ge-2/0/0.0 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000, Push 100000(top)[0] Offset: -4
    Protocol next hop: 10.245.255.63
    Push 100000 Offset: -4
    Indirect next hop: 86af0c0 298
    State: <Active Int>
    Local AS: 99
    Age: 9:52
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

### show route protocol l2vpn extensive

```

user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
  *L2VPN Preference: 7
    Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
    Label operation: Pop Offset: 4
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
  *L2VPN Preference: 7
    Next hop: via so-0/0/1.0, selected

```

```
Label operation: Push 800000 Offset: -4
Protocol next hop: 10.255.14.220
Push 800000 Offset: -4
  Indirect next hop: 85142a0 288
State: <Active Int>
Local AS: 69
Age: 7:48
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I
Communities: target:69:1 Layer2-info: encaps:PPP,
control flags:2, mtu: 0
```

### show route protocol ldp

```
user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100080            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Swap 100000
```

### show route protocol ldp extensive

```
user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          Label operation: Push 100000
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

192.168.17.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>
          Local AS: 65500
```



```

Age: 1d 23:03:58      Metric: 1
Task: LDP
Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
AS path: I

private1__inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              Label operation: Pop
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              Label operation: Swap 100000
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 192.168.16.1/32

```

### show route protocol ospf (Layer 3 VPN)

```

user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30      * [OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.39.1.8/30      [OSPF/10] 00:05:18, metric 2

```

```

> via t3-3/2/0.0
10.255.14.171/32  *[OSPF/10] 00:05:18, metric 4
> via t3-3/2/0.0
10.255.14.179/32  *[OSPF/10] 00:05:18, metric 2
> via t3-3/2/0.0
224.0.0.5/32     *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30     [OSPF/10] 00:05:43, metric 1
> via so-0/2/2.0
10.255.14.173/32  *[OSPF/10] 00:05:43, metric 1
> via so-0/2/2.0
224.0.0.5/32     *[OSPF/10] 20:26:20, metric 1

```

### show route protocol ospf detail

```

user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
  OSPF   Preference: 10
        Nexthop: via so-0/2/2.0, selected
        State: <Int>
        Inactive reason: Route Preference
        Age: 6:25      Metric: 1
        Area: 0.0.0.0
        Task: VPN-AB-OSPF
        AS path: I
        Communities: Route-Type:0.0.0.0:1:0

...

```

### show route protocol rip

```

user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32  *[RIP/100] 20:24:34, metric 2
> to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32     *[RIP/100] 00:03:59, metric 1

```

### show route protocol rip detail

```

user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
  *RIP   Preference: 100
        Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
        State: <Active Int>
        Age: 20:25:02  Metric: 2
        Task: VPN-AB-RIPv2
        Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179

```

```
AS path: I
Route learned from 10.39.1.22 expires in 96 seconds
```

### show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

### show route protocol static detail

```
user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.13.10.0/23 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
```

State: <Active NoReadvrt Int Ext>  
Age: 7w3d 21:24:25  
Validation State: unverified  
Task: RT  
Announcement bits (1): 0-KRT  
AS path: I


---

## Other Operational Commands

---

- `clear security pki local-certificate`
- `clear system services dhcp binding`
- `clear system services dhcp conflict`
- `clear system services dhcp statistics`
- `request ipsec switch`
- `request security certificate (signed)`
- `request security certificate (unsigned)`
- `request security key-pair`
- `request security pki generate-key-pair`
- `request security pki local-certificate generate-self-signed`
- `show security pki local-certificate`
- `show system services dhcp binding`
- `show system services dhcp conflict`
- `show system services dhcp global`
- `show system services dhcp pool`
- `show system services dhcp statistics`
- `show system services service-deployment`
- `ssh`
- `telnet`

## clear security pki local-certificate

<b>Syntax</b>	clear security pki local-certificate <all   certificate-id <i>certificate-id-name</i>   system-generated>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for EX Series switches.
<b>Description</b>	Delete local digital certificates, certificate requests, and the corresponding public/private key pairs from the switch.
<b>Options</b>	all—(Optional) Delete all local digital certificates, certificate requests, and the corresponding public and private key pairs from the router.
<div>  <b>NOTE:</b> This option does not delete the automatically generated self-signed certificate or its public/private key pair.         </div>	
	<b>certificate-id <i>certificate-id-name</i></b> —(Optional) Delete the specified local digital certificate and corresponding public and private key pair.  <b>system-generated</b> —(Optional) Delete the automatically generated self-signed certificate.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Deleting Self-Signed Certificates (CLI Procedure) on page 1419</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear security pki local-certificate all on page 1789</a>
<b>Output Fields</b>	This command produces no output.

## Sample Output

clear security pki local-certificate all

```
user@switch> clear security pki local-certificate all
```

## clear system services dhcp binding

---

<b>Syntax</b>	clear system services dhcp binding <address>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(J Series routers and EX Series switches only) Remove obsolete IP address bindings on a Dynamic Host Configuration Protocol (DHCP) server and return them to the IP address pool.
<b>Options</b>	<b>address</b> —(Optional) Remove a specific IP address binding and return it to the address pool.
<b>Required Privilege Level</b>	view and system
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show system services dhcp binding on page 1803</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear system services dhcp binding on page 1790</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### clear system services dhcp binding

```
user@host> clear system services dhcp binding
```

## clear system services dhcp conflict

---

<b>Syntax</b>	clear system services dhcp conflict <address>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(J Series routers and EX Series switches only) Remove IP addresses from the Dynamic Host Configuration Protocol (DHCP) server conflict list and return them to the IP address pool.
<b>Options</b>	<b>address</b> —(Optional) Remove a specific IP address from the conflict list and return it to the address pool.
<b>Required Privilege Level</b>	view and system
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show system services dhcp conflict on page 1806</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear system services dhcp conflict on page 1791</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### clear system services dhcp conflict

```
user@host> clear system services dhcp conflict
```

## clear system services dhcp statistics

---

<b>Syntax</b>	clear system services dhcp statistics
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(J Series routers and EX Series switches only) Clear Dynamic Host Configuration Protocol (DHCP) server statistics.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view and system
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show system services dhcp statistics on page 1811</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear system services dhcp statistics on page 1792</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### clear system services dhcp statistics

```
user@host> clear system services dhcp statistics
```



## request ipsec switch

---

<b>Syntax</b>	<code>request ipsec switch (interface &lt;es-fpc/pic/port&gt;   security-associations &lt;sa-name&gt;)</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(Encryption interface on M Series, PTX Series, and T Series routers and EX Series switches only) Manually switch from the primary to the backup encryption services interface, or switch from the primary to the backup IP Security (IPsec) tunnel.
<b>Options</b>	<code>interface &lt;es-fpc/pic/port&gt;</code> —Switch to the backup encryption interface. <code>security-associations &lt;sa-name&gt;</code> —Switch to the backup tunnel.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ipsec redundancy</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request ipsec switch on page 1793</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request ipsec switch

```
user@host> request ipsec switch security-associations sa-private
```

## request security certificate (signed)

---

<b>Syntax</b>	<code>request security certificate enroll filename <i>filename</i> subject <i>subject</i> alternative-subject <i>alternative-subject</i> certification-authority <i>certification-authority</i> encoding (binary   pem) key-file <i>key-file</i> domain-name <i>domain-name</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Obtain a signed certificate from a certificate authority (CA). The signed certificate validates the CA and the owner of the certificate. The results are saved in a specified file to the <code>/var/etc/ikecert</code> directory.
<b>Options</b>	<p><b>filename <i>filename</i></b>—File that stores the certificate.</p> <p><b>subject <i>subject</i></b>—Distinguished name (<b>dn</b>), which consists of a set of components—for example, an organization (<b>o</b>), an organization unit (<b>ou</b>), a country (<b>c</b>), and a locality (<b>l</b>).</p> <p><b>alternative-subject <i>alternative-subject</i></b>—Tunnel source address.</p> <p><b>certification-authority <i>certification-authority</i></b>—Name of the certificate authority profile in the configuration.</p> <p><b>encoding (binary   pem)</b>—File format used for the certificate. The format can be a binary file or privacy-enhanced mail (PEM), an ASCII base64-encoded format. The default format is binary.</p> <p><b>key-file <i>key-file</i></b>—File containing a local private key.</p> <p><b>domain-name <i>domain-name</i></b>—Fully qualified domain name.</p>
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">request security certificate (signed) on page 1794</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request security certificate (signed)

```
user@host> request security certificate enroll filename host.crt subject c=uk,o=london  
alternative-subject 10.50.1.4 certification-authority verisign key-file host-1.prv domain-name  
host.juniper.net  
CA name: juniper.net CA file: ca_verisign  
local pub/private key pair: host.prv  
subject: c=uk,o=london domain name: host.juniper.net  
alternative subject: 10.50.1.4  
Encoding: binary  
Certificate enrollment has started. To view the status of your enrollment, check  
the key management process (kmd) log file at /var/log/kmd. <-----
```



## request security certificate (unsigned)

---

<b>Syntax</b>	<code>request security certificate enroll filename <i>filename</i> ca-file <i>ca-file</i> ca-name <i>ca-name</i> encoding (binary   perm) url <i>url</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Obtain a certificate from a certificate authority (CA). The results are saved in a specified file to the <code>/var/etc/ikecert</code> directory.
<b>Options</b>	<b>filename <i>filename</i></b> —File that stores the public key certificate.  <b>ca-file <i>ca-file</i></b> —Name of the certificate authority profile in the configuration.  <b>ca-name <i>ca-name</i></b> —Name of the certificate authority.  <b>encoding (binary   pem)</b> —File format used for the certificate. The format can be a binary file or privacy-enhanced mail (PEM), an ASCII base64-encoded format. The default value is <b>binary</b> .  <b>url <i>url</i></b> —Certificate authority URL.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">request security certificate (unsigned) on page 1796</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request security certificate (unsigned)

```
user@host> request security certificate enroll filename ca_verisign ca-file verisign ca-name
juniper.net urlxyzcompany URL
http://<verisign ca-name xyzcompany url>/cgi-bin/pkiclient.exe CA name: juniper.net
CA file: verisign Encoding: binary
Certificate enrollment has started. To view the status of your enrollment, check
the key management process (kmd) log file at /var/log/kmd. <-----
```

## request security key-pair

---

<b>Syntax</b>	<code>request security key-pair <i>filename</i></code> <code>&lt;size <i>key-size</i>&gt;</code> <code>&lt;type (rsa   dsa)&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Generate a public and private key pair for a digital certificate.
<b>Options</b>	<b><i>filename</i></b> —Name of a file in which to store the key pair.  <b><i>size key-size</i></b> —(Optional) Key size, in bits. The key size can be <b>512</b> , <b>1024</b> , or <b>2048</b> . The default value is <b>1024</b> .  <b><i>type</i></b> —(Optional) Algorithm used to encrypt the key: <ul style="list-style-type: none"> <li>• <b>rsa</b>—RSA algorithm. This is the default.</li> <li>• <b>dsa</b>—Digital signature algorithm with Secure Hash Algorithm (SHA).</li> </ul>
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">request security key-pair on page 1797</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request security key-pair

```
user@host> request security key-pair security-key-file
```

## request security pki generate-key-pair

---

<b>Syntax</b>	<code>request security pki generate-key-pair certificate-id <i>certificate-id-name</i></code> <code>&lt;size (512   1024   2048)&gt;</code> <code>&lt;type (dsa   rsa)&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for EX Series switches.
<b>Description</b>	Generate a public key infrastructure (PKI) public/private key pair for a local digital certificate.
<b>Options</b>	<p><b>certificate-id <i>certificate-id-name</i></b>—Name of the local digital certificate and the public/private key pair.</p> <p><b>size</b>—(Optional) Key pair size. The key pair size can be <b>512</b>, <b>1024</b>, or <b>2048</b> bits. If a key pair size is not specified, the default value, <b>1024</b> bits, is applied.</p> <p><b>type</b>—(Optional) The algorithm to be used for encrypting the public/private key pair. The encryption algorithm can be <b>dsa</b> or <b>rsa</b>. If an encryption algorithm is not specified, the default value, <b>rsa</b>, is applied.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Manually Generating Self-Signed Certificates on Switches (CLI Procedure) on page 1417</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request security pki generate-key-pair on page 1798</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request security pki generate-key-pair

```
user@switch> request security pki generate-key-pair certificate-id billy size 2048
Generated key pair billy, key size 2048 bits
```

## request security pki local-certificate generate-self-signed

<b>Syntax</b>	<code>request security pki local-certificate generate-self-signed certificate-id <i>certificate-id-name</i> domain-name <i>domain-name</i> ip-address <i>ip-address</i> email <i>email-address</i> subject <i>subject-distinguished-name</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for EX Series switches.
<b>Description</b>	Manually generate a self-signed certificate for the given distinguished name.
<b>Options</b>	<p><b>certificate-id <i>certificate-id-name</i></b>—Name of the local digital certificate and the public/private key pair.</p> <p><b>domain-name <i>domain-name</i></b>—Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.</p> <p><b>email <i>email-address</i></b>—E-mail address of the certificate holder.</p> <p><b>ip-address <i>ip-address</i></b>—IP address of the switch.</p> <p><b>subject <i>subject-distinguished-name</i></b>—Distinguished name format that contains the common name, department, company name, state, and country:</p> <ul style="list-style-type: none"> <li>• <b>CN</b>—Common name</li> <li>• <b>OU</b>—Organizational unit name</li> <li>• <b>O</b>—Organization name</li> <li>• <b>ST</b>—State</li> <li>• <b>C</b>—Country</li> </ul>
<b>Required Privilege Level</b>	maintenance security
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Manually Generating Self-Signed Certificates on Switches (CLI Procedure) on page 1417</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request security pki local-certificate generate-self-signed on page 1799</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request security pki local-certificate generate-self-signed

```
user@switch> request security pki local-certificate generate-self-signed certificate-id self-cert
subject cn=abc domain-name abc.net email jdoe@abc.net
Self-signed certificate generated and loaded successfully
```

## show security pki local-certificate

<b>Syntax</b>	show security pki local-certificate <brief   detail> <certificate-id <i>certificate-id-name</i> > <system-generated>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for EX Series switches.
<b>Description</b>	Display information about the local digital certificates and the corresponding public keys installed in the switch.
<b>Options</b>	<p><b>none</b>—(Same as brief) Display information about all local digital certificates and corresponding public keys.</p> <p><b>brief   detail</b>—(Optional) Display information about local digital certificates and corresponding public keys for the specified level of output.</p> <p><b>certificate-id <i>certificate-id-name</i></b>—(Optional) Display information about only the specified the local digital certificate and corresponding public keys.</p> <p><b>system-generated</b>—(Optional) Display information about the automatically generated self-signed certificate.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Manually Generating Self-Signed Certificates on Switches (CLI Procedure)</a> on page 1417</li> </ul>
<b>List of Sample Output</b>	<a href="#">show security pki local-certificate</a> on page 1801 <a href="#">show security pki local-certificate detail</a> on page 1802
<b>Output Fields</b>	<a href="#">Table 159 on page 1800</a> lists the output fields for the <b>show security pki local-certificate</b> command. Output fields are listed in the approximate order in which they appear.

Table 159: show security pki local-certificate Output Fields

Field Name	Field Description	Level of Output
Certificate identifier	Name of the digital certificate.	All levels
Certificate version	Revision number of the digital certificate.	detail
Serial number	Unique serial number of the digital certificate.	detail
Issued by	Authority that issued the digital certificate.	none brief
Issued to	Device that was issued the digital certificate.	none brief



Table 159: show security pki local-certificate Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Issuer</b>	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: <ul style="list-style-type: none"> <li>• <b>Common name</b>—Name of the authority.</li> <li>• <b>Organization</b>—Organization of origin.</li> <li>• <b>Organizational unit</b>—Department within an organization.</li> <li>• <b>State</b>—State of origin.</li> <li>• <b>Country</b>—Country of origin.</li> </ul>	<b>detail</b>
<b>Subject</b>	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: <ul style="list-style-type: none"> <li>• <b>Common name</b>—Name of the authority.</li> <li>• <b>Organization</b>—Organization of origin.</li> <li>• <b>Organizational unit</b>—Department within an organization.</li> <li>• <b>State</b>—State of origin.</li> <li>• <b>Country</b>—Country of origin.</li> </ul>	<b>detail</b>
<b>Alternate subject</b>	Domain name or IP address of the device related to the digital certificate.	<b>detail</b>
<b>Validity</b>	Time period when the digital certificate is valid. Values are: <ul style="list-style-type: none"> <li>• <b>Not before</b>—Start time when the digital certificate becomes valid.</li> <li>• <b>Not after</b>—End time when the digital certificate becomes invalid.</li> </ul>	All levels
<b>Public key algorithm</b>	Encryption algorithm used with the private key, such as <b>rsaEncryption (1024 bits)</b> .	All levels
<b>Public key verification status</b>	Public key verification status: <b>Failed</b> or <b>Passed</b> . The <b>detail</b> output also provides the verification hash.	All levels
<b>Signature algorithm</b>	Encryption algorithm that the CA used to sign the digital certificate, such as <b>sha1WithRSAEncryption</b> .	<b>detail</b>
<b>Fingerprint</b>	Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate.	<b>detail</b>
<b>Distribution CRL</b>	Distinguished name information and URL for the certificate revocation list (CRL) server.	<b>detail</b>
<b>Use for key</b>	Use of the public key, such as <b>Certificate signing</b> , <b>CRL signing</b> , <b>Digital signature</b> , or <b>Key encipherment</b> .	<b>detail</b>

## Sample Output

### show security pki local-certificate

```

user@switch> show security pki local-certificate
Certificate identifier: local-entrust2
Issued to: router2.juniper.net, Issued by: juniper

```

```
Validity:
  Not before: 2005 Nov 21st, 23:28:22 GMT
  Not after: 2008 Nov 21st, 23:58:22 GMT
Public key algorithm: rsaEncryption(1024 bits)
Public key verification status: Passed
```

#### show security pki local-certificate detail

```
user@switch> show security pki local-certificate detail
Certificate identifier: local-entrust3
Certificate version: 3
Serial number: 4355 94f9
Issuer:
  Organization: juniper, Country: us
Subject:
  Organization: juniper, Country: us, Common name: switch1.juniper.net
Alternate subject: switch1.juniper.net
Validity:
  Not before: 2005 Nov 21st, 23:33:58 GMT
  Not after: 2008 Nov 22nd, 00:03:58 GMT
Public key algorithm: rsaEncryption(1024 bits)
Public key verification status: Passed
fb:79:df:d4:a9:03:0f:d3:69:7e:c1:e4:27:35:9c:d9:b1:a2:47:78
d2:6d:f3:e5:f4:68:4f:b3:04:45:88:57:99:82:39:a6:51:9e:5f:42
23:3f:d7:6e:3d:a5:54:a9:b1:2d:6e:90:dd:12:8a:bf:ef:2b:20:50
ba:f0:da:d9:0c:ad:5e:d6:c6:98:3a:ae:3f:90:dd:94:78:c1:ea:2e
7c:f0:2d:d4:79:d4:cd:f0:52:df:5e:72:f2:e7:ae:66:f7:61:f4:bc
72:57:3e:6c:6d:d3:24:58:8b:f4:ef:da:2a:6a:fa:eb:98:f8:34:84
79:54:da:4f:d3:6f:52:1f
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
  61:3a:d0:b4:7a:16:9b:39:ba:81:3f:9d:ab:34:e5:c8:be:3b:a1:6d (sha1)
  60:a0:ff:58:05:4a:65:73:9d:74:3a:e1:83:6f:1b:c8 (md5)
Distribution CRL:
  C=us, O=juniper, CN=CRL1
  http://CA-1/CRL/juniper_us_crlfile.crl
Use for key: Digital signature
```

## show system services dhcp binding

<b>Syntax</b>	show system services dhcp binding <detail> <address>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(J Series routers only) Display Dynamic Host Configuration Protocol (DHCP) server client binding information.
<b>Options</b>	<p><b>none</b>—Display brief information about all active client bindings.</p> <p><b>detail</b>—(Optional) Display detailed information about all active client bindings.</p> <p><b>address</b>—(Optional) Display detailed client binding information for the specified IP address only.</p>
<b>Required Privilege Level</b>	view and system
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear system services dhcp binding on page 1790</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show system services dhcp binding on page 1804</a> <a href="#">show system services dhcp binding address on page 1804</a> <a href="#">show system services dhcp binding address detail on page 1804</a>
<b>Output Fields</b>	<a href="#">Table 160 on page 1803</a> describes the output fields for the <b>show system services dhcp binding</b> command. Output fields are listed in the approximate order in which they appear.

**Table 160: show system services dhcp binding Output Fields**

Field Name	Field Description	Level of Output
<b>Allocated address</b>	List of IP addresses the DHCP server has assigned to clients.	All levels
<b>MAC address</b>	Corresponding media access control (MAC) hardware address of the client.	All levels
<b>Client identifier</b>	( <b>address</b> option only) Client's unique identifier (represented by an ASCII string or hexadecimal digits). This identifier is used by the DHCP server to index its database of address bindings.	All levels
<b>Binding Type</b>	Type of binding assigned to the client. DHCP servers can assign a dynamic binding from a pool of IP addresses or a static binding to one or more specific IP addresses.	All levels
<b>Lease Expires at</b>	Time the lease expires or <b>never</b> for leases that do not expire.	All levels
<b>Lease Obtained at</b>	( <b>address</b> option only) Time the client obtained the lease from the DHCP server.	<b>detail</b>

Table 160: show system services dhcp binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	Status of the binding. Bindings can be active or expired.	<b>detail</b>
<b>Pool</b>	Address pool that contains the IP address assigned to the client.	<b>detail</b>
<b>Request received on</b>	Interface on which the DHCP message exchange occurs. The IP address pool is configured based on the interface's IP address. If a relay agent is used, its IP address is also displayed.	<b>detail</b>
<b>DHCP options</b>	User-defined options created for the DHCP server. If no options have been defined, this field is blank.	<b>detail</b>

## Sample Output

### show system services dhcp binding

```
user@host> show system services dhcp binding

Allocated address  MAC address      Binding Type  Lease expires at
192.168.1.2        00:a0:12:00:12:ab  static       never
192.168.1.3        00:a0:12:00:13:02  dynamic      2004-05-03 13:01:42 PDT
```

### show system services dhcp binding address

```
user@host> show system services dhcp binding 192.168.1.3

DHCP binding information:
Allocated address: 192.168.1.3
Mac address: 00:a0:12:00:12:ab
Client identifier
61 63 65 64 2d 30 30 3a 61 30 3a 31 32 3a 30 30aced-00:a0:12:00
3a 31 33 3a 30 32:13:02

Lease information:
  Binding Type dynamic
  Obtained at 2004-05-02 13:01:42 PDT
  Expires at 2004-05-03 13:01:42 PDT
```

### show system services dhcp binding address detail

```
user@host> show system services dhcp binding 192.168.1.3 detail

DHCP binding information:
Allocated address      192.168.1.3
MAC address 00:a0:12:00:12:ab
Pool                  192.168.1.0/24
Request received on fe-0/0/0, relayed by 192.168.4.254

Lease information:
  Type                DHCP
  Obtained at         2004-05-02 13:01:42 PDT
  Expires at          2004-05-03 13:01:42 PDT
  State active

DHCP options:
  Name: name-server, Value: { 6.6.6.6, 6.6.6.7 }
```

Name: domain-name, Value: mydomain.tld  
Code: 19, Type: flag, Value: off  
Code: 40, Type: string, Value: domain.tld  
Code: 32, Type: ip-address, Value: 3.3.3.33

## show system services dhcp conflict

<b>Syntax</b>	show system services dhcp conflict
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(J Series routers only and EX Series switches) Display Dynamic Host Configuration Protocol (DHCP) client-detected conflicts for IP addresses. When a conflict is detected, the DHCP server removes the address from the address pool.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view and system
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear system services dhcp conflict on page 1791</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show system services dhcp conflict on page 1806</a>
<b>Output Fields</b>	<a href="#">Table 161 on page 1806</a> describes the output fields for the <b>show system services dhcp conflict</b> command. Output fields are listed in the approximate order in which they appear.

**Table 161: show system services dhcp conflict Output Fields**

Field Name	Field Description
Detection time	Date and time the client detected the conflict.
Detection method	How the conflict was detected.
Address	IP address where the conflict occurs. The addresses in the conflicts list remain excluded from the pool until you use a <b>clear system services dhcp conflict</b> command to manually clear the list.

## Sample Output

### show system services dhcp conflict

```
user@host> show system services dhcp conflict
```

```

Detection time      Detection method  Address
2004-08-03 19:04:00 PDT  ARP             3.3.3.5
2004-08-04 04:23:12 PDT  Ping            4.4.4.8
2004-08-05 21:06:44 PDT  Client          3.3.3.10
```

## show system services dhcp global

<b>Syntax</b>	show system services dhcp global
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(J Series routers and EX Series switches only) Display Dynamic Host Configuration Protocol (DHCP) global configuration options. Global options apply to all scopes and clients served by the DHCP server. Global options are overridden if specified otherwise in scope or client options. Scope options apply to specific subnets or ranges of addresses. Client options apply to specific clients.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view and system
<b>List of Sample Output</b>	<a href="#">show system services dhcp global on page 1808</a>
<b>Output Fields</b>	<a href="#">Table 162 on page 1807</a> describes the output fields for the <b>show system services dhcp global</b> command. Output fields are listed in the approximate order in which they appear.

**Table 162: show system services dhcp global Output Fields**

Field Name	Field Description
<b>BOOTP lease length</b>	Length of lease time assigned to BOOTP clients.
<b>Default lease time</b>	Lease time assigned to clients that do not request a specific lease time.
<b>Minimum lease time</b>	Minimum time a client retains an IP address lease on the server.
<b>Maximum lease time</b>	Maximum time a client can retain an IP address lease on the server.
<b>DHCP options</b>	User-defined options created for the DHCP server. If no options have been defined, this field is blank.

## Sample Output

### show system services dhcp global

```
user@host> show system services dhcp global

Global settings:
  BOOTP lease length      infinite

DHCP lease times:
  Default lease time      1 hour
  Minimum lease time      2 hours
  Maximum lease time      infinite

DHCP options:
  Name: name-server, Value: { 6.6.6.6, 6.6.6.7 }
  Name: domain-name, Value: mydomain.tld
  Code: 19, Type: flag, Value: off
  Code: 40, Type: string, Value: domain.tld
  Code: 32, Type: ip-address, Value: 3.3.3.33
```



## show system services dhcp pool

<b>Syntax</b>	show system services dhcp pool <detail> <subnet-address>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(J Series routers and EX Series switches only) Display Dynamic Host Configuration Protocol (DHCP) server IP address pools.
<b>Options</b>	<b>none</b> —Display brief information about all IP address pools.  <b>detail</b> —(Optional) Display detailed information.  <b>subnet-address</b> —(Optional) Display information for the specified subnet address.
<b>Required Privilege Level</b>	view and system
<b>List of Sample Output</b>	<a href="#">show system services dhcp pool on page 1810</a> <a href="#">show system services dhcp pool subnet-address on page 1810</a> <a href="#">show system services dhcp pool subnet-address detail on page 1810</a>
<b>Output Fields</b>	<a href="#">Table 163 on page 1809</a> describes the output fields for the <b>show system services dhcp pool</b> command. Output fields are listed in the approximate order in which they appear.

Table 163: show system services dhcp pool Output Fields

Field Name	Field Description	Level of Output
Pool name	Subnet on which the IP address pool is defined.	None specified
Low address	Lowest address in the IP address pool.	None specified
High address	Highest address in the IP address pool.	None specified
Excluded addresses	Addresses excluded from the address pool.	None specified
Subnet	( <i>subnet-address</i> option only) Subnet to which the specified address pool belongs.	None specified
Address range	( <i>subnet-address</i> option only) Range of IP addresses in the address pool.	None specified
Addresses assigned	Number of IP addresses in the pool that are assigned to DHCP clients and the total number of IP addresses in the pool.	detail
Active	Number of assigned IP addresses in the pool that are active.	detail
Excluded	Number of assigned IP addresses in the pool that are excluded.	detail
Default lease time	Lease time assigned to clients that do not request a specific lease time.	detail

Table 163: show system services dhcp pool Output Fields (*continued*)

Field Name	Field Description	Level of Output
Minimum lease time	Minimum time a client can retain an IP address lease on the server.	detail
Maximum lease time	Maximum time a client can retain an IP address lease on the server.	detail
DHCP options	User-defined options created for the DHCP server. If no options have been defined, this field is blank.	detail

## Sample Output

### show system services dhcp pool

```
user@host> show system services dhcp pool

Pool name      Low address    High address    Excluded addresses
3.3.3.0/24     3.3.3.2       3.3.3.254      3.3.3.1
```

### show system services dhcp pool subnet-address

```
user@host> show system services dhcp pool 3.3.3.0/24

Pool information:
  Subnet                3.3.3.0/24
  Address range         3.3.3.2 - 3.3.3.254
  Addresses assigned     2/253
```

### show system services dhcp pool subnet-address detail

```
user@host> show system services dhcp pool 3.3.3.0/24 detail

Pool information:
  Subnet                3.3.3.0/24
  Address range         3.3.3.2 - 3.3.3.254
  Addresses assigned     2/253
  Active: 1, Excluded: 1

DHCP lease times:
  Default lease time     1 hour
  Minimum lease time     2 hours
  Maximum lease time     infinite

DHCP options:
  Name: name-server, Value: { 6.6.6.6, 6.6.6.7 }
  Name: domain-name, Value: mydomain.tld
  Name: router, Value: { 3.3.3.1 }
  Name: server-identifier, Value: 3.3.3.1
  Code: 19, Type: flag, Value: off
  Code: 40, Type: string, Value: domain.tld
  Code: 32, Type: ip-address, Value: 3.3.3.333.3.3.254 3.3.3.1
```

## show system services dhcp statistics

<b>Syntax</b>	show system services dhcp statistics
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(J Series routers and EX Series switches only) Display Dynamic Host Configuration Protocol (DHCP) server statistics.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view and system
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear system services dhcp statistics on page 1792</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show system services dhcp statistics on page 1812</a>
<b>Output Fields</b>	<a href="#">Table 164 on page 1811</a> describes the output fields for the <b>show system services dhcp statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 164: show system services dhcp statistics Output Fields**

Field Name	Field Description
<b>Default lease time</b>	Lease time assigned to clients that do not request a specific lease time.
<b>Minimum lease time</b>	Minimum time a client can retain an IP address lease on the server.
<b>Maximum lease time</b>	Maximum time a client can retain an IP address lease on the server.
<b>Packets dropped</b>	Total number of packets dropped and number of packets dropped because of: <ul style="list-style-type: none"> <li>• Invalid hardware address</li> <li>• Invalid opcode</li> <li>• Invalid server address</li> <li>• No available address</li> <li>• No interface match</li> <li>• No routing instance match</li> <li>• No valid local addresses</li> <li>• Packet too short</li> <li>• Read error</li> <li>• Send error</li> </ul>

Table 164: show system services dhcp statistics Output Fields (*continued*)

Field Name	Field Description
<b>Messages received</b>	<p>Number of the following message types sent from DHCP clients and received by the DHCP server:</p> <ul style="list-style-type: none"> <li>• BOOTREQUEST</li> <li>• DHCPDECLINE</li> <li>• DHCPDISCOVER</li> <li>• DHCPINFORM</li> <li>• DHCPRELEASE</li> <li>• DHCPREQUEST</li> </ul>
<b>Messages sent</b>	<p>Number of the following message types sent from the DHCP server to DHCP clients:</p> <ul style="list-style-type: none"> <li>• BOOTREPLY</li> <li>• DHCPACK</li> <li>• DHCPOFFER</li> <li>• DHCPNAK</li> </ul>

## Sample Output

### show system services dhcp statistics

```
user@host> show system services dhcp statistics
```

```
DHCP lease times:
  Default lease time      1 hour
  Minimum lease time     2 hours
  Maximum lease time     infinite
```

```
Packets dropped:
  Total                  0
  Bad hardware address   0
  Bad opcode             0
  Invalid server address 0
  No available addresses 0
  No interface match     0
  No routing instance match 0
  No valid local address 0
  Packet too short       0
  Read error             0
  Send error             0
```

```
Messages received:
  BOOTREQUEST           0
  DHCPDECLINE           0
  DHCPDISCOVER          0
  DHCPINFORM            0
  DHCPRELEASE           0
  DHCPREQUEST           0
```

```
Messages sent:
  BOOTREPLY             0
  DHCPACK               0
  DHCPOFFER             0
  DHCPNAK               0
```



## show system services service-deployment

<b>Syntax</b>	show system services service-deployment
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about a Session and Resource Control (SRC) client.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	system view
<b>List of Sample Output</b>	<a href="#">show system services service-deployment on page 1814</a>
<b>Output Fields</b>	<a href="#">Table 165 on page 1814</a> lists the output fields for the <b>show system services service-deployment</b> command. Output fields are listed in the approximate order in which they appear.

**Table 165: show system services service-deployment Output Fields**

Field Name	Field Description
PDT Keepalive settings	Configured PDT keepalive interval, in seconds.
Keepalives sent	Number of keepalives sent.
Notifications sent	Number of notifications sent.
Last update from peer	Time at which the last update from a peer was received.

## Sample Output

### show system services service-deployment

```

user@host> show system services service-deployment
Connected to 192.4.4.4 port 10288 since 2004-05-03 11:04:34 PDT Keepalive settings:
Interval 15 seconds Keepalives sent: 750 Notifications sent: 0 Last update from
peer: 00:00:06 ago

```

## ssh

**List of Syntax**    [Syntax on page 1815](#)  
                           [Syntax \(EX Series Switch and the QFX Series\) on page 1815](#)

**Syntax**    `ssh host`  
                   `<bypass-routing>`  
                   `<inet | inet6>`  
                   `<interface interface-name>`  
                   `<logical-system logical-system-name>`  
                   `<routing-instance routing-instance-name>`  
                   `<source address>`  
                   `<v1 | v2>`

**Syntax (EX Series Switch and the QFX Series)**    `ssh host`  
   `<bypass-routing>`  
   `<inet | inet6>`  
   `<interface interface-name>`  
   `<routing-instance routing-instance-name>`  
   `<source address>`  
   `<v1 | v2>`

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                   Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                   Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Use the SSH program to open a connection between a local router or switch and a remote system and execute commands on the remote system. You can issue the **ssh** command from the Junos OS CLI to log in to a remote system or from a remote system to log in to the local router or switch. When executing this command, you include one or more CLI commands by enclosing them in quotation marks and separating the commands with semicolons:

```
ssh address 'cli-command1 ; cli-command2 '
```

**Options**    **host**—Name or address of the remote system.

**bypass-routing**—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

**inet | inet6**—(Optional) Create an IPv4 or IPv6 connection, respectively.

**interface interface-name**—(Optional) Interface name for the SSH session. (This option does not work when **default-address-selection** is configured at the **[edit system]** hierarchy level, because this configuration uses the loopback interface as the source address for all locally generated IP packets.)

**logical-system logical-system-name**—(Optional) Name of a particular logical system for the SSH attempt.

**routing-instance** *routing-instance-name*—(Optional) Name of the routing instance for the SSH attempt.

**source address**—(Optional) Source address of the SSH connection.

**v1 | v2**—(Optional) Use SSH version 1 or 2, respectively, when connecting to a remote host.

**Additional Information** To configure an SSH (version 1) key for your user account, include the **authentication ssh-rsa** statement at the **[edit system login user *user-name*]** hierarchy level. To configure an SSH (version 2) key for your user account, include the **authentication dsa-rsa** statement at the **[edit system login user *user-name*]** hierarchy level.

You can limit the number of times a user can attempt to enter a password while logging in through SSH. To specify the number of times a user can attempt to enter a password to log in through SSH, include the **retry-options** statement at the **[edit system login]** hierarchy level. For details, see the .

**Required Privilege Level** network

**Related Documentation**

- *Configuring SSH Host Keys for Secure Copying of Data*

**List of Sample Output** [ssh on page 1816](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

ssh

```
user@switch> ssh cree
Host key not found from the list of known hosts.
Are you sure you want to continue connecting (yes/no)? yes

Host ?cree' added to the list of known hosts.
boojun@cree's password:
Last login: Sun Jun 21 10:43:42 1998 from junos-router
% ...
```



## telnet

**List of Syntax**   [Syntax on page 1817](#)  
                               [Syntax \(EX Series Switches\) on page 1817](#)

**Syntax**   `telnet host`  
                   `<8bit>`  
                   `<bypass-routing>`  
                   `<inet | inet6>`  
                   `<interface interface-name>`  
                   `<logical-system logical-system-name>`  
                   `<no-resolve>`  
                   `<port port-number>`  
                   `<routing-instance routing-instance-name>`  
                   `<source source-address>`

**Syntax (EX Series Switches)**   `telnet host`  
   `<8bit>`  
   `<bypass-routing>`  
   `<inet | inet6>`  
   `<interface interface-name>`  
   `<no-resolve>`  
   `<port port-number>`  
   `<routing-instance routing-instance-name>`  
   `<source source-address>`

**Release Information**   Command introduced before Junos OS Release 7.4.  
                                   Command introduced in Junos OS Release 9.0 for EX Series switches.

**Description**   Open a telnet session to a remote system. Type Ctrl+] to escape from the telnet session to the telnet command level, and then type **quit** to exit from telnet.

**Options**   *host*—Name or address of the remote system.

**8bit**—(Optional) Use an 8-bit data path.

**bypass-routing**—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

**inet | inet6**—(Optional) Open an IPv4 or IPv6 session, respectively.

**interface *interface-name***—(Optional) Interface name for the telnet session. (This option does not work when **default-address-selection** is configured at the **[edit system]** hierarchy level, because this configuration uses the loopback interface as the source address for all locally generated IP packets.)

**logical-system *logical-system-name***—(Optional) Name of a particular logical system for the telnet attempt.

**no-resolve**—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

**port** *port-number*—(Optional) Port number or service name on the remote system.

**routing-instance** *routing-instance-name*—(Optional) Name of the routing instance for the telnet attempt.

**source** *source-address*—(Optional) Source address of the telnet connection.

**Additional Information** You can limit the number of times a user can attempt to enter a password while logging in through telnet. To specify the number of times a user can attempt to enter a password to log in through telnet, include the **retry-options** statement at the [edit system login] hierarchy level. For details, see the *Junos OS Administration Library for Routing Devices*.

**Required Privilege Level** network

**List of Sample Output** [telnet on page 1818](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### telnet

```
user@host> telnet 192.154.1.254
Trying 192.154.169.254...
Connected to level5.company.net.
Escape character is '^]'.
ttypa
login:
```

## PART 9

# Access Control

- [Overview on page 1821](#)
- [Configuration on page 1843](#)
- [Administration on page 1993](#)



## CHAPTER 27

# Overview

- [Access Control Overview on page 1821](#)

## Access Control Overview

---

- [802.1X for EX Series Switches Overview on page 1821](#)
- [Understanding Authentication on EX Series Switches on page 1824](#)
- [Understanding Guest VLANs for 802.1X on EX Series Switches on page 1830](#)
- [Understanding 802.1X and RADIUS Accounting on EX Series Switches on page 1830](#)
- [Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831](#)
- [Understanding 802.1X and VoIP on EX Series Switches on page 1834](#)
- [Understanding 802.1X and VSAs on EX Series Switches on page 1837](#)
- [Understanding Dynamic VLANs for 802.1X on EX Series Switches on page 1837](#)
- [Understanding Server Fail Fallback and Authentication on EX Series Switches on page 1838](#)
- [Authentication Process Flow for EX Series Switches on page 1839](#)
- [Understanding Authentication Session Timeout on page 1841](#)

## 802.1X for EX Series Switches Overview

IEEE 802.1X provides network edge security, protecting Ethernet LANs from unauthorized user access.

### How 802.1X Authentication Works

---

802.1X authentication works by using an *Authenticator Port Access Entity* (the switch) to block all traffic to and from a supplicant (end device) at the port until the supplicant's credentials are presented and matched on the *Authentication server* (a RADIUS server). When authenticated, the switch stops blocking traffic and opens the port to the supplicant.

The end device is authenticated in either *single* mode, *single-secure* mode, or *multiple* mode:

- **single**—Authenticates only the first end device. All other end devices that connect later to the port are allowed full access without any further authentication. They effectively “piggyback” on the end devices' authentication.

- **single-secure**—Allows only one end device to connect to the port. No other end device is allowed to connect until the first logs out.
- **multiple**—Allows multiple end devices to connect to the port. Each end device will be authenticated individually.

Network access can be further defined using VLANs and firewall filters, which both act as filters to separate and match groups of end devices to the areas of the LAN they require. For example, you can configure VLANs to handle different categories of authentication failures depending upon:

- Whether or not the end device is 802.1X-enabled.
- Whether or not MAC RADIUS authentication has been configured on the switch interfaces to which the hosts are connected.
- Whether the RADIUS authentication server becomes unavailable or sends a RADIUS access-reject message. See [“Configuring Server Fail Fallback \(CLI Procedure\)” on page 1921](#).

---

## 802.1X Features Overview

---



**NOTE:** The 802.1X features available on the switches depend upon which switch you are using. See *EX Series Switch Software Features Overview* for a complete list of the Junos OS 802.1X features that are supported on specific switches.

---

802.1X features on Juniper Networks EX Series Ethernet Switches are:

- **Guest VLAN**—Provides limited access to a LAN, typically just to the Internet, for nonresponsive end devices that are not 802.1X-enabled when MAC RADIUS authentication has not been configured on the switch interfaces to which the hosts are connected. Also, a guest VLAN can be used to provide limited access to a LAN for guest users. Typically, the guest VLAN provides access just to the Internet and to other guests' end devices.
- **Server-reject VLAN**—Provides limited access to a LAN, typically just to the Internet, for responsive end devices that are 802.1X-enabled but that have sent the wrong credentials.
- **Server-fail VLAN**—Provides limited access to a LAN, typically just to the Internet, for 802.1X end devices during a RADIUS server timeout.
- **Dynamic VLAN**—Enables an end device, after authentication, to be a member of a VLAN dynamically.
- **Private VLAN**—Enables configuration of 802.1X authentication on interfaces that are members of private VLANs (PVLANS).
- **Dynamic changes to a user session**—Allows the switch administrator to terminate an already authenticated session. This feature is based on support of the RADIUS Disconnect Message defined in RFC 3576.

- Support for VoIP—Supports IP telephones. If the phone is 802.1X-enabled, it is authenticated like any other supplicant. If the phone is not 802.1X-enabled, but has another 802.1X-compatible device connected to its data port, that device is authenticated, and then VoIP traffic can flow to and from the phone (providing that the interface is configured in single mode and not in single-secure mode).



**NOTE:** Configuring a VoIP VLAN on private VLAN (PVLAN) interfaces is not supported.

- RADIUS accounting—Sends accounting information to the RADIUS accounting server. Accounting information is sent to the server whenever a subscriber logs in or logs out and whenever a subscriber activates or deactivates a subscription.
- Vendor Specific Attributes (VSAs)—Supports the **Juniper-Switching-Filter** attribute on the RADIUS authentication server that can be used further define a supplicant's access during the 802.1X authentication process. Centrally configuring VSAs on the authentication server does away with the need to configure these same attributes in the form of firewall filters on every switch in the LAN to which the supplicant may connect to the LAN. This feature is based on RLI 4583, AAA RADIUS BRAS VSA Support.

### Supported Features Related to 802.1X Authentication

802.1X does not replace other security technologies. 802.1X works together with port security features, such as DHCP snooping, dynamic ARP inspection (DAI), and MAC limiting, to guard against spoofing.

Supported features related to authentication include:

- Static MAC bypass—Provides a bypass mechanism to authenticate devices that are not 802.1X-enabled (such as printers). Static MAC bypass connects these devices to 802.1X-enabled ports, bypassing 802.1X authentication.
- MAC RADIUS authentication—Provides a means to enable or disable MAC authentication independently of whether 802.1X authentication is enabled.



**NOTE:** You cannot configure 802.1X authentication on redundant trunk groups (RTGs).

#### Related Documentation

- [Understanding Authentication on EX Series Switches on page 1824](#)
- [Understanding 802.1X and VoIP on EX Series Switches on page 1834](#)
- [Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831](#)
- [Understanding 802.1X and RADIUS Accounting on EX Series Switches on page 1830](#)
- [Understanding Guest VLANs for 802.1X on EX Series Switches on page 1830](#)
- [Understanding 802.1X and VSAs on EX Series Switches on page 1837](#)

- [Understanding Server Fail Fallback and Authentication on EX Series Switches on page 1838](#)

## Understanding Authentication on EX Series Switches

You can control access to your network through a Juniper Networks EX Series Ethernet Switch using several different authentication methods—802.1X, MAC RADIUS, or captive portal. Authentication prevents unauthorized devices and users from gaining access to your LAN. For 802.1X and MAC RADIUS authentication, end devices must be authenticated before they receive an IP address from a DHCP server. For captive portal authentication, the switch allows the end devices to get an IP address and allows forwarding of DHCP, DNS, and ARP packets.

You can allow end devices to access the network without authentication by including the MAC address of the end device in the static MAC bypass list or, for captive portal, by including the MAC address of the end device in the authentication whitelist.

You can configure 802.1X, MAC RADIUS, and captive portal on the same interface and in any combination, except that you cannot configure MAC RADIUS and captive portal on an interface without also configuring 802.1X. If you configure multiple authentication methods on a single interface, the switch falls back to another method if the first method is unsuccessful. For a description of the process flow when multiple authentication methods are configured on an interface, see [“Authentication Process Flow for EX Series Switches” on page 1839](#).

This topic covers:

- [Sample Basic Authentication Topology on page 1824](#)
- [802.1X Authentication on page 1825](#)
- [MAC RADIUS Authentication on page 1827](#)
- [Captive Portal Authentication on page 1827](#)
- [Static MAC Bypass of Authentication on page 1828](#)
- [Fallback of Authentication Methods on page 1828](#)


---

### Sample Basic Authentication Topology

[Figure 6 on page 1825](#) illustrates a basic deployment topology for authentication on an EX Series switch:



Figure 6: Example Authentication Topology



### 802.1X Authentication

802.1X is an IEEE standard for port-based network access control (PNAC). It provides an authentication mechanism to allow devices to access a LAN. The 802.1X authentication feature on an EX Series switch is based upon the IEEE 802.1D standard *Port-Based Network Access Control*.

The communication protocol between the end device and the switch is Extensible Authentication Protocol Over LAN (EAPoL). EAPoL is a version of EAP designed to work with Ethernet networks. The communication protocol between the authentication server and the switch is RADIUS.

During the authentication process, the switch completes multiple message exchanges between the end device and the authentication server. While 802.1X authentication is in process, only 802.1X traffic is allowed. Other traffic, such as DHCP and HTTP, is blocked at the data link layer.



**NOTE:** You can configure both the maximum number of times an EAPoL request packet is retransmitted and the timeout period between attempts. For information, see [“Configuring 802.1X Interface Settings \(CLI Procedure\)” on page 1908](#).

An 802.1X authentication configuration for a LAN contains three basic components:

- *Supplicant* (also called end device)—Supplicant is the IEEE term for an end device that requests to join the network. The end device can be responsive or nonresponsive. A responsive end device is 802.1X-enabled and provides authentication credentials—specifically, a username and password for EAP MD5 or a username and client certificates for EAP-TLS, EAP-TTLS, and EAP-PEAP.

You can configure a server-reject VLAN to provide limited LAN access for responsive end devices that are 802.1X-enabled but that have sent the wrong credentials. A server-reject VLAN can provide a remedial connection, typically just to the Internet, for these devices. See [“Example: Configuring Fallback Options on EX Series Switches for EAP-TTLS Authentication and Odyssey Access Clients” on page 1903](#) for additional information.



**NOTE:** If the end device that is authenticated using the server-reject VLAN is an IP phone, voice traffic is not allowed.

A nonresponsive end device is not 802.1X-enabled, but it can be authenticated through MAC RADIUS authentication.

- *Authenticator port access entity*—The IEEE term for the authenticator. The EX Series switch is the authenticator, and it controls access by blocking all traffic to and from end devices until they are authenticated.
- *Authentication server*—The authentication server contains the backend database that makes authentication decisions. It contains credential information for each end device that is allowed to connect to the network. The authenticator forwards credentials supplied by the end device to the authentication server. If the credentials forwarded by the authenticator match the credentials in the authentication server database, access is granted. If the credentials forwarded do not match, access is denied. The EX Series switches support RADIUS authentication servers.



**NOTE:** You cannot configure 802.1X authentication on redundant trunk groups (RTGs). For more information on RTGs, see [“Understanding Redundant Trunk Links” on page 2276](#).

---

## MAC RADIUS Authentication

---

You can configure MAC RADIUS authentication on interfaces that are connected to end devices that are not 802.1X-enabled but that you want to allow to access the LAN.

The EAP method supported for MAC RADIUS authentication on EX Series switches is EAP-MD5.

If both 802.1X-enabled end devices and end devices that are not 802.1X-enabled connect to an interface, you can configure both 802.1X and MAC RADIUS authentication methods on the interface. In this case, the switch first attempts to authenticate using 802.1X, and if that method fails, it attempts to authenticate the end device using MAC RADIUS authentication.

If you know that only non-802.1X-enabled end devices connect on that interface, you can eliminate the delay that occurs while the switch determines that the end device is non-802.1X-enabled by configuring the **mac-radius restrict** option. When this option is configured, the switch does not attempt to authenticate the end device through 802.1X but instead immediately sends a request to the RADIUS server for authentication of the MAC address of the end device. If the MAC address of an end device is configured as permitted on the RADIUS server, the switch opens LAN access to the end device on the interface to which it is connected.

This option is useful when no other 802.1X authentication methods, such as guest VLAN, are needed on the interface. When you configure **mac-radius restrict** on an interface to eliminate this delay, the switch drops all 802.1X packets.

---

## Captive Portal Authentication

---

Captive portal authentication (hereafter referred to as captive portal) allows you to authenticate users on EX Series switches by redirecting Web browser requests to a login page that requires users to input a username and password before they are allowed access to the network. Captive portal controls network access by requiring users to provide information that is authenticated against a RADIUS server database using EAP-MD5. You can also use captive portal to display an acceptable-use policy to users before they access your network.

Juniper Networks Junos operating system (Junos OS) for EX Series switches provides a template that allows you to easily design and modify the look of the captive portal login page. You enable specific interfaces for captive portal. The first time an end device connected to a captive portal interface attempts to access a web page, the switch presents the captive portal login page. Upon successful authentication, the user is allowed access to the network and to continue to the original page requested.



**NOTE:** If Hypertext Transfer Protocol over Secure Sockets Layer (HTTPS) is enabled, Hypertext Transfer Protocol (HTTP) requests are redirected to an HTTPS connection for the captive portal authentication process. After authentication, the end device is returned to the HTTP connection.

---

If there are end devices that are not HTTP-enabled connected to the captive portal interface, you can allow them to bypass captive portal authentication by adding their MAC addresses to an authentication whitelist.

When the user is authenticated by the RADIUS server, any per-user policies (attributes) associated with that user are also sent to the switch.

Captive portal on EX Series switches has the following limitations:

- The captive portal interface must be configured for **family ethernet-switching** and set to port mode **access**.
- Captive portal does not support dynamic assignment of VLANs downloaded from the RADIUS server.
- If the user is idle for more than about 5 minutes and there is no traffic passed, the user must log back in to the captive portal.

### Static MAC Bypass of Authentication

---

You can allow end devices to access the LAN without authentication on a RADIUS server by including their MAC addresses in the static MAC bypass list (also known as the exclusion list).

You might choose to include a device in the bypass list to:

- Allow non-802.1X-enabled devices access to the LAN.
- Eliminate the delay that occurs while the switch determines that a connected device is a non-802.1X-enabled host.

When you configure static MAC on the switch, the MAC address of the end device is first checked in a local database (a user-configured list of MAC addresses). If a match is found, the end device is successfully authenticated and the interface is opened up for it. No further authentication is done for that end device. If a match is not found and 802.1X authentication is enabled on the switch, the switch attempts to authenticate the end device through the RADIUS server.

For each MAC address, you can also configure the VLAN to which the end device is moved or the interfaces on which the host connects.



**CAUTION:** When you clear the learned MAC addresses from an interface using the `clear dot1x interface` command, all MAC addresses are cleared, including those in the static MAC bypass list.

---

### Fallback of Authentication Methods

---

You can configure multiple authentication methods on a single interface to enable fallback to another method if one method fails.

If an interface is configured in multiple supplicant mode, all end devices connecting through the interface must use either captive portal or a combination of 802.1X and MAC

RADIUS, captive portal cannot be mixed with 802.1X or MAC RADIUS. Therefore, if there is already an end device on the interface that was authenticated through 802.1X or MAC RADIUS authentication, then additional end devices authenticating do not fall back to captive portal. If only 802.1X authentication or MAC RADIUS authentication is configured, some end devices can be authenticated using 802.1X and others can still be authenticated using MAC RADIUS.

Fallback of authentication methods occurs in the following order:

1. 802.1X authentication—If 802.1X is configured on the interface, the switch sends EAPoL requests to the end device and attempts to authenticate the end device through 802.1X authentication. If the end device does not respond to the EAP requests, the switch checks whether MAC RADIUS authentication is configured on the interface.
2. MAC RADIUS authentication—If MAC RADIUS authentication is configured on the interface, the switch sends the MAC RADIUS address of the end device to the authentication server. If MAC RADIUS authentication is not configured, the switch checks whether captive portal is configured on the interface.
3. Captive portal authentication—If captive portal is configured on the interface, the switch attempts to authenticate using this method after attempting any other configured authentication methods. If an end device is authenticated on the interface using captive portal, this becomes the active authentication method on the interface. When captive portal is the active authentication method, the switch falls back to 802.1X authentication if there are no sessions in the authenticated state and if the interface receives an EAP packet.

#### Related Documentation

- [802.1X for EX Series Switches Overview on page 1821](#)
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)
- [Configuring MAC RADIUS Authentication \(CLI Procedure\) on page 1923](#)
- [Configuring Captive Portal Authentication \(CLI Procedure\)](#)
- [Configuring Static MAC Bypass of Authentication \(CLI Procedure\) on page 1924](#)
- [Controlling Authentication Session Timeouts \(CLI Procedure\) on page 1930](#)
- [Authentication Process Flow for EX Series Switches on page 1839](#)

## Understanding Guest VLANs for 802.1X on EX Series Switches

Guest VLANs can be configured on switches that are using 802.1X authentication to provide limited access—typically only to the Internet—for:

- Corporate guests
- End devices that are not 802.1X-enabled
- Nonresponsive end devices when MAC RADIUS authentication has not been configured on the switch interfaces to which the hosts are connected

A guest VLAN is not used for supplicants sending incorrect credentials. Those supplicants are directed to the server-reject VLAN instead.

For end devices that are not 802.1X-enabled, a guest VLAN can allow limited access to a server from which the non-802.1X-enabled end device can download the supplicant software and attempt authentication again.

A guest VLAN is not used when MAC RADIUS authentication has been configured on the switch interfaces to which the hosts are connected. Some end devices, such as a printer, cannot be enabled for 802.1X. The hosts for such devices should be connected to switch interfaces that are configured for MAC RADIUS authentication. See [“Configuring MAC RADIUS Authentication \(CLI Procedure\)” on page 1923](#).

### Related Documentation

- [Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an EX Series Switch on page 1847](#)
- [Understanding Dynamic VLANs for 802.1X on EX Series Switches on page 1837](#)
- [Understanding Authentication on EX Series Switches on page 1824](#)

## Understanding 802.1X and RADIUS Accounting on EX Series Switches

Juniper Networks EX Series Ethernet Switches support IETF RFC 2866, *RADIUS Accounting*. Configuring RADIUS accounting on an EX Series switch permits statistical data about users logging onto or off a LAN to be collected and sent to a RADIUS accounting server. The statistical data gathered can be used for general network monitoring, to analyze and track usage patterns, or to bill a user based upon the amount of time or type of services accessed.

To configure RADIUS accounting, specify one or more RADIUS accounting servers to receive the statistical data from the switch, and select the type of accounting data to be collected.

The RADIUS accounting server you specify can be the same server used for RADIUS authentication, or it can be a separate RADIUS server. You can specify a list of RADIUS accounting servers. In the event that the primary server (the first one configured) is unavailable, each RADIUS server in the list is tried in the order in which they are configured in the Juniper Networks Junos operating system (Junos OS).

The RADIUS accounting process between a switch and a RADIUS server works like this:

1. A RADIUS accounting server listens for User Datagram Protocol (UDP) packets on a specific port. For example, on FreeRADIUS, the default port is 1813.
2. The switch forwards an *accounting-request* packet containing an event record to the accounting server. For example, a supplicant is authenticated through 802.1X authentication and connected to the LAN. The event record associated with this supplicant contains an *Acct-Status-Type* attribute whose value indicates the beginning of user service for this supplicant. When the supplicant's session ends, the accounting request will contain an *Acct-Status-Type* attribute value indicating the end of user service. The RADIUS accounting server records this as a stop-accounting record containing session information and the length of the session.
3. The RADIUS accounting server logs these events as start-accounting or stop-accounting records. The records are in a file. On FreeRADIUS, the file name is the server's address; for example, 122.69.1.250.
4. The accounting server sends an *accounting-response* packet back to the switch confirming it has received the accounting request.
5. If the switch does not receive a response from the server, it continues to send accounting requests until an accounting response is returned from the accounting server.

The statistics collected through this process can be displayed from the RADIUS server; to see those statistics, the user accesses the log file configured to receive them.

**Related  
Documentation**

- [Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch on page 1843](#)
- [802.1X for EX Series Switches Overview on page 1821](#)
- [Configuring 802.1X RADIUS Accounting \(CLI Procedure\) on page 1909](#)

## Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches

Juniper Networks EX Series Ethernet Switches use Link Layer Discovery Protocol (LLDP) and Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) to learn and distribute device information on network links. The information allows the switch to quickly identify a variety of devices, resulting in a LAN that interoperates smoothly and efficiently.

LLDP-capable devices transmit information in type, length, and value (TLV) messages to neighbor devices. Device information can include information such as chassis and port identification and system name and system capabilities. The TLVs leverage this information from parameters that have already been configured in the Juniper Networks Junos operating system (Junos OS).

LLDP-MED goes one step further than LLDP, exchanging IP-telephony messages between the switch and the IP telephone.



**NOTE:** If your IP telephone is configured for voice over IP (VoIP), the switch automatically detects the configuration and assigns the telephone to the voice VLAN. The implementation of a voice VLAN on an IP telephone is vendor-specific. Consult the documentation that came with your IP telephone for instructions on configuring a voice VLAN. For example, on an Avaya phone, you can ensure that the phone gets the correct VoIP VLAN ID even in the absence of LLDP-MED by enabling DHCP option 176.

LLDP and LLDP-MED also provide PoE power management capabilities. LLDP power negotiation allows the switch to manage PoE power by negotiating with LLDP-enabled powered devices to dynamically allocate PoE power as needed. LLDP power priority allows an LLDP-enabled powered device to set the PoE power priority on the switch interface to which it connects.

The switch also uses these protocols to ensure that voice traffic gets tagged and prioritized with the correct values at the source itself. For example, 802.1p CoS and 802.1Q tag information can be sent to the IP telephone.

EX Series switches support the following basic TLVs:

- **Chassis Identifier**—The MAC address associated with the local system.



**NOTE:** The Chassis ID TLV has a subtype for Network Address Family. LLDP frames are validated only if this subtype has a value of 1 (IPv4) or 2 (IPv6). For any other value, the transmitting device is detected by LLDP as a neighbor and displayed in the output of the "show lldp neighbors" command, but is not assigned to the VLAN.

- **Port Identifier**—The port identification for the specified port in the local system.
- **Port Description**—Textual description of the interface or the logical unit. The description for the logical unit is used, if available; otherwise, the Port Description TLV will contain the description configured on the physical interface. For example, LAG member interfaces do not contain a logical unit, so only the description configured on the physical interface can be used.
- **System Name**—The user-configured name of the local system. The system name can be a maximum of 256 characters.
- **System Description**—The system description containing information about the software and current image running on the system. This information is not configurable, but taken from the software.
- **System Capabilities**—The primary function performed by the system. The capabilities that system supports; for example, bridge or router. This information is not configurable, but based on the model of the product.
- **Management Address**—The IPv4 or IPv6 management address of the local system.

EX Series switches support the following 802.3 TLVs:



- **Power via MDI**—A TLV that advertises MDI power support, PSE power pair, and power class information.
- **MAC/PHY Configuration Status**—A TLV that advertises information about the physical interface, such as autonegotiation status and support and MAU type. The information is not configurable, but based on the physical interface structure.



**NOTE:** The MAC/PHY Configuration Status TLV has a subtype for the PMD Auto-Negotiation Advertised Capability field. This field will contain a value of **other** or **unknown** if the LLDP packet was transmitted from a 10-gigabit SFP+ port.

- **Link Aggregation**—A TLV that advertises if the port is aggregated and its aggregated port ID.
- **Maximum Frame Size**—A TLV that advertises the Maximum Transmission Unit (MTU) of the interface sending LLDP frames.
- **Port Vlan**—A TLV that advertises the VLAN name configured on the interface.

EX Series switches support the following LLDP-MED TLVs:

- **LLDP MED Capabilities**—A TLV that advertises the primary function of the port. The capabilities values range 0 through 15:
  - **0**— Capabilities
  - **1**— Network Policy
  - **2**— Location Identification
  - **3**— Extended Power via MDI-PSE
  - **4**— Inventory
  - **5–15**— Reserved
- **LLDP-MED Device Class Values:**
  - **0**— Class not defined.
  - **1**— Class 1 Device.
  - **2**— Class 2 Device.
  - **3**— Class 3 Device.
  - **4**— Network Connectivity Device
  - **5–255**— Reserved.
- **Network Policy**—A TLV that advertises the port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier, application types, such as voice or streaming video, 802.1Q VLAN tagging, and 802.1p priority bits and Diffserv code points.

- **Endpoint Location**— A TLV that advertises the physical location of the endpoint.
- **Extended Power via MDI**— A TLV that advertises the power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port.

**Related  
Documentation**

- *Understanding Layer 2 Protocol Tunneling on EX Series Switches*
- *Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch*
- [Configuring LLDP \(CLI Procedure\) on page 1913](#)
- [Configuring LLDP-MED \(CLI Procedure\) on page 1917](#)
- [Understanding PoE on EX Series Switches on page 4423](#)

## Understanding 802.1X and VoIP on EX Series Switches


When you use Voice over IP (VoIP), you can connect IP telephones to the switch and configure IEEE 802.1X authentication for 802.1X-compatible IP telephones. The 802.1X authentication provides network edge security, protecting Ethernet LANs from unauthorized user access.

VoIP is a protocol used for the transmission of voice through packet-switched networks. VoIP transmits voice calls using a network connection instead of an analog phone line.

When VoIP is used with 802.1X, the RADIUS server authenticates the phone, and Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) provides the class-of-service (CoS) parameters to the phone.


You can configure 802.1X authentication to work with VoIP in multiple supplicant or single supplicant mode. In *multiple-supplicant* mode, the 802.1X process allows multiple supplicants to connect to the interface. Each supplicant will be authenticated individually. For an example of a VoIP multiple supplicant topology, see [Figure 7 on page 1835](#).

Figure 7: VoIP Multiple Supplicant Topology



If an 802.1X-compatible IP telephone does not have an 802.1X host but has another 802.1X-compatible device connected to its data port, you can connect the phone to an interface in single-supplicant mode. In *single-supplicant* mode, the 802.1X process authenticates only the first supplicant. All other supplicants who connect later to the interface are allowed full access without any further authentication. They effectively “piggyback” on the first supplicant’s authentication. For an example of a VoIP single supplicant topology, see [Figure 8 on page 1836](#).

Figure 8: VoIP Single Supplicant Topology



If an IP telephone does not support 802.1X, you can configure VoIP to bypass 802.1X and LLDP-MED and have the packets forwarded to a VoIP VLAN,

**Related Documentation**

- [Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831](#)
- *Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch*
- *Example: Configuring VoIP on an EX Series Switch Without Including 802.1X Authentication*
- *Example: Configuring VoIP on an EX Series Switch Without Including LLDP-MED Support*

## Understanding 802.1X and VSAs on EX Series Switches

Juniper Networks EX Series Ethernet Switches support the configuration of RADIUS server attributes specific to Juniper Networks. These attributes are known as vendor-specific attributes (VSAs) and are described in RFC 2138, *Remote Authentication Dial In User Service* (RADIUS). Through VSAs, you can configure port-filtering attributes on the RADIUS server. VSAs are clear text fields sent from the RADIUS server to the switch as a result of the 802.1X authentication success or failure. The 802.1X authentication prevents unauthorized user access by blocking a supplicant at the port until the supplicant is authenticated by the RADIUS server. The VSA attributes are interpreted by the switch during authentication, and the switch takes appropriate actions. Implementing port-filtering attributes with 802.1X authentication on the RADIUS server provides a central location for controlling LAN access for supplicants.

These port-filtering attributes specific to Juniper Networks are encapsulated in a RADIUS server VSA with the vendor ID set to the Juniper Networks ID number, 2636.

As well as configuring port-filtering attributes through VSAs, you can apply a port firewall filter that has already been configured on the switch directly to the RADIUS server. Like port-filtering attributes, the filter is applied during the 802.1X authentication process, and its actions are applied at the switch port. Adding a port firewall filter to a RADIUS server eliminates the need to add the filter to multiple ports and switches. For more information, see [“Example: Applying a Firewall Filter to 802.1X-Authenticated Supplicants Using RADIUS Server Attributes on an EX Series Switch”](#) on page 1887.

VSAs are only supported for 802.1X single-supplicant configurations and multiple-supplicant configurations.

### Related Documentation

- [Understanding Authentication on EX Series Switches on page 1824](#)
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 1910](#)
- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
- [VSA Match Conditions and Actions on page 1918](#)

## Understanding Dynamic VLANs for 802.1X on EX Series Switches

Dynamic VLANs, in conjunction with the 802.1X authentication process, provide secure access to the LAN for end devices belonging to different VLANs on a single port.

When this feature is configured on the RADIUS server, an end device or user authenticating on the RADIUS server is assigned to the VLAN configured for it. The end device or user becomes a member of a VLAN dynamically after successful 802.1X authentication. For information on configuring dynamic VLANs on your RADIUS server, see the documentation for your RADIUS server.

Successful authentication requires that the VLAN ID or VLAN name exist on the switch and match the VLAN ID or VLAN name sent by the RADIUS server during authentication.

If neither exists, the end device is unauthenticated. If a guest VLAN is established, the unauthenticated end device is automatically moved to the guest VLAN.

**Related  
Documentation**

- [Example: Configuring MAC RADIUS Authentication on an EX Series Switch on page 1881](#)
- [Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an EX Series Switch on page 1847](#)
- [Understanding Guest VLANs for 802.1X on EX Series Switches on page 1830](#)

## Understanding Server Fail Fallback and Authentication on EX Series Switches

Server fail fallback allows you to specify how end devices connected to the switch are supported if the RADIUS authentication server becomes unavailable or sends a RADIUS access-reject message.

Juniper Networks EX Series Ethernet Switches use authentication to implement access control in an enterprise network. If 802.1X, MAC RADIUS, or captive portal authentication are configured on the interface, end devices are evaluated at the initial connection by an authentication (RADIUS) server. If the end device is configured on the authentication server, the device is granted access to the LAN and the EX Series switch opens the interface to permit access.

A RADIUS server timeout occurs if no RADIUS authentication servers are reachable when an end device logs in and attempts to access the LAN. Server fail fallback allows you to specify one of four actions to be taken toward end devices awaiting authentication when the server is timed out:

- *Permit* authentication, allowing traffic to flow from the end device through the interface as if the end device were successfully authenticated by the RADIUS server.
- *Deny* authentication, preventing traffic from flowing from the end device through the interface. This is the default.
- *Move* the end device to a specified VLAN. (The VLAN must already exist on the switch.)
- *Sustain* authenticated end devices that already have LAN access and *deny* unauthenticated end devices. If the RADIUS servers time out during reauthentication, previously authenticated end devices are reauthenticated and new users are denied LAN access.

Server fail fallback is triggered most often during reauthentication when the already configured and in-use RADIUS server becomes inaccessible. However, server fail fallback can also be triggered by an end device's first attempt at authentication through the RADIUS server.

Server fail fallback allows you to specify that an end device be moved to a specified VLAN if the switch receives a RADIUS access-reject message. The configured VLAN name overrides any attributes sent by the server.

**Related  
Documentation**

- [802.1X for EX Series Switches Overview on page 1821](#)
- [Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to an EX Series Switch on page 1875](#)

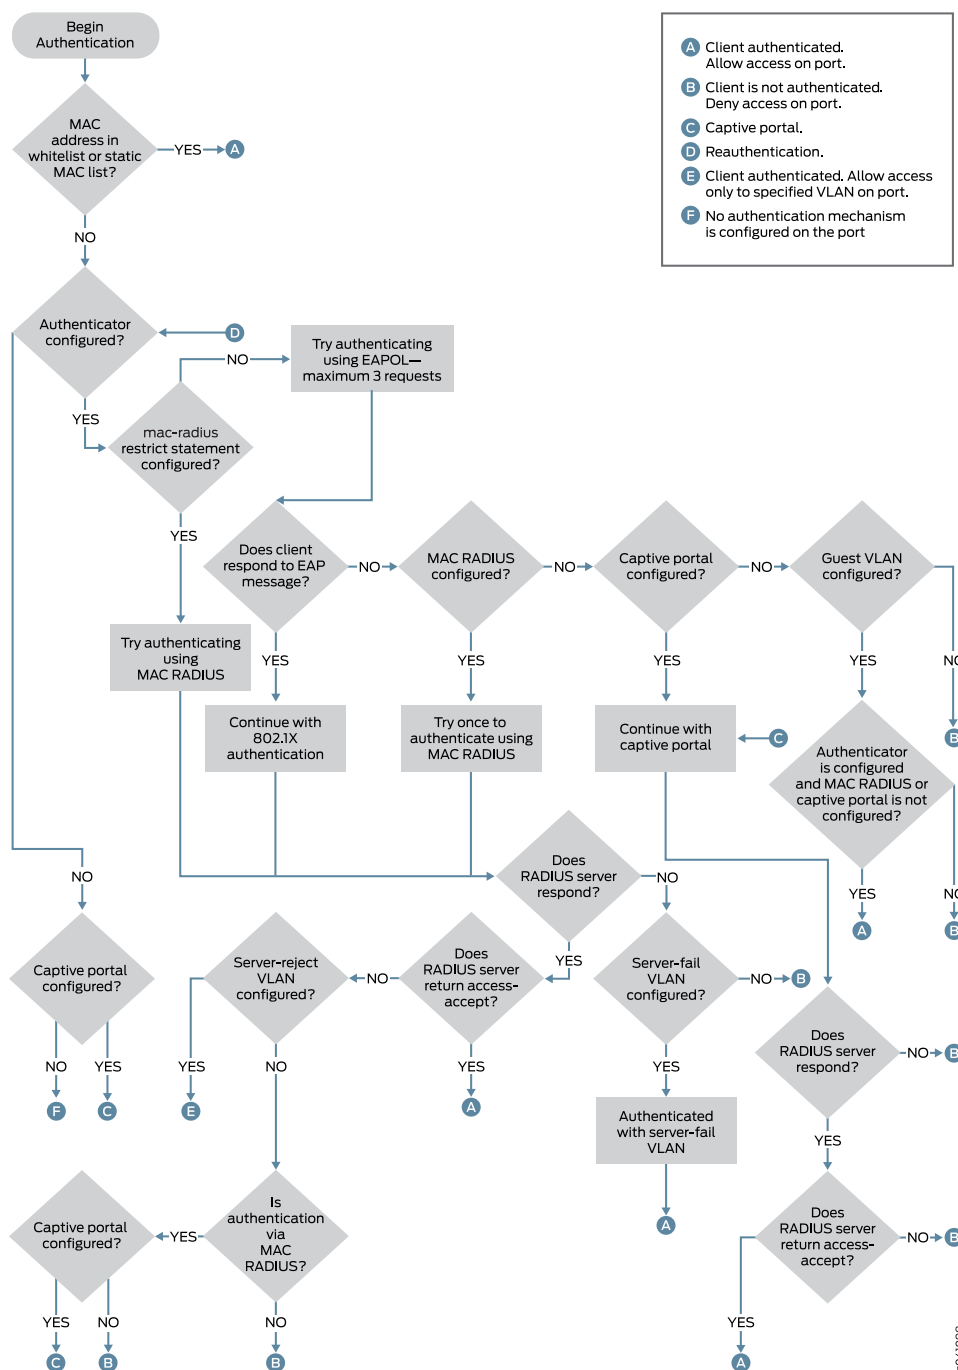
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Configuring Server Fail Fallback \(CLI Procedure\) on page 1921](#)
- [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)

## Authentication Process Flow for EX Series Switches

You can control access to your network through an EX Series switch by using several different authentication methods—including 802.1X, MAC RADIUS, or captive portal.

[Figure 9 on page 1840](#) illustrates the authentication process:

### Figure 9: Authentication Process Flow for an EX Series Switch



## Related Documentation

- [Understanding Authentication on EX Series Switches on page 1824](#)
- [Understanding Server Fail Fallback and Authentication on EX Series Switches on page 1838](#)



- [Understanding Guest VLANs for 802.1X on EX Series Switches on page 1830](#)
- [Understanding Dynamic VLANs for 802.1X on EX Series Switches on page 1837](#)
- *Example: Setting Up Captive Portal Authentication on an EX Series Switch*

## Understanding Authentication Session Timeout

You can specify authentication session timeout values for captive portal authentication sessions and 802.1X and MAC RADIUS authentication sessions.

For captive portal authentication, the length of the session depends on the value configured for the **session-expiry** statement. The remainder of this topic pertains only to 802.1X and MAC RADIUS authentication sessions.

For 802.1X and MAC RADIUS authentication sessions, the timeout of the session depends on the value of **reauthentication interval** for **dot1x authentication**. The authentication session might also end when the MAC table aging time expires because, unless you configure it not to, the session is removed from the authentication session table when the MAC address is removed from the Ethernet switching table.

Information about each 802.1X and MAC RADIUS authentication session—including the associated interfaces and VLANs for each MAC address that is authenticated by 802.1X authentication or MAC RADIUS authentication—is stored in the authentication session table. The authentication session table is tied to the Ethernet switching table (also called the MAC table). Each time the switch detects traffic from a MAC address, it updates the timestamp for that network node in the Ethernet switching table. A timer on the switch periodically checks the timestamp and if its value exceeds the user-configured **mac-table-aging-time** value, the switch removes the MAC address from the Ethernet switching table. When a MAC address ages out of the Ethernet switching table, the entry for that MAC address is also removed from the authentication database, with the result that the session ends.

You can control variables affecting timeout of authentication sessions in the following ways:

- Set the authentication session timeout on all interfaces or on selected interfaces using the **reauthentication** statement.
- Disassociate the authentication session table from the Ethernet switching table using the **no-mac-table-binding** statement. This setting prevents the termination of the authentication session when the associated MAC address ages out of the Ethernet switching table.

### Related Documentation

- [Understanding Authentication on EX Series Switches on page 1824](#)
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Controlling Authentication Session Timeouts \(CLI Procedure\) on page 1930](#)
- [Configuring MAC Table Aging \(CLI Procedure\)](#)



## CHAPTER 28

# Configuration

- [Configuration Examples on page 1843](#)
- [Configuration Tasks on page 1907](#)
- [Configuration Statements on page 1931](#)

### Configuration Examples

---

- [Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch on page 1843](#)
- [Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an EX Series Switch on page 1847](#)
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch on page 1858](#)
- [Example: Configuring VoIP on an EX Series Switch Without Including 802.1X Authentication on page 1866](#)
- [Example: Configuring Static MAC Bypass of Authentication on an EX Series Switch on page 1872](#)
- [Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to an EX Series Switch on page 1875](#)
- [Example: Configuring MAC RADIUS Authentication on an EX Series Switch on page 1881](#)
- [Example: Applying a Firewall Filter to 802.1X-Authenticated Supplicants Using RADIUS Server Attributes on an EX Series Switch on page 1887](#)
- [Example: Applying Firewall Filters to Multiple Supplicants on Interfaces Enabled for 802.1X or MAC RADIUS Authentication on page 1893](#)
- [Example: Setting Up Captive Portal Authentication on an EX Series Switch on page 1898](#)
- [Example: Configuring Fallback Options on EX Series Switches for EAP-TTLS Authentication and Odyssey Access Clients on page 1903](#)

### Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch

802.1X is the IEEE standard for Port-Based Network Access Control (PNAC). You use 802.1X to control network access. Only users and devices providing credentials that have been verified against a user database are allowed access to the network. You can use a

RADIUS server as the user database for 802.1X authentication, as well as for MAC RADIUS authentication.

This example describes how to connect a RADIUS server to an EX Series switch, and configure it for 802.1X:

- [Requirements on page 1844](#)
- [Overview and Topology on page 1844](#)
- [Configuration on page 1846](#)
- [Verification on page 1847](#)

### Requirements

---

This example uses the following hardware and software components:

- Junos OS Release 9.0 or later for EX Series switches
- One EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- One RADIUS authentication server that supports 802.1X. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you connect the server to the switch, be sure you have:

- Performed basic bridging and VLAN configuration on the switch. See the documentation that describes setting up basic bridging and a VLAN for your switch. If you are using a switch that supports the Enhanced Layer 2 Software (ELS) configuration style, see [“Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch” on page 2281](#). For all other switches, see *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*.



**NOTE:** For more about ELS, see: [“Getting Started with Enhanced Layer 2 Software” on page 3](#)

---

- Configured users on the RADIUS authentication server.

### Overview and Topology

---

The EX Series switch acts as an authenticator Port Access Entity (PAE). It blocks all traffic and acts as a control gate until the supplicant (client) is authenticated by the server. All other users and devices are denied access.

[Figure 10 on page 1845](#) shows one EX4200 switch that is connected to the devices listed in [Table 166 on page 1845](#).

Figure 10: Topology for Configuration

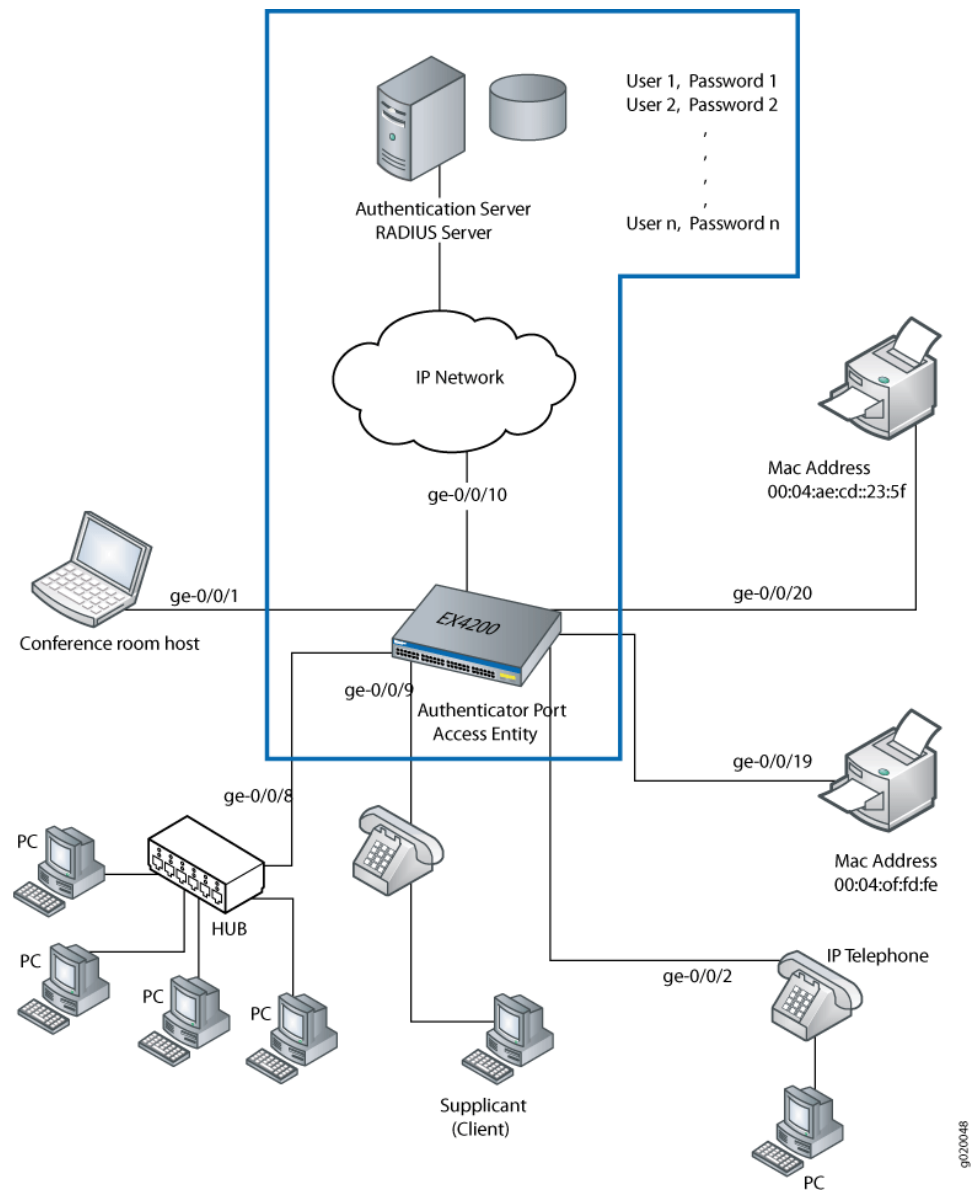


Table 166: Components of the Topology

Property	Settings
Switch hardware	EX4200 access switch, 24 Gigabit Ethernet ports: 8 PoE ports ( <b>ge-0/0/0</b> through <b>ge-0/0/7</b> ) and 16 non-PoE ports ( <b>ge-0/0/8</b> through <b>ge-0/0/23</b> )
VLAN name	<b>default</b>
One RADIUS server	Backend database with an address of <b>10.0.0.100</b> connected to the switch at port <b>ge-0/0/10</b>

In this example, connect the RADIUS server to access port **ge-0/0/10** on the EX4200 switch. The switch acts as the authenticator and forwards credentials from the supplicant to the user database on the RADIUS server. You must configure connectivity between the EX4200 and the RADIUS server by specifying the address of the server and configuring the secret password. This information is configured in an access profile on the switch.



**NOTE:** For more information about authentication, authorization, and accounting (AAA) services, see the [Junos OS System Basics Configuration Guide](#).

---

### Configuration

---

#### CLI Quick Configuration

To quickly connect the RADIUS server to the switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set access radius-server 10.0.0.100 secret juniper
set access radius-server 10.0.0.200 secret juniper
set access profile profile1 authentication-order radius
set access profile profile1 radius authentication-server [10.0.0.100 10.0.0.200]
```

#### Step-by-Step Procedure

To connect the RADIUS server to the switch:

1. Define the address of the servers, and configure the secret password. The secret password on the switch must match the secret password on the server:

```
[edit]
user@switch# set access radius-server 10.0.0.100 secret juniper
user@switch# set access radius-server 10.0.0.200 secret juniper
```

2. Configure the authentication order, making **radius** the first method of authentication:

```
[edit]
user@switch# set access profile profile1 authentication-order radius
```

3. Configure a list of server IP addresses to be tried in order to authenticate the supplicant:

```
[edit]
user@switch# set access profile profile1 radius authentication-server [10.0.0.100 10.0.0.200]
```

**Results** Display the results of the configuration:

```
user@switch> show configuration access
radius-server {
  10.0.0.100
  port 1812;
  secret "$9$qPT3ApBSrv69rvWLVb.P5"; ## SECRET-DATA
}
profile profile1{
  authentication-order radius;
  radius {
    authentication-server 10.0.0.100 10.0.0.200;
  }
}
```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verify That the Switch and RADIUS Server are Properly Connected on page 1847](#)

### *Verify That the Switch and RADIUS Server are Properly Connected*

<b>Purpose</b>	Verify that the RADIUS server is connected to the switch on the specified port.
<b>Action</b>	<p>Ping the RADIUS server to verify the connection between the switch and the server:</p> <pre>user@switch&gt; ping 10.0.0.100 PING 10.0.0.100 (10.0.0.100): 56 data bytes 64 bytes from 10.93.15.218: icmp_seq=0 ttl=64 time=9.734 ms 64 bytes from 10.93.15.218: icmp_seq=1 ttl=64 time=0.228 ms</pre>
<b>Meaning</b>	ICMP echo request packets are sent from the switch to the target server at 10.0.0.100 to test whether it is reachable across the IP network. ICMP echo responses are being returned from the server, verifying that the switch and the server are connected.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852</a></li> <li>• <a href="#">Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an EX Series Switch on page 1847</a></li> <li>• <a href="#">Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch</a></li> <li>• <a href="#">Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 1909</a></li> <li>• <a href="#">Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 1910</a></li> </ul>

## Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an EX Series Switch

802.1X on EX Series switches provides LAN access to users who do not have credentials in the RADIUS database. These users, referred to as guests, are authenticated and typically provided with access to the Internet.

This example describes how to create a guest VLAN and configure 802.1X authentication for it.

- [Requirements on page 1848](#)
- [Overview and Topology on page 1848](#)
- [Configuration of a Guest VLAN That Includes 802.1X Authentication on page 1850](#)
- [Verification on page 1850](#)

## Requirements

---

This example uses the following hardware and software components:

- Junos OS Release 9.0 or later for EX Series switches
- One EX Series switch acting as an authenticator interface access entity (PAE). The interfaces on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- One RADIUS authentication server that supports 802.1X. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you configure guest VLAN authentication, be sure you have:

- Installed your EX Series switch. See the installation information for your switch.
- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (CLI Procedure)*.
- Performed basic bridging and VLAN configuration on the switch. See the documentation that describes setting up basic bridging and a VLAN for your switch. If you are using a switch that supports the Enhanced Layer 2 Software (ELS) configuration style, see [“Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch” on page 2281](#). For all other switches, see *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*.



**NOTE:** For more about ELS, see: [“Getting Started with Enhanced Layer 2 Software” on page 3](#)

---

## Overview and Topology

---

As part of IEEE 802.1X Port-Based Network Access Control (PNAC), you can provide limited network access to supplicants who do not belong to a VLAN authentication group by configuring authentication to a guest VLAN. Typically, guest VLAN access is used to provide Internet access to visitors to a corporate site. However, you can also use the guest VLAN feature to provide supplicants that fail 802.1X authentication to a corporate LAN with access to a VLAN with limited resources.

[Figure 11 on page 1849](#) shows the conference room connected to the switch at interface `ge-0/0/1`.



Figure 11: Topology for Guest VLAN Example

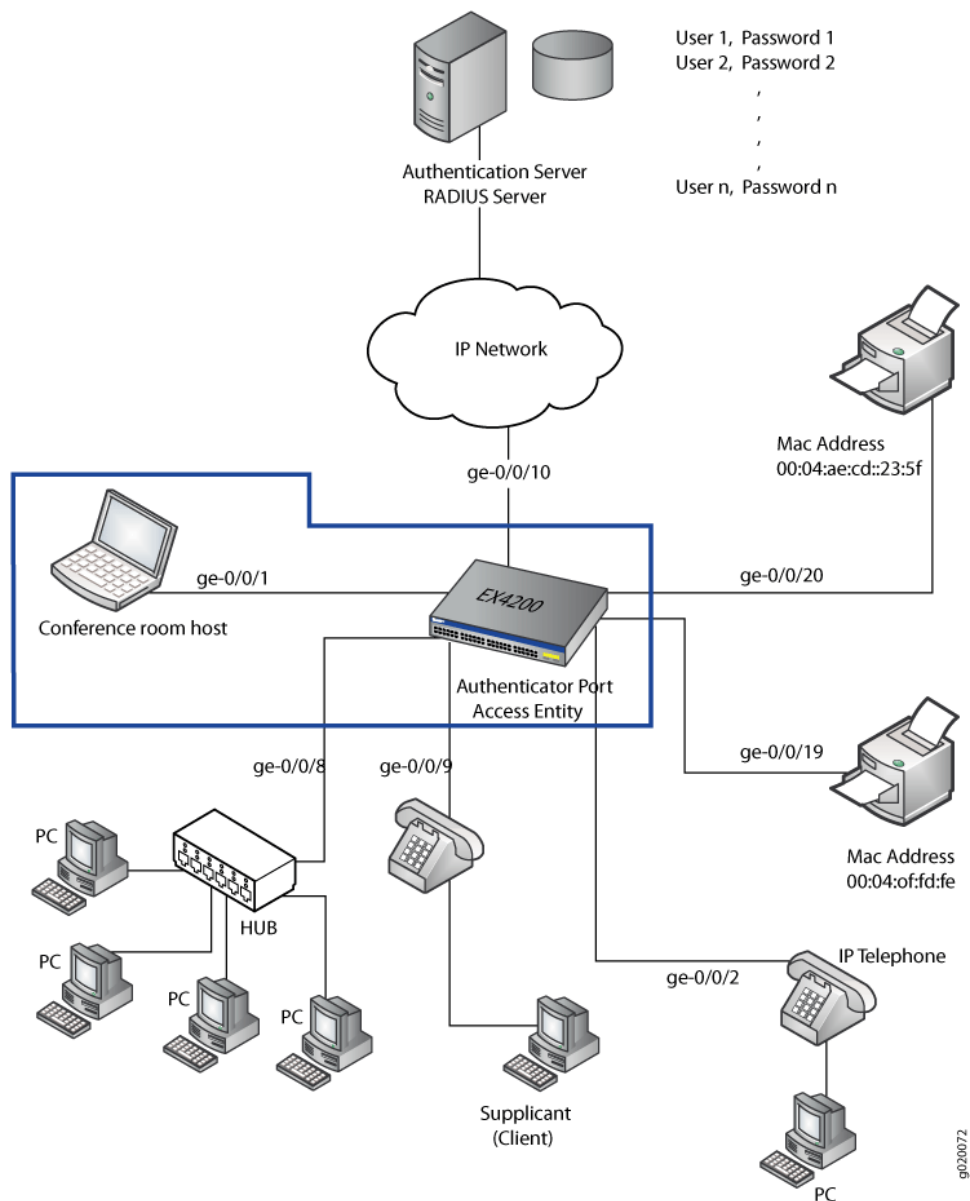


Table 167: Components of the Guest VLAN Topology

Property	Settings
Switch hardware	EX4200 switch, 24 Gigabit Ethernet interfaces: 8 PoE interfaces ( <b>ge-0/0/0</b> through <b>ge-0/0/7</b> ) and 16 non-PoE interfaces ( <b>ge-0/0/8</b> through <b>ge-0/0/23</b> )
VLAN names and tag IDs	<b>sales</b> , tag 100 <b>support</b> , tag 200  <b>guest-vlan</b> , tag 300
One RADIUS server	Backend database connected to the switch through interface <b>ge-0/0/10</b>

In this example, access interface **ge-0/0/1** provides LAN connectivity in the conference room. Configure this access interface to provide LAN connectivity to visitors in the conference room who are not authenticated by the corporate VLAN.

### Configuration of a Guest VLAN That Includes 802.1X Authentication

---

#### CLI Quick Configuration

To quickly configure a guest VLAN, with 802.1X authentication, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans guest-vlan vlan-id 300
set protocols dot1x authenticator interface all guest-vlan guest-vlan
```

#### Step-by-Step Procedure

To configure a guest VLAN that includes 802.1X authentication on an EX Series switch:

1. Configure the VLAN ID for the guest VLAN:

```
[edit]
user@switch# set vlans guest-vlan vlan-id 300
```

2. Configure the guest VLAN under **dot1x** protocols:

```
[edit]
user@switch# set protocols dot1x authenticator interface all guest-vlan guest-vlan
```

#### Results

Check the results of the configuration:

```
user@switch> show configuration
protocols {
  dot1x {
    authenticator {
      interface {
        all {
          guest-vlan {
            guest-vlan;
          }
        }
      }
    }
  }
}
vlangs {
  guest-vlan {
    vlan-id 300;
  }
}
```

### Verification

---

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Guest VLAN is Configured on page 1850](#)

#### *Verifying That the Guest VLAN is Configured*

#### Purpose

Verify that the guest VLAN is created and that an interface has failed authentication and been moved to the guest VLAN.



**NOTE:** On switches running Junos OS for EX Series with support for the Enhanced Layer 2 Software (ELS), the output for the `show vlans` command will contain additional information. If your switch runs software that supports ELS, see [show vlans](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software”](#) on page 3.

**Action** Use the operational mode commands:

```
user@switch> show vlans
```

Name	Tag	Interfaces
default		ge-0/0/3.0*
dynamic	40	None
guest	30	None
guest-vlan	300	ge-0/0/1.0*
vlan_dyn		None

```
user@switch> show dot1x interface ge-0/0/1.0 detail
ge-0/0/1.0
```

```

  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: guest-vlan
  Number of connected supplicants: 1
    Supplicant: user1, 00:00:00:00:13:23
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: vo11
      Dynamic Filter: match source-dot1q-tag 10 action deny
      Session Reauth interval: 60 seconds
      Reauthentication due in 50 seconds
```

**Meaning** The output from the `show vlans` command shows **guest-vlan** as the the name of the VLAN and the VLAN ID as **300**.

The output from the `show dot1x interface ge-0/0/1.0 detail` command displays the **Guest VLAN membership** field, indicating that a supplicant at this interface failed 802.1X authentication and was passed through to the **guest-vlan**.

**Related Documentation**

- [Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch on page 1843](#)
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch](#)
- [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)

## Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch

802.1x Port-Based Network Access Control (PNAC) authentication on EX Series switches provides three types of authentication to meet the access needs of your enterprise LAN:

- Authenticate the first end device (supplicant) on an authenticator port, and allow all others also connecting to have access.
- Authenticate only one end device on an authenticator port at one time.
- Authenticate multiple end devices on an authenticator port. Multiple supplicant mode is used in VoIP configurations.

This example configures an EX4200 switch to use IEEE 802.1X to authenticate end devices that use three different administrative modes:

- [Requirements on page 1852](#)
- [Overview and Topology on page 1853](#)
- [Configuration of 802.1X to Support Multiple Supplicant Modes on page 1855](#)
- [Verification on page 1856](#)

### Requirements

---

This example uses the following hardware and software components:

- Junos OS Release 9.0 or later for EX Series switches
- One EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from end devices until they are authenticated.
- One RADIUS authentication server that supports 802.1X. The authentication server acts as the backend database and contains credential information for end devices (supplicants) that have permission to connect to the network.

Before you configure the ports for 802.1X authentication, be sure you have:

- Installed your EX Series switch.
- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (CLI Procedure)*.
- Performed basic bridging and VLAN configuration on the switch. See the documentation that describes setting up basic bridging and a VLAN for your switch. If you are using a

switch that supports the Enhanced Layer 2 Software (ELS) configuration style, see [“Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch” on page 2281](#). For all other switches, see *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*.



**NOTE:** For more about ELS, see: [“Getting Started with Enhanced Layer 2 Software” on page 3](#)

- Configured users on the authentication server.

### Overview and Topology

As shown in [Figure 12 on page 1854](#), the topology contains an EX4200 access switch connected to the authentication server on port **ge-0/0/10**. Interfaces **ge-0/0/8**, **ge-0/0/9**, and **ge-0/0/11** will be configured for three different administrative modes.

Figure 12: Topology for Configuring Supplicant Modes




Table 168: Components of the Supplicant Mode Configuration Topology

Property	Settings
Switch hardware	EX4200 switch, 24 Gigabit Ethernet ports: 8 PoE ports ( <b>ge-0/0/0</b> through <b>ge-0/0/7</b> ) and 16 non-PoE ports ( <b>ge-0/0/8</b> through <b>ge-0/0/23</b> )
Connections to Avaya phones—with integrated hub, to connect phone and desktop PC to a single port; (requires PoE)	<b>ge-0/0/8</b> , <b>ge-0/0/9</b> , and <b>ge-0/0/11</b>

To configure the administrative modes to support supplicants in different areas of the Enterprise network:

- Configure access port **ge-0/0/8** for single supplicant mode authentication.
- Configure access port **ge-0/0/9** for single secure supplicant mode authentication.
- Configure access port **ge-0/0/11** for multiple supplicant mode authentication.

*Single supplicant mode* authenticates only the first end device that connects to an authenticator port. All other end devices connecting to the authenticator port after the first has connected successfully, whether they are 802.1X-enabled or not, are permitted free access to the port without further authentication. If the first authenticated end device logs out, all other end devices are locked out until an end device authenticates.

*Single-secure supplicant mode* authenticates only one end device to connect to an authenticator port. No other end device can connect to the authenticator port until the first logs out.

*Multiple supplicant mode* authenticates multiple end devices individually on one authenticator port. If you configure a maximum number of devices that can be connected to a port through port security, the lesser of the configured values is used to determine the maximum number of end devices allowed per port.

### Configuration of 802.1X to Support Multiple Supplicant Modes

**CLI Quick Configuration** To quickly configure the ports with different 802.1X authentication modes, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols dot1x authenticator interface ge-0/0/8 supplicant single
set protocols dot1x authenticator interface ge-0/0/9 supplicant single-secure
set protocols dot1x authenticator interface ge-0/0/11 supplicant multiple
```

**Step-by-Step Procedure** Configure the administrative mode on the interfaces:

1. Configure the supplicant mode as single on interface **ge-0/0/8**:  

```
[edit protocols]
user@switch# set dot1x authenticator interface ge-0/0/8 supplicant single
```
2. Configure the supplicant mode as single secure on interface **ge-0/0/9**:  

```
[edit protocols]
user@switch# set dot1x authenticator interface ge-0/0/9 supplicant single-secure
```
3. Configure multiple supplicant mode on interface **ge-0/0/11**:  

```
[edit protocols]
user@switch# set dot1x authenticator interface ge-0/0/11 supplicant multiple
```

### Results

Check the results of the configuration:

```
[edit]
user@access-switch> show configuration
protocols {
  dot1x {
    authenticator {
      interface {
        ge-0/0/8.0 {
          supplicant single;
```

```
    )
    ge-0/0/9.0 {
        supplicant single-secure;
    }
    ge-0/0/11.0 {
        supplicant multiple;
    }
}
}
```

---

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying the 802.1X Configuration on page 1856](#)

#### *Verifying the 802.1X Configuration*

**Purpose** Verify the 802.1X configuration on interfaces **ge-0/0/8**, **ge-0/0/9**, and **ge-0/0/5**.



**Action** Verify the 802.1X configuration with the operational mode command **show dot1x interface**:

```
user@switch> show dot1x interface ge-0/0/8.0 detail
ge-0/0/8.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Disabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
```

```
user@switch> show dot1x interface ge-0/0/9.0 detail
ge-0/0/9.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single-Secure
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Disabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 0
```

```
user@switch> show dot1x interface ge-0/0/11.0 detail
ge-0/0/11.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Multiple
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Disabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 0
```

**Meaning** The **Supplicant mode** output field displays the configured administrative mode for each interface. Interface **ge-0/0/8.0** displays **Single** supplicant mode. Interface **ge-0/0/9.0** displays **Single Secure** supplicant mode. Interface **ge-0/0/11.0** displays **Multiple** supplicant mode.

**Related Documentation**

- [Controlling Authentication Session Timeouts \(CLI Procedure\) on page 1930](#)
- [Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch on page 1843](#)
- [Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an EX Series Switch on page 1847](#)
- [Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch](#)
- [Configuring 802.1X RADIUS Accounting \(CLI Procedure\) on page 1909](#)
- [Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 1910](#)
- [Understanding Authentication on EX Series Switches on page 1824](#)

## Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch

---



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

---

You can configure voice over IP (VoIP) on an EX Series switch to support IP telephones. The Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) protocol forwards VoIP parameters from the switch to the phone. You also configure 802.1X authentication to allow the telephone access to the LAN. Authentication is done through a backend RADIUS server.

This example describes how to configure VoIP on an EX Series switch to support an Avaya IP phone, as well as the LLDP-MED protocol and 802.1X authentication:

- [Requirements on page 1858](#)
- [Overview and Topology on page 1859](#)
- [Configuration on page 1861](#)
- [Verification on page 1863](#)

### Requirements

---

This example uses the following hardware and software components:

- Junos OS Release 13.2X50 or later for EX Series switches
- One EX4300 switch acting as an authenticator port access entity (PAE). The interfaces on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- An Avaya IP telephone that supports LLDP-MED and 802.1X

Before you configure VoIP, be sure you have:

- Installed your EX Series switch. See the installation information for your switch.
- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (CLI Procedure)*.
- Performed basic bridging and VLAN configuration on the switch. See [“Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch” on page 2281](#).
- Configured the RADIUS server for 802.1X authentication and set up the access profile. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).
- (Optional) Configured the interface ge-0/0/2 for Power over Ethernet (PoE). The PoE configuration is not necessary if the VoIP supplicant is using a power adapter. For information about configuring PoE, see [“Configuring PoE on EX Series Switches \(CLI Procedure\)” on page 4440](#).




**NOTE:** If the IP address is not configured on the Avaya IP phone, the phone exchanges LLDP-MED information to get the VLAN ID for the voice VLAN. You must configure the `voip` statement on the interface to designate the interface as a VoIP interface and allow the switch to forward the VLAN name and VLAN ID for the voice VLAN to the IP telephone. The IP telephone then uses the voice VLAN (that is, it references the voice VLAN's ID) to send a DHCP discover request and exchange information with the DHCP server (voice gateway).

## Overview and Topology

Instead of using a regular telephone, you connect an IP telephone directly to the switch. An IP phone has all the hardware and software needed to handle VoIP. You also can power an IP telephone by connecting it to one of the Power over Ethernet (PoE) interfaces on the switch.

In this example, the access interface ge-0/0/2 on the EX4300 switch is connected to an Avaya IP telephone. Avaya phones have a built-in bridge that allows you to connect a desktop PC to the phone, so the desktop and phone in a single office require only one interface on the switch. The EX Series switch is connected to a RADIUS server on the ge-0/0/10 interface (see [Figure 13 on page 1860](#)).

Figure 13: VoIP Topology



In this example, you configure VoIP parameters and specify the forwarding class **assured-forward** for voice traffic to provide the highest quality of service.

[Table 169 on page 1860](#) describes the components used in this VoIP configuration example.

Table 169: Components of the VoIP Configuration Topology

Property	Settings
Switch hardware	EX4300 switch

Table 169: Components of the VoIP Configuration Topology (*continued*)

Property	Settings
VLAN names	<b>data-vlan</b> <b>voice-vlan</b>
Connection to Avaya phone—with integrated hub, to connect phone and desktop PC to a single interface (requires PoE)	<b>ge-0/0/2</b>
One RADIUS server	Provides backend database connected to the switch through interface <b>ge-0/0/10</b> .

As well as configuring a VoIP for interface ge-0/0/2, you configure:

- 802.1X authentication. Authentication is set to **multiple** supplicant mode to support more than one supplicant's access to the LAN through interface ge-0/0/2.
- LLDP-MED protocol information. The switch uses LLDP-MED to forward VoIP parameters to the phone. Using LLDP-MED ensures that voice traffic gets tagged and prioritized with the correct values at the source itself. For example, 802.1p class of service and 802.1Q tag information can be sent to the IP telephone.



**NOTE:** A PoE configuration is not necessary if an IP telephone is using a power adapter.

### Configuration

#### CLI Quick Configuration

To quickly configure VoIP, LLDP-MED, and 802.1X, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans data-vlan vlan-id 77
set vlans voice-vlan vlan-id 99
set vlans data-vlan switch-options interface ge-0/0/2.0
set interfaces ge-0/0/2 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members data-vlan
set switch-options voip interface ge-0/0/2.0 vlan voice-vlan
set switch-options voip interface ge-0/0/2.0 forwarding-class assured-forwarding
set protocols lldp-med interface ge-0/0/2
set protocols dot1x authenticator interface ge-0/0/2.0 supplicant multiple
```

#### Step-by-Step Procedure

To configure VoIP with LLDP-MED and 802.1X:

1. Configure the VLANs for voice and data:
 

```
[edit vlans]
user@switch# set data-vlan vlan-id 77
user@switch# set voice-vlan vlan-id 99
```
2. Associate the VLAN **data-vlan** with the interface:
 

```
[edit vlans]
user@switch# set data-vlan switch-options interface ge-0/0/2.0
```
3. Configure the interface as an access interface, configure support for Ethernet switching, and add the **data-vlan** VLAN:

- ```
[edit interfaces]
user@switch# set ge-0/0/2 unit 0 family ethernet-switching interface-mode access
user@switch# set ge-0/0/2 unit 0 family ethernet-switching vlan members data-vlan
```
4. Configure VoIP on the interface and specify the **assured-forwarding** forwarding class to provide the most dependable class of service:
 

```
[edit switch-options]
user@switch# set voip interface ge-0/0/2.0 vlan voice-vlan
user@switch# set voip interface ge-0/0/2.0 forwarding-class assured-forwarding
```
  5. Configure LLDP-MED protocol support:
 

```
[edit protocols]
user@switch# set lldp-med interface ge-0/0/2
```
  6. To authenticate an IP phone and a PC connected to the IP phone on the interface, configure 802.1X authentication support and specify **multiple** supplicant mode:



**NOTE:** If you do not want to authenticate any device, skip the 802.1X configuration on this interface.

```
[edit protocols]
user@switch# set dot1x authenticator interface ge-0/0/2.0 supplicant multiple
```

**Results** Display the results of the configuration:

```
[edit]
user@switch# show configuration
interfaces {
  ge-0/0/2 {
    unit 0 {
      family ethernet-switching {
        interface-mode access;
        vlan {
          members data-vlan;
        }
      }
    }
  }
}
protocols {
  lldp-med {
    interface ge-0/0/2;
  }
  dot1x {
    authenticator {
      interface {
        ge-0/0/2.0 {
          supplicant multiple;
        }
      }
    }
  }
}
vllans {
  data-vlan {
```

```
    vlan-id 77;
    switch-options {
        interface ge-0/0/2.0;
    }
}
voice-vlan {
    vlan-id 99;
}
}
switch-options {
    voip {
        interface ge-0/0/2.0 {
            vlan voice-vlan;
            forwarding-class assured-forwarding;
        }
    }
}
```

---

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying LLDP-MED Configuration on page 1863](#)
- [Verifying 802.1X Authentication for IP Phone and Desktop PC on page 1864](#)
- [Verifying the VLAN Association with the Interface on page 1865](#)

#### ***Verifying LLDP-MED Configuration***

**Purpose** Verify that LLDP-MED is enabled on the interface.

**Action** user@switch> **show lldp detail**

```
LLDP : Enabled
Advertisement interval : 30 seconds
Transmit delay : 2 seconds
Hold timer : 120 seconds
Notification interval : 0 Second(s)
Config Trap Interval : 0 seconds
Connection Hold timer : 300 seconds
```

```
LLDP MED : Enabled
MED fast start count : 3 Packets
```

```
Port ID TLV subtype : locally-assigned
```

| Interface      | Parent Interface | LLDP    | LLDP-MED | Power Negotiation |
|----------------|------------------|---------|----------|-------------------|
| Neighbor count |                  |         |          |                   |
| all            | -                | Enabled | Enabled  | Enabled           |
| 0              |                  |         |          |                   |
| ge-0/0/2       | -                | -       | Enabled  | -                 |
| 0              |                  |         |          |                   |

| Interface | Parent Interface | Vlan-id | Vlan-name |
|-----------|------------------|---------|-----------|
| ge-0/0/0  | -                | 1       | vlan-1    |
| ge-0/0/1  | -                | 1       | vlan-1    |
| ge-0/0/2  | -                | 77      | vlan-77   |
| ge-0/0/2  | -                | 99      | vlan-99   |
| ge-0/0/3  | -                | 1       | vlan-1    |
| ge-0/0/4  | -                | 1       | vlan-1    |
| ge-0/0/5  | -                | 1       | vlan-1    |
| ge-0/0/6  | -                | 1       | vlan-1    |
| ge-0/0/7  | -                | 1       | vlan-1    |
| ge-0/0/8  | -                | 1       | vlan-1    |
| ge-0/0/9  | -                | 1       | vlan-1    |
| ge-0/0/10 | -                | 1       | vlan-1    |

Basic Management TLVs supported:  
End Of LLDPDU, Chassis ID, Port ID, Time To Live, Port Description, System Name,  
System Description, System Capabilities, Management Address

Organizationally Specific TLVs supported:  
MAC/PHY configuration/status, Power via MDI, Link aggregation, Maximum Frame Size,  
Port VLAN tag, Port VLAN name.

**Meaning** The **show lldp detail** output shows that both LLDP and LLDP-MED are configured on the ge-0/0/2 interface. The end of the output shows the list of supported LLDP basic management TLVs and organizationally specific TLVs that are supported.

#### **Verifying 802.1X Authentication for IP Phone and Desktop PC**

**Purpose** Display the 802.1X configuration to confirm that the VoIP interface has access to the LAN.



**Action** user@switch> **show dot1x** interface ge-0/0/2.0 detail  
ge-0/0/2.0  
Role: Authenticator  
Administrative state: Auto  
Supplicant mode: Multiple  
Number of retries: 3  
Quiet period: 60 seconds  
Transmit period: 30 seconds  
Mac Radius: Disabled  
Mac Radius Restrict: Disabled  
Reauthentication: Enabled  
Configured Reauthentication interval: 3600 seconds  
Supplicant timeout: 30 seconds  
Server timeout: 30 seconds  
Maximum EAPOL requests: 2  
Guest VLAN member: <not configured>  
Number of connected supplicants: 1  
Supplicant: user101, 00:04:0f:fd:ac:fe  
Operational state: Authenticated  
Authentication method: Radius  
Authenticated VLAN: vo11  
Dynamic Filter: match source-dot1q-tag 10 action deny  
Session Reauth interval: 60 seconds  
Reauthentication due in 50 seconds

**Meaning** The field **Role** shows that the ge-0/0/2.0 interface is in the authenticator state. The **Supplicant** field shows that the interface is configured in multiple supplicant mode, permitting multiple supplicants to be authenticated on this interface. The MAC addresses of the supplicants currently connected are displayed at the bottom of the output.

#### *Verifying the VLAN Association with the Interface*

**Purpose** Display the interface's VLAN membership.

**Action** user@switch> **show ethernet-switching interface ge-0/0/2.0**  
Routing Instance Name : default-switch  
Logical Interface flags (DL - disable learning, AD - packet action drop,  
LH - MAC limit hit, DN - interface down )  

| Logical interface | Vlan members  | TAG | MAC limit | STP state  | Logical interface flags | Tagging  |
|-------------------|---------------|-----|-----------|------------|-------------------------|----------|
| ge-0/0/2.0        |               |     | 65535     |            |                         | untagged |
|                   | voice-vlan 99 |     | 65535     | Discarding |                         |          |
|                   | data-vlan 77  |     | 65535     | Discarding |                         |          |

**Meaning** The field **VLAN members** shows that the ge-0/0/2.0 interface supports both the **data-vlan** VLAN and **voice-vlan** VLAN.

**Related Documentation**

- [Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch on page 1843](#)
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Defining CoS Forwarding Classes \(CLI Procedure\) on page 2107](#)

- [Defining CoS Forwarding Classes \(J-Web Procedure\) on page 2107](#)
- [Configuring LLDP-MED \(CLI Procedure\) on page 1917](#)

## Example: Configuring VoIP on an EX Series Switch Without Including 802.1X Authentication

---



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring VoIP on an EX Series Switch Without Including 802.1X Authentication*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

---

You can configure voice over IP (VoIP) on an EX Series switch to support IP telephones.

To configure VoIP on an EX Series switch to support an IP phone that does not support 802.1X authentication, you must either add the MAC address of the phone to the static MAC bypass list or enable MAC RADIUS authentication on the switch.

This example describes how to configure VoIP on an EX Series switch without 802.1X authentication by using static MAC bypass of authentication:

- [Requirements on page 1866](#)
- [Overview on page 1867](#)
- [Configuration on page 1867](#)
- [Verification on page 1869](#)

### Requirements

---

This example uses the following hardware and software components:

- One EX4300 switch.
- Junos OS Release 13.2 or later for EX Series switches
- An Avaya IP telephone

Before you configure VoIP, be sure you have:

- Installed your EX Series switch. See the installation information for your switch.
- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (CLI Procedure)*.
- Performed basic bridging and VLAN configuration on the switch. See [“Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch” on page 2281](#).
- Configured the RADIUS server for 802.1X authentication and set up the access profile. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).
- (Optional) Configured the interface ge-0/0/2 for Power over Ethernet (PoE). The PoE configuration is not necessary if the VoIP supplicant is using a power adapter. For

information about configuring PoE, see “Configuring PoE on EX Series Switches (CLI Procedure)” on page 4440.



**NOTE:** If the IP address is not configured on the Avaya IP phone, the phone exchanges LLDP-MED information to get the VLAN ID for the voice VLAN. You must configure the `voip` statement on the interface to designate the interface as a VoIP interface and allow the switch to forward the VLAN name and VLAN ID for the voice VLAN to the IP telephone. The IP telephone then uses the voice VLAN (that is, it references the voice VLAN's ID) to send a DHCP discover request and exchange information with the DHCP server (voice gateway).

## Overview

Instead of using a regular telephone, you connect an IP telephone directly to the switch. An IP phone has all the hardware and software needed to handle VoIP. You also can power an IP telephone by connecting it to one of the Power over Ethernet (PoE) interfaces on the switch.

In this example, the access interface `ge-0/0/2` on the EX4300 switch is connected to a non-802.1X IP phone.

To configure VoIP on an EX Series switch to support an IP phone that does not support 802.1X authentication, add the MAC address of the phone as a static entry in the authenticator database and set the supplicant mode to multiple.

## Configuration

### CLI Quick Configuration

To quickly configure VoIP without using 802.1X authentication, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans data-vlan vlan-id 77
set vlans voice-vlan vlan-id 99
set vlans data-vlan switch-options interface ge-0/0/2.0
set interfaces ge-0/0/2 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members data-vlan
set switch-options voip interface ge-0/0/2.0 vlan voice-vlan
set switch-options voip interface ge-0/0/2.0 forwarding-class assured-forwarding
set protocols lldp-med interface ge-0/0/2
set protocols dot1x authenticator authentication-profile-name auth-profile
set protocols dot1x authenticator static 00:04:f2:11:aa:a7
set protocols dot1x authenticator interface ge-0/0/2.0 supplicant multiple
```

### Step-by-Step Procedure

To configure VoIP without 802.1X authentication:

1. Configure the VLANs for voice and data:
 

```
[edit vlans]
user@switch# set data-vlan vlan-id 77
user@switch# set voice-vlan vlan-id 99
```
2. Associate the VLAN `data-vlan` with the interface:
 

```
[edit vlans]
```

- ```

user@switch# set data-vlan switch-options interface ge-0/0/2.0

```
3. Configure the interface as an access interface, configure support for Ethernet switching, and add the **data-vlan** VLAN:
 

```

[edit interfaces]
user@switch# set ge-0/0/2 unit 0 family ethernet-switching interface-mode access
user@switch# set ge-0/0/2 unit 0 family ethernet-switching vlan members data-vlan

```
  4. Configure VoIP on the interface and specify the **assured-forwarding** forwarding class to provide the most dependable class of service:
 

```

[edit switch-options]
user@switch# set voip interface ge-0/0/2.0 vlan voice-vlan
user@switch# set voip interface ge-0/0/2.0 forwarding-class assured-forwarding

```
  5. Configure LLDP-MED protocol support:
 

```

[edit protocols]
user@switch# set lldp-med interface ge-0/0/2

```
  6. Set the authentication profile (see ["Configuring 802.1X Interface Settings \(CLI Procedure\)" on page 1908](#) and ["Configuring 802.1X RADIUS Accounting \(CLI Procedure\)" on page 1909](#)):
 

```

[edit protocols]
set dot1x authenticator authentication-profile-name auth-profile

```
  7. Add the MAC address of the phone to the static MAC bypass list:
 

```

[edit protocols]
set dot1x authenticator static 00:04:f2:11:aa:a7

```
  8. Set the supplicant mode to **multiple**:
 

```

[edit protocols]
set dot1x authenticator interface ge-0/0/2.0 supplicant multiple

```

**Results** Display the results of the configuration:

```

[edit]
user@switch# show configuration
interfaces {
  ge-0/0/2 {
    unit 0 {
      family ethernet-switching {
        interface-mode access;
        vlan {
          members data-vlan;
        }
      }
    }
  }
}
protocols {
  lldp-med {
    interface ge-0/0/2;
  }
  dot1x {
    authenticator {
      authentication-profile-name auth-profile;
      static {
        00:04:f2:11:aa:a7;
      }
    }
  }
}

```

```

    }
    interface {
        ge-0/0/2.0 {
            supplicant multiple;
        }
    }
}
vllans {
    data-vlan {
        vlan-id 77;
        switch-options {
            interface ge-0/0/2.0;
        }
    }
    voice-vlan {
        vlan-id 99;
    }
}
switch-options {
    voip {
        interface ge-0/0/2.0 {
            vlan voice-vlan;
            forwarding-class assured-forwarding;
        }
    }
}
}

```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying LLDP-MED Configuration on page 1869](#)
- [Verifying Authentication for the Desktop PC on page 1870](#)
- [Verifying the VLAN Association with the Interface on page 1871](#)

#### **Verifying LLDP-MED Configuration**

**Purpose** Verify that LLDP-MED is enabled on the interface.

**Action** user@switch> **show lldp detail**

```
LLDP : Enabled
Advertisement interval : 30 seconds
Transmit delay : 2 seconds
Hold timer : 120 seconds
Notification interval : 0 Second(s)
Config Trap Interval : 0 seconds
Connection Hold timer : 300 seconds
```

```
LLDP MED : Enabled
MED fast start count : 3 Packets
```

```
Port ID TLV subtype : locally-assigned
```

Interface	Parent Interface	LLDP	LLDP-MED	Power Negotiation
Neighbor count				
all	-	Enabled	Enabled	Enabled
0				
ge-0/0/2	-	-	Enabled	-
0				

Interface	Parent Interface	Vlan-id	Vlan-name
ge-0/0/0	-	1	vlan-1
ge-0/0/1	-	1	vlan-1
ge-0/0/2	-	77	vlan-77
ge-0/0/2	-	99	vlan-99
ge-0/0/3	-	1	vlan-1
ge-0/0/4	-	1	vlan-1
ge-0/0/5	-	1	vlan-1
ge-0/0/6	-	1	vlan-1
ge-0/0/7	-	1	vlan-1
ge-0/0/8	-	1	vlan-1
ge-0/0/9	-	1	vlan-1
ge-0/0/10	-	1	vlan-1

Basic Management TLVs supported:  
End Of LLDPDU, Chassis ID, Port ID, Time To Live, Port Description, System Name,  
System Description, System Capabilities, Management Address

Organizationally Specific TLVs supported:  
MAC/PHY configuration/status, Power via MDI, Link aggregation, Maximum Frame Size,  
Port VLAN tag, Port VLAN name.

**Meaning** The **show lldp detail** output shows that both LLDP and LLDP-MED are configured on the ge-0/0/2 interface. The end of the output shows the list of supported LLDP basic management TLVs and organizationally specific TLVs that are supported.

#### ***Verifying Authentication for the Desktop PC***

**Purpose** Display the 802.1X configuration for the desktop PC connected to the VoIP interface through the IP phone.

**Action** user@switch> `show dot1x interface ge-0/0/2.0 detail`  
 ge-0/0/2.0  
 Role: Authenticator  
 Administrative state: Auto  
 Supplicant mode: Multiple  
 Number of retries: 3  
 Quiet period: 60 seconds  
 Transmit period: 30 seconds  
 Mac Radius: Disabled  
 Mac Radius Restrict: Disabled  
 Reauthentication: Enabled  
 Configured Reauthentication interval: 3600 seconds  
 Supplicant timeout: 30 seconds  
 Server timeout: 30 seconds  
 Maximum EAPOL requests: 2  
 Guest VLAN member: <not configured>  
 Number of connected supplicants: 1  
 Supplicant: user101, 00:04:0f:fd:ac:fe  
 Operational state: Authenticated  
 Authentication method: Radius  
 Authenticated VLAN: vo11  
 Dynamic Filter: match source-dot1q-tag 10 action deny  
 Session Reauth interval: 60 seconds  
 Reauthentication due in 50 seconds

**Meaning** The field **Role** shows that the ge-0/0/2.0 interface is in the authenticator state. The **Supplicant** field shows that the interface is configured in multiple supplicant mode, permitting multiple supplicants to be authenticated on this interface. The MAC addresses of the supplicants currently connected are displayed at the bottom of the output.

#### *Verifying the VLAN Association with the Interface*

**Purpose** Display the interface's VLAN membership.

**Action** user@switch> `show ethernet-switching interface ge-0/0/2.0`  
 Routing Instance Name : default-switch  
 Logical Interface flags (DL - disable learning, AD - packet action drop,  
 LH - MAC limit hit, DN - interface down )  

Logical interface	Vlan members	TAG	MAC limit	STP state	Logical interface flags	Tagging
ge-0/0/2.0			65535			untagged
	voice-vlan 99		65535	Discarding		
	data-vlan 77		65535	Discarding		

**Meaning** The field **VLAN members** shows that the ge-0/0/2.0 interface supports both the **data-vlan** VLAN and **voice-vlan** VLAN.

**Related Documentation**

- [Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch on page 1858](#)
- [Understanding 802.1X and VoIP on EX Series Switches on page 1834](#)
- [Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831](#)

## Example: Configuring Static MAC Bypass of Authentication on an EX Series Switch

To allow devices to access your LAN through 802.1X-configured interfaces without authentication, you can configure a static MAC bypass list on the EX Series switch. The static MAC bypass list, also known as the *exclusion list*, specifies MAC addresses that are allowed on the switch without a request to an authentication server.

You can use static MAC bypass of authentication to allow connection for devices that are not 802.1X-enabled, such as printers. If a host's MAC address is compared and matched against the static MAC address list, the nonresponsive host is authenticated and an interface opened for it.

This example describes how to configure static MAC bypass of authentication for two printers:

- [Requirements on page 1872](#)
- [Overview and Topology on page 1872](#)
- [Configuration on page 1874](#)
- [Verification on page 1875](#)

---

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.0 or later for EX Series switches
- One EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.

Before you configure static MAC authentication, be sure you have:

- Performed basic bridging and VLAN configuration on the switch. See the documentation that describes setting up basic bridging and a VLAN for your switch. If you are using a switch that supports the Enhanced Layer 2 Software (ELS) configuration style, see [“Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch” on page 2281](#). For all other switches, see *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*.

For more about ELS, see: [“Getting Started with Enhanced Layer 2 Software” on page 3](#)

- Specified the RADIUS server connections and configured an access profile on the switch. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).

---


### Overview and Topology

To permit printers access to the LAN, add them to the static MAC bypass list. The MAC addresses on this list are permitted access without authentication from the RADIUS server.

[Figure 14 on page 1873](#) shows the two printers connected to the EX4200.



Figure 14: Topology for Static MAC Authentication Configuration



The interfaces shown in [Table 170 on page 1873](#) will be configured for static MAC authentication.

Table 170: Components of the Static MAC Authentication Configuration Topology

Property	Settings
Switch hardware	EX4200, 24 Gigabit Ethernet ports: 8 PoE ports ( <b>ge-0/0/0</b> through <b>ge-0/0/23</b> )
VLAN name	<b>default</b>
Connections to integrated printer/fax/copier machines (no PoE required)	<b>ge-0/0/19</b> , MAC address 00:04:0f:fd:ac:fe <b>ge-0/0/20</b> , MAC address 00:04:ae:cd:23:5f

The printer with the MAC address 00:04:0f:fd:ac:fe is connected to access interface **ge-0/0/19**. A second printer with the MAC address 00:04:ae:cd:23:5f is connected to access interface **ge-0/0/20**. Both printers will be added to the static list and bypass 802.1X authentication.

### Configuration

---

#### CLI Quick Configuration

To quickly configure static MAC authentication, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols dot1x authenticator static [00:04:0f:fd:ac:fe 00:04:ae:cd:23:5f]
set protocols dot1x authenticator interface all supplicant multiple
set protocols dot1x authenticator authentication-profile-name profile1
```

#### Step-by-Step Procedure

Configure static MAC authentication:

1. Configure MAC addresses **00:04:0f:fd:ac:fe** and **00:04:ae:cd:23:5f** as static MAC addresses:

```
[edit protocols]
user@switch# set dot1x authenticator static [00:04:0f:fd:ac:fe 00:04:ae:cd:23:5f]
```

2. Configure the 802.1X authentication method:

```
[edit protocols]
user@switch# set dot1x authenticator interface all supplicant multiple
```

3. Configure the authentication profile name (access profile name) to use for authentication:

```
[edit protocols]
user@switch# set dot1x authenticator authentication-profile-name profile1
```



**NOTE:** Access profile configuration is required only for 802.1X clients, not for static MAC clients.

---

**Results** Display the results of the configuration:

```
user@switch> show
interfaces {
  ge-0/0/19 {
    unit 0 {
      family ethernet-switching {
        vlan members default;
      }
    }
  }
  ge-0/0/20 {
    unit 0 {
      family ethernet-switching {
        vlan members default;
      }
    }
  }
}
```

```

protocols {
  dot1x {
    authenticator {
      authentication-profile-name profile1
      static [00:04:0f:fd:ac:fe 00:04:ae:cd:23:5f];
      interface {
        all {
          supplicant multiple;
        }
      }
    }
  }
}

```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying Static MAC Bypass of Authentication on page 1875](#)

#### *Verifying Static MAC Bypass of Authentication*

**Purpose** Verify that the MAC address for both printers is configured and associated with the correct interfaces.

**Action** Use the operational mode command:

```
user@switch> show dot1x static-mac-address
```

MAC address	VLAN-Assignment	Interface
00:04:0f:fd:ac:fe	default	ge-0/0/19.0
00:04:ae:cd:23:5f	default	ge-0/0/20.0

**Meaning** The output field **MAC address** shows the MAC addresses of the two printers.

The output field **Interface** shows that the MAC address **00:04:0f:fd:ac:fe** can connect to the LAN through interface **ge-0/0/19.0** and that the MAC address **00:04:ae:cd:23:5f** can connect to the LAN through interface **ge-0/0/20.0**.

- Related Documentation**
- [Configuring 802.1X Authentication \(J-Web Procedure\)](#)
  - [Configuring Static MAC Bypass of Authentication \(CLI Procedure\) on page 1924](#)
  - [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)
  - [Understanding Authentication on EX Series Switches on page 1824](#)

### Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to an EX Series Switch

Server fail fallback allows you to specify how 802.1X supplicants connected to the switch are supported if the RADIUS authentication server becomes unavailable or sends a RADIUS access-reject message.

You use 802.1X to control network access. Only users and devices (supplicants) providing credentials that have been verified against a user database are allowed access to the network. You use a RADIUS server as the user database.

This example describes how to configure an interface to move a supplicant to a VLAN in the event of a RADIUS server timeout:

- [Requirements on page 1876](#)
- [Overview and Topology on page 1876](#)
- [Configuration on page 1878](#)
- [Verification on page 1879](#)

---

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.3 or later for EX Series switches
- One EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- One RADIUS authentication server that supports 802.1X. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you connect the server to the switch, be sure you have:

- Performed basic bridging and VLAN configuration on the switch. See the documentation that describes setting up basic bridging and a VLAN for your switch. If you are using a switch that supports the Enhanced Layer 2 Software (ELS) configuration style, see [“Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch” on page 2281](#). For all other switches, see *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*.



**NOTE:** For more about ELS, see: [“Getting Started with Enhanced Layer 2 Software” on page 3](#)

---

- Set up a connection between the switch and the RADIUS server. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).
- Disable firewall filters on the interface. Firewall filters interfere with server fail fallback operation.
- Configured users on the authentication server.

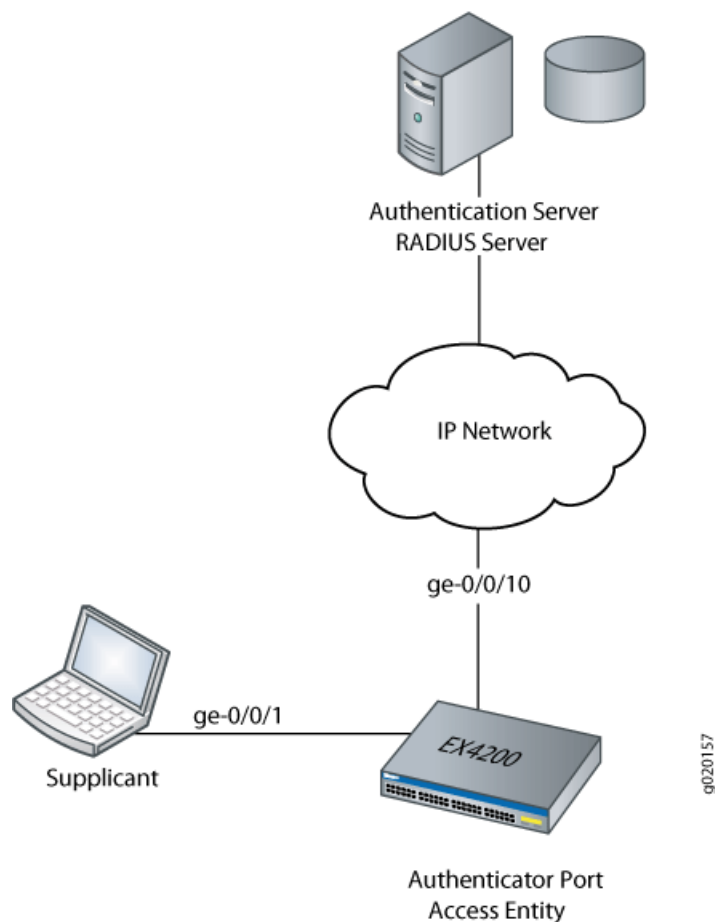
---

### Overview and Topology

A RADIUS server timeout occurs if no authentication RADIUS servers are reachable when a supplicant logs in and attempts to access the LAN. Using server fail fallback, configure alternative options for supplicants attempting LAN access. You can configure the switch

to accept or deny access to supplicants or to maintain the access already granted towards supplicants before the RADIUS server timeout. Additionally, you can configure the switch to move supplicants to a specific VLAN if a RADIUS timeout occurs or if the RADIUS server sends an EAP Access-Reject message. [Figure 15 on page 1877](#) shows the topology used for this example. The RADIUS server is connected to the EX4200 switch on access port **ge-0/0/10**. The switch acts as the authenticator Port Access Entity (PAE) and forwards credentials from the supplicant to the user database on the RADIUS server. The switch blocks all traffic and acts as a control gate until the supplicant is authenticated by the authentication server. A supplicant is connected to the switch through interface **ge-0/0/1**.

**Figure 15: Topology for Configuration**



[Table 171 on page 1877](#) describes the components in this topology.

**Table 171: Components of the Topology**

Property	Settings
Switch hardware	EX4200 access switch, 24 Gigabit Ethernet ports: 8 PoE ports.
VLAN names	<b>default</b> VLAN <b>vlan-sf</b> VLAN

Table 171: Components of the Topology (*continued*)

Property	Settings
Supplicant	Supplicant attempting access on interface <b>ge-0/0/1</b>
One RADIUS server	Backend database with an address of <b>10.0.0.100</b> connected to the switch at port <b>ge-0/0/10</b>

In this example, configure interface **ge-0/0/1** to move a supplicant attempting access to the LAN during a RADIUS timeout to another VLAN. A RADIUS timeout prevents the normal exchange of EAP messages that carry information from the RADIUS server to the switch and permit the authentication of a supplicant. The **default** VLAN is configured on interface **ge-0/0/1**. When a RADIUS timeout occurs, supplicants on the interface will be moved from the **default** VLAN to the VLAN named **vlan-sf**.



**NOTE:** For more information about authentication, authorization, and accounting (AAA) services, see [Junos OS System Basics Configuration Guide](#).

### Configuration

**CLI Quick Configuration** To quickly configure server fail fallback on the switch, copy the following commands and paste them into the switch terminal window:

```
[edit protocols dot1x authenticator]
set interface ge-0/0/1 server-fail vlan-name vlan-sf
```

**Step-by-Step Procedure** To configure an interface to divert supplicants to a specific VLAN when a RADIUS timeout occurs (here, the VLAN is **vlan-sf**):

1. Define the VLAN to which supplicants are diverted:

```
[edit protocols dot1x authenticator]
user@switch# set interface ge-0/0/1 server-fail vlan-name vlan-sf
```

**Results** Display the results of the configuration:

```
user@switch> show configuration
interfaces {
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members default;
        }
      }
    }
  }
}
protocols {
  dot1x {
    authenticator {
      authentication-profile-name profile52;
      interface {
        ge-0/0/1.0 {
```

```

server-fail vlan-name vlan-sf;
    }
  }
}

```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Supplicants Are Moved to an Alternative VLAN During a RADIUS Timeout on page 1879](#)

#### *Verifying That the Supplicants Are Moved to an Alternative VLAN During a RADIUS Timeout*

**Purpose** Verify that the interface moves supplicants to an alternative VLAN during a RADIUS timeout.



**NOTE:** On switches running Junos OS for EX Series with support for the Enhanced Layer 2 Software (ELS), the output for the `show vlans` command will contain additional information. If your switch runs software that supports ELS, see [show vlans](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#)

**Action** Display the VLANs configured on the switch; the interface **ge-0/0/1.0** is a member of the **default** VLAN:

```
user@switch> show vlans
Name      Tag      Interfaces
default
          ge-0/0/0.0, ge-0/0/1.0*, ge-0/0/5.0*, ge-0/0/10.0,
          ge-0/0/12.0*, ge-0/0/14.0*, ge-0/0/15.0, ge-0/0/20.0
v2        77
          None
vlan-sf   50
          None
mgmt      me0.0*
```

Display 802.1X protocol information on the switch to view supplicants that are authenticated on interface **ge-0/0/1.0**:

```
user@switch> show dot1x interface brief
802.1X Information:
Interface  Role      State      MAC address      User
ge-0/0/1.0  Authenticator  Authenticated  00:00:00:00:00:01  abc
ge-0/0/10.0 Authenticator  Initialize
ge-0/0/14.0 Authenticator  Connecting
ge-0/0/15.0 Authenticator  Initialize
ge-0/0/20.0 Authenticator  Initialize
```

A RADIUS server timeout occurs. Display the Ethernet switching table to show that the supplicant with the MAC address **00:00:00:00:00:01** previously accessing the LAN through the **default** VLAN is now being learned on the VLAN named **vlan-sf**:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 3 entries, 1 learned
VLAN      MAC address      Type      Age Interfaces
v1        *                Flood     - All-members
vlan-sf   00:00:00:00:00:01 Learn     1:07 ge-0/0/1.0
default   *                Flood     - All-members
```

Display 802.1X protocol information to show that interface **ge-0/0/1.0** is connecting and will open LAN access to supplicants:

```
user@switch> show dot1x interface brief
802.1X Information:
Interface  Role      State      MAC address      User
ge-0/0/1.0  Authenticator  Connecting
ge-0/0/10.0 Authenticator  Initialize
ge-0/0/14.0 Authenticator  Connecting
ge-0/0/15.0 Authenticator  Initialize
ge-0/0/20.0 Authenticator  Initialize
```

**Meaning** The command **show vlans** displays interface **ge-0/0/1.0** as a member of the **default** VLAN. The command **show dot1x interface brief** shows that a supplicant (**abc**) is authenticated on interface **ge-0/0/1.0** and has the MAC address **00:00:00:00:00:01**. A RADIUS server timeout occurs, and the authentication server cannot be reached by the



switch. The command **show-ethernet-switching table** shows that MAC address **00:00:00:00:00:01** is learned on VLAN **vlan-sf**. The supplicant has been moved from the **default** VLAN to the **vlan-sf** VLAN. The supplicant is then connected to the LAN through the VLAN named **vlan-sf**.

**Related Documentation**

- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Configuring Server Fail Fallback \(CLI Procedure\) on page 1921](#)
- [Configuring 802.1X RADIUS Accounting \(CLI Procedure\) on page 1909](#)
- [Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 1910](#)
- [Understanding Server Fail Fallback and Authentication on EX Series Switches on page 1838](#)

## Example: Configuring MAC RADIUS Authentication on an EX Series Switch

To permit hosts that are not 802.1X-enabled to access the LAN, you can configure MAC RADIUS authentication on the switch interfaces to which the non-802.1X-enabled hosts are connected. When MAC RADIUS authentication is configured, the switch will attempt to authenticate the host with the RADIUS server using the host's MAC address.

This example describes how to configure MAC RADIUS authentication for two non-802.1X-enabled hosts:

- [Requirements on page 1881](#)
- [Overview and Topology on page 1882](#)
- [Configuration on page 1884](#)
- [Verification on page 1885](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.3 or later for EX Series switches.
- An EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- A RADIUS authentication server. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you configure MAC RADIUS authentication, be sure you have:

- Configured basic access between the EX Series switch and the RADIUS server. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).
- Performed basic bridging and VLAN configuration on the switch. See the documentation that describes setting up basic bridging and a VLAN for your switch. If you are using a

switch that supports the Enhanced Layer 2 Software (ELS) configuration style, see [“Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch” on page 2281](#). For all other switches, see *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*.



**NOTE:** For more about ELS, see: [“Getting Started with Enhanced Layer 2 Software” on page 3](#)

- Performed basic 802.1X configuration. See [“Configuring 802.1X Interface Settings \(CLI Procedure\)” on page 1908](#).

---

### Overview and Topology

---

IEEE 802.1X Port-Based Network Access Control (PNAC) authenticates and permits devices access to a LAN if the devices can communicate with the switch using the 802.1X protocol (are 802.1X-enabled). To permit non-802.1X-enabled end devices to access the LAN, you can configure MAC RADIUS authentication on the interfaces to which the end devices are connected. When the MAC address of the end device appears on the interface, the switch consults the RADIUS server to check whether it is a permitted MAC address. If the MAC address of the end device is configured as permitted on the RADIUS server, the switch opens LAN access to the end device.

You can configure both MAC RADIUS authentication and 802.1X authentication methods on an interface configured for multiple supplicants. Additionally, if an interface is only connected to a non-802.1X-enabled host, you can enable MAC RADIUS and not enable 802.1X authentication using the **mac-radius restrict** option, and thus avoid the delay that occurs while the switch determines that the device does not respond to EAP messages.

[Figure 16 on page 1883](#) shows the two printers connected to the switch.

Figure 16: Topology for MAC RADIUS Authentication Configuration

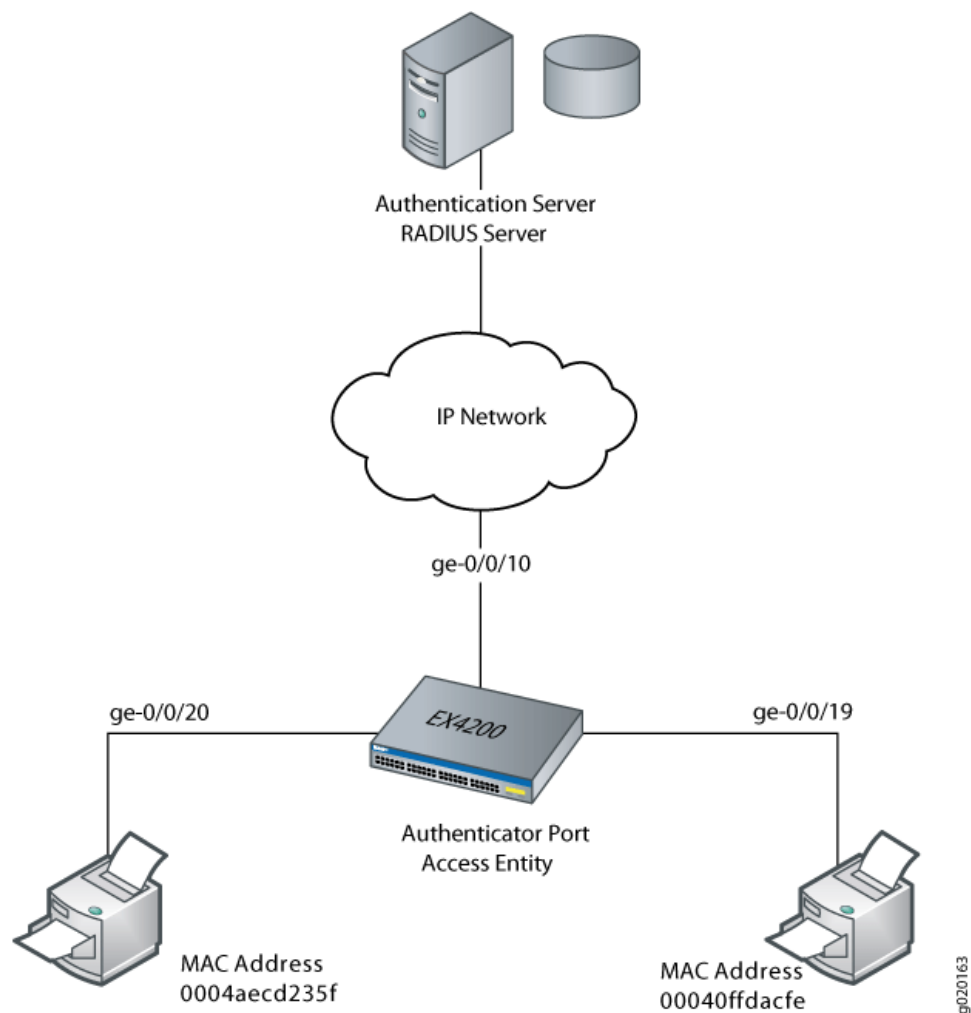


Table 172 on page 1883 shows the components in the example for MAC RADIUS authentication.

Table 172: Components of the MAC RADIUS Authentication Configuration Topology

Property	Settings
Switch hardware	EX4200 ports (ge-0/0/0 through ge-0/0/23)
VLAN name	sales
Connections to printers (no PoE required)	ge-0/0/19, MAC address 00040ffdacfe ge-0/0/20, MAC address 0004aec235f
RADIUS server	Connected to the switch on interface <b>ge-0/0/10</b>

The printer with the MAC address 00040ffdacfe is connected to access interface ge-0/0/19. A second printer with the MAC address 0004aec235f is connected to access interface ge-0/0/20. In this example, both interfaces are configured for MAC RADIUS authentication on the switch, and the MAC addresses (without colons) of both printers are configured on the RADIUS server. Interface ge-0/0/20 is configured to eliminate the normal delay while the switch attempts 802.1X authentication; MAC RADIUS authentication is enabled and 802.1X authentication is disabled using the **mac radius restrict** option.

### Configuration

---

#### CLI Quick Configuration

To quickly configure MAC RADIUS authentication, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols dot1x authenticator interface ge-0/0/19 mac-radius
set protocols dot1x authenticator interface ge-0/0/20 mac-radius restrict
```



**NOTE:** You must also configure the two MAC addresses as usernames and passwords on the RADIUS server, as is done in step 2 of the Step-by-Step Procedure.

---

#### Step-by-Step Procedure

Configure MAC RADIUS authentication on the switch and on the RADIUS server:

1. On the switch, configure the interfaces to which the printers are attached for MAC RADIUS authentication, and configure the **restrict** option on interface **ge-0/0/20**, so that only MAC RADIUS authentication is used:

```
[edit]
user@switch# set protocols dot1x authenticator interface ge-0/0/19 mac-radius
user@switch# set protocols dot1x authenticator interface ge-0/0/20 mac-radius restrict
```

2. On the RADIUS server, configure the MAC addresses **00040ffdacfe** and **0004aec235f** as usernames and passwords:

```
[root@freeradius]#
edit /etc/raddb
vi users
00040ffdacfe Auth-type:=EAP, User-Password = "00040ffdacfe"
0004aec235f Auth-type:=EAP, User-Password = "0004aec235f"
```

#### Results

Display the results of the configuration on the switch:

```
user@switch> show configuration
protocols {
  dot1x {
    authenticator {
      authentication-profile-name profile52;
    }
    interface {
      ge-0/0/19.0 {
        mac-radius;
      }
      ge-0/0/20.0 {
        mac-radius {
```

```
restrict;  
}  
}  
}  
}  
}
```

### Verification

---

Verify that the supplicants are authenticated:

- [Verifying That the Supplicants Are Authenticated on page 1885](#)

#### ***Verifying That the Supplicants Are Authenticated***

**Purpose** After supplicants are configured for MAC RADIUS authentication on the switch and on the RADIUS server, verify that they are authenticated and display the method of authentication:

**Action** Display information about 802.1X-configured interfaces **ge-0/0/19** and **ge-0/0/20**:

```
user@switch> show dot1x interface ge-0/0/19.0 detail
ge-0/0/19.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 1
    Supplicant: user101, 00:04:0f:fd:ac:fe
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: vo11
      Dynamic Filter: match source-dot1q-tag 10 action deny
      Session Reauth interval: 60 seconds
      Reauthentication due in 50 seconds

user@switch> show dot1x interface ge-0/0/20.0 detail
ge-0/0/20.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Enabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 1
    Supplicant: user102, 00:04:ae:cd:23:5f
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: vo11
      Dynamic Filter: match source-dot1q-tag 10 action deny
      Session Reauth interval: 60 seconds
      Reauthentication due in 50 seconds
```

**Meaning** The sample output from the **show dot1x interface detail** command displays the MAC address of the connected end device in the **Supplicant** field. On interface **ge-0/0/19**, the MAC address is **00:04:0f:fd:ac:fe**, which is the MAC address of the first printer configured for MAC RADIUS authentication. The **Authentication method** field displays the authentication method as **MAC Radius**. On interface **ge-0/0/20**, the MAC address is **00:04:ae:cd:23:5f**, which is the MAC address of the second printer configured for MAC

RADIUS authentication. The **Authentication method** field displays the authentication method as **MAC Radius**.

**Related Documentation**

- [Configuring MAC RADIUS Authentication \(CLI Procedure\) on page 1923](#)
- [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)
- [Configuring 802.1X Authentication \(J-Web Procedure\)](#)
- [Understanding Authentication on EX Series Switches on page 1824](#)

### Example: Applying a Firewall Filter to 802.1X-Authenticated Supplicants Using RADIUS Server Attributes on an EX Series Switch

You can use RADIUS server attributes and a port firewall filter to centrally apply terms to multiple supplicants (end devices) connected to an EX Series switch in your enterprise. Terms are applied after a device is successfully authenticated through 802.1X. If the firewall filter configuration is modified after end devices are authenticated using the 802.1X authentication, then the established 802.1X authentication session must be terminated and re-established for the firewall filter changes to take effect.

EX Series switches support port firewall filters. Port firewall filters are configured on a single EX Series switch, but in order for them to operate throughout an enterprise, they have to be configured on multiple switches. To reduce the need to configure the same port firewall filter on multiple switches, you can instead apply the filter centrally on the RADIUS server using RADIUS server attributes.

The following example uses FreeRADIUS to apply a port firewall filter on a RADIUS server. For specifics on configuring your server, consult the documentation that was included with your RADIUS server.

This example describes how to configure a port firewall filter with terms, create counters to count packets for the supplicants, apply the filter to user profiles on the RADIUS server, and display the counters to verify the configuration:

- [Requirements on page 1887](#)
- [Overview and Topology on page 1888](#)
- [Configuring the Port Firewall Filter and Counters on page 1890](#)
- [Applying the Port Firewall Filter to the Supplicant User Profiles on the RADIUS Server on page 1892](#)
- [Verification on page 1892](#)

#### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.3 or later for EX Series switches
- One EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.

- One RADIUS authentication server. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you connect the server to the switch, be sure you have:

- Set up a connection between the switch and the RADIUS server. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch”](#) on page 1843.
- Configured 802.1X authentication on the switch, with the authentication mode for interface `ge-0/0/2` set to **multiple**. See [“Configuring 802.1X Interface Settings \(CLI Procedure\)”](#) on page 1908 and [“Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch”](#) on page 1852.
- Configured users on the RADIUS authentication server (in this example, the user profiles for Supplicant 1 and Supplicant 2 in the topology are modified on the RADIUS server).

---

### Overview and Topology

When the 802.1X configuration on an interface is set to **multiple** supplicant mode, you can apply a single port firewall filter configured through the Junos OS CLI on the EX Series switch to any number of end devices (supplicants) on one interface by adding the filter centrally to the RADIUS server. Only a single filter can be applied to an interface; however, the filter can contain multiple terms for separate end devices.

For more information about firewall filters, see [“Firewall Filters for EX Series Switches Overview”](#) on page 4696.

RADIUS server attributes are applied to end devices after the devices are successfully authenticated using 802.1X. To authenticate an end device, the switch forwards the end device's credentials to the RADIUS server. The RADIUS server matches the credentials against preconfigured information about the supplicant located in the supplicant's user profile on the RADIUS server. If a match is found, the RADIUS server instructs the switch to open an interface to the end device. Traffic then flows from and to the end device on the LAN. Further instructions configured in the port firewall filter and added to the end device's user profile using a RADIUS server attribute further define the access that the end device is granted. Filtering terms configured in the port firewall filter are applied to the end device after 802.1X authentication is complete.



**NOTE:** If you modify the port firewall filter after an end device is successfully authenticated using 802.1X, you must terminate and re-establish the 802.1X authentication session for the firewall filter configuration changes to be effective.

---

[Figure 17 on page 1889](#) shows the topology used for this example. The RADIUS server is connected to an EX4200 switch on access port `ge-0/0/10`. Two end devices (supplicants) are accessing the LAN on interface `ge-0/0/2`. Supplicant 1 has the MAC address `00:50:8b:6f:60:3a`. Supplicant 2 has the MAC address `00:50:8b:6f:60:3b`.



Figure 17: Topology for Firewall Filter and RADIUS Server Attributes Configuration




Table 173 on page 1889 describes the components in this topology.

Table 173: Components of the Firewall Filter and RADIUS Server Attributes Topology

Property	Settings
Switch hardware	EX4200 access switch, 24 Gigabit Ethernet ports, 8 PoE ports.
One RADIUS server	Backend database with the address 10.0.0.100 connected to the switch at port ge-0/0/10.
802.1X supplicants connected to the switch on interface ge-0/0/2	<ul style="list-style-type: none"> <li>Supplicant 1 has MAC address 00:50:8b:6f:60:3a.</li> <li>Supplicant 2 has MAC address 00:50:8b:6f:60:3b.</li> </ul>

Table 173: Components of the Firewall Filter and RADIUS Server Attributes Topology (*continued*)

Property	Settings
Port firewall filter to be applied on the RADIUS server	<b>filter1</b>
Counters	<b>counter1</b> counts packets from Supplicant 1, and <b>counter2</b> counts packets from Supplicant 2.
Policer	<b>policer p1</b>
User profiles on the RADIUS server	<ul style="list-style-type: none"> <li>Supplicant 1 has the user profile <b>supplicant1</b>.</li> <li>Supplicant 2 has the user profile <b>supplicant2</b>.</li> </ul>

In this example, you configure a port firewall filter named **filter1**. The filter contains terms that will be applied to the end devices based on the MAC addresses of the end devices. When you configure the filter, you also configure the counters **counter1** and **counter2**. Packets from each end device are counted, which helps you verify that the configuration is working. Policer **policer p1** limits the traffic rate based on the values for **exceeding** and **discard** parameters. Then, you check to see that the RADIUS server attribute is available on the RADIUS server and apply the filter to the user profiles of each end device on the RADIUS server. Finally, you verify the configuration by displaying output for the two counters.



**NOTE:** For more information about authentication, authorization, and accounting (AAA) services, see the [Junos OS System Basics Configuration Guide](#).

### Configuring the Port Firewall Filter and Counters

#### CLI Quick Configuration

To quickly configure a port firewall filter with terms for Supplicant 1 and Supplicant 2 and create parallel counters for each supplicant, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family ethernet-switching filter filter1 term supplicant1 from source-mac-address 00:50:8b:6f:60:3a
set firewall family ethernet-switching filter filter1 term supplicant2 from source-mac-address 00:50:8b:6f:60:3b
set firewall policer p1 if-exceeding bandwidth-limit 1m
set firewall policer p1 if-exceeding burst-size-limit 1k
set firewall policer p1 then discard
set firewall family ethernet-switching filter filter1 term supplicant1 then count counter1
set firewall family ethernet-switching filter filter1 term supplicant1 then policer p1
set firewall family ethernet-switching filter filter1 term supplicant2 then count counter2
```

#### Step-by-Step Procedure

To configure a port firewall filter and counters on the switch:

1. Configure a port firewall filter (here, **filter1**) with terms for each end device based upon the MAC address of each end device:

```
[edit firewall family ethernet-switching]
user@switch# set filter filter1 term supplicant1 from source-mac-address 00:50:8b:6f:60:3a
user@switch# set filter filter1 term supplicant2 from source-mac-address 00:50:8b:6f:60:3b
```

2. Set policer definition:

```
user@switch# show policer p1 |display set
set firewall policer p1 if-exceeding bandwidth-limit 1m
set firewall policer p1 if-exceeding burst-size-limit 1k
set firewall policer p1 then discard
```

3. Create two counters that will count packets for each end device and a policer which limits the traffic rate:

```
[edit firewall family ethernet-switching]
user@switch# set filter filter1 term supplicant1 then count counter1
user@switch# set filter filter1 term supplicant1 then policer p1
user@switch# set filter filter1 term supplicant2 then count counter2
```

**Results** Display the results of the configuration:

```
user@switch> show configuration
firewall {
  family ethernet-switching {
    filter filter1 {
      term supplicant1 {
        from {
          source-mac-address {
            00:50:8b:6f:60:3a;
          }
        }
        then count counter1;
        then policer p1;
      }
      term supplicant2 {
        from {
          source-mac-address {
            00:50:8b:6f:60:3b;
          }
        }
        then count counter2;
      }
    }
  }
}
policer p1 {
  if-exceeding {
    bandwidth-limit 1m;
    burst-size-limit 1k;
  }
  then discard;
}
```

## Applying the Port Firewall Filter to the Supplicant User Profiles on the RADIUS Server

---

**Step-by-Step Procedure** To verify that the RADIUS server attribute **Filter-ID** is on the RADIUS server and to apply the filter to the user profiles:

1. Display the dictionary **dictionary.rfc2865** on the RADIUS server, and verify that the attribute **Filter-ID** is in the dictionary:

```
[root@freeradius]# cd usr/share/freeradius/dictionary.rfc2865
```

2. Close the dictionary file.

3. Display the local user profiles of the end devices to which you want to apply the filter (here, the user profiles are called **supplicant1** and **supplicant2**):

```
[root@freeradius]# cat /usr/local/etc/raddb/users
```

The output shows:

```
supplicant1 Auth-Type := EAP, User-Password == "supplicant1"
    Tunnel-Type = VLAN,
    Tunnel-Medium-Type = IEEE-802,
    Tunnel-Private-Group-Id = "1005"
```

```
supplicant2 Auth-Type := EAP, User-Password == "supplicant2"
    Tunnel-Type = VLAN,
    Tunnel-Medium-Type = IEEE-802,
    Tunnel-Private-Group-Id = "1005"
```

4. Apply the filter to both user profiles by adding the line **Filter-Id = "filter1"** to each profile, and then close the file:

```
[root@freeradius]# cat /usr/local/etc/raddb/users
```

After you paste the line into the files, the files look like this:

```
supplicant1 Auth-Type := EAP, User-Password == "supplicant1"
    Tunnel-Type = VLAN,
    Tunnel-Medium-Type = IEEE-802,
    Tunnel-Private-Group-Id = "1005",
    Filter-Id = "filter1"
```

```
supplicant2 Auth-Type := EAP, User-Password == "supplicant2"
    Tunnel-Type = VLAN,
    Tunnel-Medium-Type = IEEE-802,
    Tunnel-Private-Group-Id = "1005",
    Filter-Id = "filter1"
```

## Verification

---

### Verifying That the Filter Has Been Applied to the Supplicants

**Purpose** After the end devices are authenticated, verify that the filter has been configured on the switch and added to each end device's user profile on the RADIUS server:

**Action** Display information about firewall filter **filter1**:

```
user@switch> show firewall filter filter1
Filter: filter1
Counters:
Name                               Bytes      Packets
counter1                           128         2
counter2                            64         1
```

**Meaning** The output of the command **show firewall filter filter1** displays **counter1** and **counter2**. Packets from Supplicant 1 are counted using **counter1**, and packets from Supplicant 2 are counted using **counter2**. The output displays packets incrementing for both counters. The filter has been applied to both end devices.

- Related Documentation**
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
  - [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
  - [Configuring 802.1X RADIUS Accounting \(CLI Procedure\) on page 1909](#)
  - [Understanding Authentication on EX Series Switches on page 1824](#)
  - [Understanding 802.1X and VSAs on EX Series Switches on page 1837](#)

### Example: Applying Firewall Filters to Multiple Supplicants on Interfaces Enabled for 802.1X or MAC RADIUS Authentication



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [“Example: Applying Firewall Filters to Multiple Supplicants on Interfaces Enabled for 802.1X or MAC RADIUS Authentication” on page 4798](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

On EX Series switches, firewall filters that you apply to interfaces enabled for 802.1X or MAC RADIUS authentication are dynamically combined with the per-user policies sent to the switch from the RADIUS server. The switch uses internal logic to dynamically combine the interface firewall filter with the user policies from the RADIUS server and create an individualized policy for each of the multiple users or nonresponsive hosts that are authenticated on the interface.

This example describes how dynamic firewall filters are created for multiple supplicants on an 802.1X-enabled interface (the same principles shown in this example apply to interfaces enabled for MAC RADIUS authentication):

- [Requirements on page 1894](#)
- [Overview and Topology on page 1894](#)
- [Configuration on page 1896](#)
- [Verification on page 1898](#)

---

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.2 or later for EX Series switches
- One EX4300 switch
- One RADIUS authentication server. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you apply firewall filters to an interface for use with multiple supplicants, be sure you have:

- Set up a connection between the switch and the RADIUS server. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).
- Configured 802.1X authentication on the switch, with the authentication mode for the interface ge-0/0/2 set to **multiple**. See [“Configuring 802.1X Interface Settings \(CLI Procedure\)” on page 1908](#) and [“Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch” on page 1852](#).
- Configured users on the RADIUS authentication server.


---

## Overview and Topology

When the 802.1X configuration on an interface is set to multiple supplicant mode, the system dynamically combines the interface firewall filter with the user policies sent to the switch from the RADIUS server during authentication and creates separate terms for each user. Because there are separate terms for each user authenticated on the interface, you can, as shown in this example, use counters to view the activities of individual users that are authenticated on the same interface.

When a new user (or a nonresponsive host) is authenticated on an interface, the system adds a term to the firewall filter associated with the interface, and the term (policy) for each user is associated with the MAC address of the user. The term for each user is based on the user-specific filters set on the RADIUS server and the filters configured on the interface. For example, as shown in [Figure 18 on page 1895](#), when User 1 is authenticated by the EX Series switch, the system creates the firewall filter **dynamic-filter-example**. When User 2 is authenticated, another term is added to the firewall filter, and so on.

Figure 18: Conceptual Model: Dynamic Filter Updated for Each New User




This is a conceptual model of the internal process—you cannot access or view the dynamic filter.



**NOTE:** If the firewall filter on the interface is modified after the user (or nonresponsive host) is authenticated, the modifications are not reflected in the dynamic filter unless the user is reauthenticated.

In this example, you configure a firewall filter to count the requests made by each endpoint authenticated on interface ge-0/0/2 to the file server, which is located on subnet 192.0.2.16/28, and set policer definitions to rate-limit the traffic. [Figure 19 on page 1896](#) shows the network topology for this example.

Figure 19: Multiple Supplicants on an 802.1X-Enabled Interface Connecting to a File Server



### Configuration

#### Configuring Firewall Filters on Interfaces with Multiple Supplicants

##### CLI Quick Configuration

To quickly configure firewall filters for multiple supplicants on an 802.1X-enabled interface copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family ethernet-switching filter filter1 term term1 from ip-destination-address 192.0.2.16/28
set firewall family ethernet-switching filter filter1 term term2 from ip-destination-address 192.0.2.16/28
set firewall policer p1 if-exceeding bandwidth-limit 1m
set firewall policer p1 if-exceeding burst-size-limit 1500
set firewall policer p1 then discard
set firewall family ethernet-switching filter filter1 term term1 then count counter1
set firewall family ethernet-switching filter filter1 term term2 then policer p1
```

##### Step-by-Step Procedure

To configure firewall filters on an interface enabled for multiple supplicants:

1. Set the policer definition:

```
user@switch# show policer p1 |display set
set firewall policer p1 if-exceeding bandwidth-limit 1m
set firewall policer p1 if-exceeding burst-size-limit 1500
set firewall policer p1 then discard
```



2. Configure a firewall filter to count packets from each user and a policer that limits the traffic rate. As each new user is authenticated on the multiple supplicant interface, this filter term will be included in the dynamically created term for the user:

```
[edit firewall family ethernet-switching]
user@switch# set filter filter1 term term1 from ip-destination-address 192.0.2.16/28
user@switch# set filter filter1 term term2 from ip-destination-address 192.0.2.16/28
user@switch# set filter filter1 term term1 then count counter1
user@switch# set filter filter1 term term2 then policer p1
```

**Results** Check the results of the configuration:

```
user@switch> show configuration
```

```
firewall {
  family ethernet-switching {
    filter filter1 {
      term term1 {
        from {
          ip-destination-address {
            192.0.2.16/28;
          }
        }
        then count counter1;
      }
      term term2 {
        from {
          ip-destination-address {
            192.0.2.16/28;
          }
        }
        then policer p1;
      }
    }
  }
  policer p1 {
    if-exceeding {
      bandwidth-limit 1m;
      burst-size-limit 1500;
    }
    then discard;
  }
}
protocols {
  dot1x {
    authenticator
    interface ge-0/0/2 {
      supplicant multiple;
    }
  }
}
```

## Verification

---

### Verifying Firewall Filters on Interfaces with Multiple Supplicants

- Purpose** Verify that firewall filters are functioning on the interface with multiple supplicants.
- Action**
1. Check the results with one user authenticated on the interface. In this case, User 1 is authenticated on ge-0/0/2:  

```
user@switch> show dot1x firewall
Filter: dot1x_ge-0/0/2
Counters
counter1_dot1x_ge-0/0/2_user1 100
```
  2. When a second user, User 2, is authenticated on the same interface, ge-0/0/2, you can verify that the filter includes the results for both of the users authenticated on the interface:  

```
user@switch> show dot1x firewall
Filter: dot1x-filter-ge-0/0/0
Counters
counter1_dot1x_ge-0/0/2_user1 100
counter1_dot1x_ge-0/0/2_user2 400
```
- Meaning** The results displayed by the **show dot1x firewall** command output reflect the dynamic filter created with the authentication of each new user. User 1 accessed the file server located at the specified destination address 100 times, while User 2 accessed the same file server 400 times.
- Related Documentation**
- [Example: Applying a Firewall Filter to 802.1X-Authenticated Supplicants Using RADIUS Server Attributes on an EX Series Switch on page 1887](#)
  - [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
  - [Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 1910](#)

## Example: Setting Up Captive Portal Authentication on an EX Series Switch



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Setting Up Captive Portal Authentication on an EX Series Switch*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

You can set up captive portal authentication (hereafter referred to as captive portal) on a switch to redirect Web browser requests to a login page that requires the user to input a username and password. Upon successful authentication, the user is allowed to continue with the original page request and subsequent access to the network.

This example describes how to set up captive portal on an EX4300 switch:

- [Requirements on page 1899](#)
- [Overview and Topology on page 1899](#)
- [Configuration on page 1899](#)
- [Verification on page 1901](#)
- [Troubleshooting on page 1902](#)

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.2X50 or later for EX Series switches
- An EX4300 Series switch

Before you begin, be sure you have:

- Performed basic bridging and VLAN configuration on the switch. See [“Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch” on page 2281](#).
- Generated an SSL certificate and installed it on the switch. See [“Generating SSL Certificates to Be Used for Secure Web Access” on page 532](#).
- Configured basic access between the EX Series switch and the RADIUS server. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).
- Designed your captive portal login page. See [“Designing a Captive Portal Authentication Login Page on an EX Series Switch” on page 1928](#).

## Overview and Topology

This example shows the configuration required on the switch to enable captive portal on an interface. To permit a printer connected to the captive portal interface to access the LAN, add its MAC address to the authentication whitelist and assign to a VLAN, vlan1. The MAC addresses on this list are permitted access on the interface without captive portal authentication.

The topology for this example consists of one EX4300 switch connected to a RADIUS authentication server. One interface on the switch is configured for captive portal. In this example, the interface is configured in multiple supplicant mode.

## Configuration

To configure captive portal on your switch:

### CLI Quick Configuration

To quickly configure captive portal on the switch after completing the tasks in the Requirements section, copy the following commands and paste them into the switch terminal window:

```
[edit]
set system services web-management https local-certificate my-signed-cert
set services captive-portal secure-authentication https
set services captive-portal interface ge-0/0/10.0 supplicant multiple
set switch-options authentication-whitelist 00:10:12:e0:28:22 vlan-assignment vlan1
```

```
set custom-options post-authentication-url http://www.my-home-page.com
```

### Step-by-Step Procedure

To configure captive portal on the switch:

1. To create a secure channel for Web access to the switch, configure captive portal for HTTPS:
  - a. Associate the security certificate with the Web server and enable HTTPS on the switch:

```
[edit]
user@switch# set system services web-management https local-certificate
my-signed-cert
```



**NOTE:** You can enable HTTP instead of HTTPS, but we recommend HTTPS for security purposes.

- b. Configure captive portal to use HTTPS:

```
[edit]
user@switch# set services captive-portal secure-authentication https
```

2. Enable an interface for captive portal:

```
[edit]
user@switch# set services captive-portal interface ge-0/0/10 supplicant multiple
```

3. (Optional) Allow specific clients to bypass captive portal authentication:



**NOTE:** If the client is already attached to the switch, you must clear its MAC address from the captive portal authentication by using the `clear captive-portal mac-address mac-address` command after adding its MAC address to the whitelist. Otherwise the new entry for the MAC address will not be added to the Ethernet switching table and the authentication bypass will not be allowed.

```
[edit]
user@switch# set switch-options authentication-whitelist 00:10:12:e0:28:22
vlan-assignment vlan1
```



**NOTE:** Optionally, you can use `set switch-options authentication-whitelist 00:10:12:e0:28:22 vlan-assignment vlan1 interface ge-0/0/10.0` to limit the scope to the interface.

4. (Optional) To redirect clients to a specified page rather than the page they originally requested, configure the post-authentication URL:

```
[edit services captive-portal]
user@switch# set custom-options post-authentication-url http://www.my-home-page.com
```

**Results** Display the results of the configuration:

```

[edit]
user@switch# show
system {
  services {
    web-management {
      https {
        local-certificate my-signed-cert;
      }
    }
  }
}
security {
  certificates {
    local {
      my-signed-cert {
        "-----BEGIN RSA PRIVATE KEY-----\nMIICXwIBAAKBgQDk8sUggnXdDUMr7T
vLv63yJq/LRpDASfIDZlX3z9ZDe1Kfk5C9\nr/tkyvzv
...
Pt5YmvWDoGo0mSjoE/liH0BqYdh9YGqv3T2IEUfflSTQQHEOShS0ogWDHF\
nnyOb1O/vQtjk20X9NVQg JHBwidssY9eRp\n-----END CERTIFICATE-----\n";
        ## SECRET-DATA
      }
    }
  }
}
services {
  captive-portal {
    interface {
      ge-0/0/10.0 {
        supplicant multiple;
      }
    }
    secure-authentication https;
    custom-options {
      post-authentication-url http://www.my-home-page.com;
    }
  }
}
switch-options {
  authentication-whitelist {
    00:10:12:e0:28:22/48 {
      vlan-assignment vlan1;
    }
  }
}
}

```

### Verification

To confirm that captive portal authentication is configured and working properly, perform these tasks:

- [Verifying That Captive Portal Is Enabled on the Interface on page 1902](#)
- [Verify That Captive Portal Is Working Correctly on page 1902](#)

***Verifying That Captive Portal Is Enabled on the Interface***

**Purpose** Verify that captive portal is configured on the interface ge-0/0/10.

**Action** Use the operational mode command **show captive-portal interface *interface-name* detail**:

```
user@switch> show captive-portal interface ge-0/0/10.0 detail
ge-0/0/10.0
  Supplicant mode: Multiple
  Number of retries: 3
  Quiet period: 60 seconds
  Configured CP session timeout: 3600 seconds
  Server timeout: 15 seconds
```

**Meaning** The output confirms that captive portal is configured on the interface ge-0/0/10, with the default settings for number of retries, quiet period, CP session timeout, and server timeout.

***Verify That Captive Portal Is Working Correctly***

**Purpose** Verify that captive portal is working on the switch.

**Action** Connect a client to the interface ge-0/0/10. From the client, open a Web browser and request a webpage. The captive portal login page that you designed should be displayed. After you enter your login information and are authenticated against the RADIUS server, the Web browser should display either the page you requested or the post-authentication URL that you configured.

---

**Troubleshooting**

To troubleshoot captive portal, perform these tasks:

- [Troubleshooting Captive Portal on page 1902](#)

***Troubleshooting Captive Portal***

**Problem** The switch does not return the captive portal login page when a user connected to a captive portal interface on the switch requests a webpage.

**Solution** You can examine the ARP, DHCP, HTTPS, and DNS counters—if one or more of these counters are not incrementing, this provides an indication of where the problem lies. For example, if the client cannot get an IP address, you might check the switch interface to determine whether the DHCP counter is incrementing—if the counter increments, the DHCP packet was received by the switch.

```
user@switch> show captive-portal firewall ge-0/0/10.0
ge-0/0/10.0
  Filter name: dot1x_ge-0/0/10
  Counters:
  Name                               Bytes      Packets
  dot1x_ge-0/0/10_CP_arp              7616         119
  dot1x_ge-0/0/10_CP_dhcp              0             0
  dot1x_ge-0/0/10_CP_http              0             0
```

dot1x_ge-0/0/10_CP_https	0	0
dot1x_ge-0/0/10_CP_t_dns	0	0
dot1x_ge-0/0/10_CP_u_dns	0	0

**Related  
Documentation**

- [Configuring Captive Portal Authentication \(CLI Procedure\) on page 1926](#)
- [Designing a Captive Portal Authentication Login Page on an EX Series Switch on page 1928](#)

## Example: Configuring Fallback Options on EX Series Switches for EAP-TTLS Authentication and Odyssey Access Clients

For 802.1X user authentication, EX Series switches support RADIUS authentication servers that are using Extensible Authentication Protocol–Tunneled TLS (EAP-TTLS) to authenticate Odyssey Access Client (OAC) supplicants. OAC networking software runs on endpoint computers (desktop, laptop, or notepad computers and supported wireless devices) and provides secure access to both wired and wireless networks.

This example describes how to configure an 802.1X-enabled interface on the switch to provide fallback support for OAC users who have entered incorrect login credentials:

- [Requirements on page 1903](#)
- [Overview and Topology on page 1904](#)
- [Configuration on page 1905](#)
- [Verification on page 1906](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.2 or later for EX Series switches
- One EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- One RADIUS authentication server that supports 802.1X. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.
- One OAC end device acting as a supplicant.

Before you begin configuring the fallback option, ensure that you have:

- Set up a connection between the switch and the RADIUS server. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).
- Configured EAP-TTLS on the server. See your RADIUS server documentation.
- Configured users on the RADIUS server. See your RADIUS server documentation.

## Overview and Topology

---

OAC is networking software that runs on endpoint computers (desktop, laptop, or notepad) and supported wireless devices. OAC provides full support for EAP, which is required for secure wireless LAN access.

In this topology, OAC is deployed with an 802.1X-enabled switch and a RADIUS server. The switch functions as an enforcement point in the network security architecture. This topology:

- Ensures that only authorized users can connect.
- Maintains privacy of login credentials.
- Maintains data privacy over the wireless link.

This example includes the configuration of a server-reject VLAN on the switch, which can be used to prevent accidental lockout for users who have entered incorrect login credentials. These users can be given limited LAN access.

However, this fallback configuration is complicated by the fact that the OAC supplicant and RADIUS server are using EAP-TTLS. EAP-TTLS creates a secure encrypted tunnel between the server and the end device to complete the authentication process. When the user enters an incorrect login, the RADIUS server sends EAP failure messages directly to the client through this tunnel. The EAP failure message causes the client to restart the authentication procedure, so that the switch's 802.1X authentication process tears down the session that was established with the switch using the server-reject VLAN. You can enable the remedial connection to continue by configuring:

- **eapol-block**—Enable the EAPoL block timer on the 802.1X interface that is configured to belong to the server-reject VLAN. The block timer causes the authentication port access entity to ignore EAP start messages from the client, attempting to restart the authentication procedure.



**NOTE:** The EAPoL block timer is triggered only after the retries on the 802.1X interface have been exhausted. You can configure retries to specify the number of times the switch attempts to authenticate the port after an initial failure. The default is three retries.

- **block-interval**—Configure the amount of time that you want the EAPoL block timer to continue to ignore EAP start messages. If you do not configure the block interval, the EAPoL block timer defaults to 120 seconds.

When the 802.1X interface ignores the EAP start messages from the client, the switch allows the existing remedial session that was established through the server-reject VLAN to remain open.

These configuration options apply to **single**, **single-secure**, and **multiple** supplicant authentication modes. In this example, the 802.1X interface is configured in single-supplicant mode.



Figure 20 on page 1905 shows an EX Series switch connecting an OAC end device to a RADIUS server, and indicates the protocols being used to connect the network entities.

**Figure 20: EX Series Switch Connecting OAC to RADIUS Server Using EAP-TTLS Authentication**




Table 174 on page 1905 describes the components in this OAC deployment:

**Table 174: Components of the OAC Deployment**

Property	Settings
Switch hardware	EX Series switch
VLANs	<b>default</b> <b>server-reject-vlan:</b> VLAN name is <b>remedial</b> and VLAN ID is <b>700</b>
802.1X interface	<b>ge-0/0/8</b>
OAC supplicant	EAP-TTLS
One RADIUS authentication server	EAP-TTLS

### Configuration

To configure fallback options for EAP-TTLS and OAC supplicants, perform this task:

#### CLI Quick Configuration

To quickly configure the fallback options for EAP-TTLS and OAC supplicants, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans remedial vlan-id 700
set protocols dot1x authenticator interface ge-0/0/8 retries 4
set protocols dot1x authenticator interface ge-0/0/8 server-reject-vlan remedial
set protocols dot1x authenticator interface ge-0/0/8 server-reject-vlan eapol-block
set protocols dot1x authenticator interface ge-0/0/8 server-reject-vlan block-interval 130
```

**Step-by-Step Procedure** To configure the fallback options for EAP-TTLS and OAC supplicants:

**TIP:** In this example, the switch has only one server-reject VLAN. Therefore, the configuration specifies `eapol-block` and `block-interval` directly after `server-reject-vlan`. However, if you have configured multiple VLANs on the switch, you should include the VLAN name or VLAN ID directly after `server-reject-vlan` to indicate which VLAN is being modified.

1. Configure a VLAN that will function as the server-reject VLAN to provide limited LAN access for users who have entered incorrect login credentials:  

```
[edit]
user@switch# set vlans remedial vlan-id 700
```
2. Configure the number of times for the client to be prompted for username and password before an incorrect login is directed to the server-reject VLAN:  

```
[edit protocols dot1x authenticator interface ge-0/0/8]
user@switch# set retries 4
```
3. Configure the 802.1X authenticator interface to use the server-reject VLAN as a fallback for incorrect logins:  

```
[edit protocols dot1x authenticator interface ge-0/0/8]
user@switch# set server-reject-vlan remedial
```
4. Enable the EAPoL block timer on the 802.1X interface that is configured to belong to the server-reject VLAN.  

```
[edit protocols dot1x authenticator interface ge-0/0/8]
user@switch# set server-reject-vlan eapol-block
```
5. Configure the amount of time for the EAPoL block to remain in effect:  

```
[edit protocols dot1x authenticator interface ge-0/0/8]
user@switch# set server-reject-vlan block-interval 130
```

**Results**

Check the results of the configuration:

```
user@switch> show configuration
protocols {
  dot1x {
    authenticator {
      interface {
        ge-0/0/8.0 {
          supplicant single;
          retries 4;
          server-reject-vlan remedial block-interval 130 eapol-block;
        }
      }
    }
  }
}
```

**Verification**

To confirm that the configuration and the fallback options are working correctly, perform this task:

- [Verifying the Configuration of the 802.1X Interface on page 1907](#)

### *Verifying the Configuration of the 802.1X Interface*

**Purpose** Verify that the 802.1X interface is configured with the desired options:

**Action** user@switch> **show dot1x** interface ge-0/0/8.0 detail  
 ge-0/0/8.0  
 Role: Authenticator  
 Administrative state: Auto  
 Supplicant mode: Single  
 Number of retries: 4  
 Quiet period: 60 seconds  
 Transmit period: 30 seconds  
 Mac Radius: Disabled  
 Mac Radius Restrict: Disabled  
 Reauthentication: Enabled  
 Configured Reauthentication interval: 120 seconds  
 Supplicant timeout: 30 seconds  
 Server timeout: 30 seconds  
 Maximum EAPoL requests: 2  
 Guest VLAN member: guest  
 Number of connected supplicants: 1  
 Supplicant: tem, 2A:92:E6:F2:00:00  
 Operational state: Authenticated  
 Backend Authentication state: Idle  
 Authentication method: Radius  
 Authenticated VLAN: remedial  
 Session Reauth interval: 120 seconds  
 Reauthentication due in 68 seconds

**Meaning** The **show dot1x ge-0/0/8 detail** output shows that the **ge-0/0/8** interface is in the **Authenticated** state and that it is using the **remedial** VLAN.

**Related Documentation**

- [Understanding Authentication on EX Series Switches on page 1824](#)

## Configuration Tasks

- [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)
- [Configuring 802.1X RADIUS Accounting \(CLI Procedure\) on page 1909](#)
- [Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 1910](#)
- [Configuring LLDP \(CLI Procedure\) on page 1913](#)
- [Configuring LLDP-MED \(CLI Procedure\) on page 1917](#)
- [VSA Match Conditions and Actions on page 1918](#)
- [Configuring Server Fail Fallback \(CLI Procedure\) on page 1921](#)
- [Configuring MAC RADIUS Authentication \(CLI Procedure\) on page 1923](#)
- [Configuring Static MAC Bypass of Authentication \(CLI Procedure\) on page 1924](#)
- [Specifying RADIUS Server Connections on an EX Series Switch \(CLI Procedure\) on page 1925](#)
- [Configuring Captive Portal Authentication \(CLI Procedure\) on page 1926](#)

- [Designing a Captive Portal Authentication Login Page on an EX Series Switch](#) on page 1928
- [Controlling Authentication Session Timeouts \(CLI Procedure\)](#) on page 1930

## Configuring 802.1X Interface Settings (CLI Procedure)

IEEE 802.1X authentication provides network edge security, protecting Ethernet LANs from unauthorized user access by blocking all traffic to and from a supplicant (client) at the interface until the supplicant's credentials are presented and matched on the *authentication server* (a RADIUS server). When the supplicant is authenticated, the switch stops blocking access and opens the interface to the supplicant.



### NOTE:

- You can also specify an 802.1X exclusion list to specify supplicants that can bypass authentication and be automatically connected to the LAN. See [“Configuring Static MAC Bypass of Authentication \(CLI Procedure\)”](#) on page 1924.
- You cannot configure 802.1X user authentication on interfaces that have been enabled for Q-in-Q tunneling.
- You cannot configure 802.1X user authentication on redundant trunk groups (RTGs). For more information on RTGs, see [“Understanding Redundant Trunk Links”](#) on page 2276.

Before you begin, specify the RADIUS server or servers to be used as the authentication server. See [“Specifying RADIUS Server Connections on an EX Series Switch \(CLI Procedure\)”](#) on page 1925.

To configure 802.1X on an interface:

1. Configure the supplicant mode as **single** (authenticates the first supplicant), **single-secure** (authenticates only one supplicant), or **multiple** (authenticates multiple supplicants):

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5 supplicant multiple
```

2. Enable reauthentication and specify the reauthentication interval:

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5/0 reauthentication interval 5
```

3. Configure the interface timeout value for the response from the supplicant:

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5 supplicant-timeout 5
```

4. Configure the timeout for the interface before it resends an authentication request to the RADIUS server:

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5 server-timeout 5
```

5. Configure how long, in seconds, the interface waits before retransmitting the initial EAPOL PDUs to the supplicant:

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5 transmit-period 60
```

6. Configure the maximum number of times an EAPOL request packet is retransmitted to the supplicant before the authentication session times out:

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5 maximum-requests 5
```

7. Configure the number of times the switch attempts to authenticate the port after an initial failure. The port remains in a wait state during the quiet period after the authentication attempt.

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5 retries 1
```



**NOTE:** This setting specifies the number of tries before the switch puts the interface in a “HELD” state.

#### Related Documentation

- [Configuring 802.1X Authentication \(J-Web Procedure\)](#)
- [Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch](#)
- [Monitoring 802.1X Authentication on page 1993](#)
- [Verifying 802.1X Authentication on page 1994](#)
- [Configuring LLDP \(CLI Procedure\) on page 1913](#)
- [Understanding Authentication on EX Series Switches on page 1824](#)

## Configuring 802.1X RADIUS Accounting (CLI Procedure)

RADIUS accounting permits statistical data about users logging onto or off a LAN to be collected and sent to a RADIUS accounting server. The statistical data gathered can be used for general network monitoring, to analyze and track usage patterns, or to bill a user based upon the amount of time or type of services accessed.

To configure basic RADIUS accounting using the CLI:

1. Specify the accounting servers to which the switch will forward accounting statistics:

```
[edit access]
user@switch# set profile profile1 radius accounting-server [122.69.1.250 122.69.1.252]
```

2. Define the RADIUS accounting servers:

```
[edit access]
user@switch# set radius-server 122.69.1.250 secret juniper
user@switch# set radius-server 122.69.1.252 secret juniper1
```

3. Enable accounting for an access profile:

```
[edit access]
user@switch# set profile profile1 accounting (Access Profile)
```

4. Configure the RADIUS servers to use while sending accounting messages and updates:

```
[edit access]
user@switch# set profile profile1 accounting order radius
```

5. Configure the statistics to be collected on the switch and forwarded to the accounting server:

```
[edit access]
```

```
user@switch# set profile profile1 accounting accounting-stop-on-access-deny
user@switch# set profile profile1 accounting accounting-stop-on-failure
```

6. Display accounting statistics collected on the switch:

```
user@switch> show network-access aaa statistics accounting
Accounting module statistics
  Requests received: 1
  Accounting Response failures: 0
  Accounting Response Success: 1
  Requests timedout: 0
```

7. Open an accounting log on the RADIUS accounting server using the server's address, and view accounting statistics:

```
[root@freeradius]# cd /usr/local/var/log/radius/radacct/122.69.1.250
[root@freeradius 122.69.1.250]# ls
```

```
detail-20071214
```

```
[root@freeradius 122.69.1.250]# vi details-20071214
```

```
User-Name = "000347e1bab9"
NAS-Port = 67
Acct-Status-Type = Stop
Acct-Session-Id = "802.1x811912"
Acct-Input-Octets = 17454
Acct-Output-Octets = 4245
Acct-Session-Time = 1221041249
Acct-Input-Packets = 72
Acct-Output-Packets = 53
Acct-Terminate-Cause = Lost-Carrier
Acct-Input-Gigawords = 0
Acct-Output-Gigawords = 0
Called-Station-Id = "00-19-e2-50-52-60"
Calling-Station-Id = "00-03-47-e1-ba-b9"
Event-Timestamp = "Sep 10 2008 16:52:39 PDT"
NAS-Identifier = "esp48t-1b-01"
NAS-Port-Type = Virtual
```

```
User-Name = "000347e1bab9"
NAS-Port = 67
Acct-Status-Type = Start
Acct-Session-Id = "802.1x811219"
Called-Station-Id = "00-19-e2-50-52-60"
Calling-Station-Id = "00-03-47-e1-ba-b9"
Event-Timestamp = "Sep 10 2008 18:58:52 PDT"
NAS-Identifier = "esp48t-1b-01"
NAS-Port-Type = Virtual
```

- Related Documentation**
- [Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch on page 1843](#)
  - [Understanding 802.1X and RADIUS Accounting on EX Series Switches on page 1830](#)

## Filtering 802.1X Supplicants Using RADIUS Server Attributes

There are two ways to configure the RADIUS server with port firewall filters:

- Include a match statement and corresponding action in the **Juniper-Firewall-Filter** attribute. The **Juniper-Firewall-Filter** attribute is a vendor-specific attribute (VSA) in

the Juniper dictionary on the RADIUS server. Use this attribute to configure simple filter conditions for authenticated users. Nothing needs to be configured on the switch; all of the configuration is on the RADIUS server.

- Apply a local firewall filter to users authenticated through the RADIUS server. Use this method for more complex filters. The firewall filter must be configured on each switch.



**NOTE:** If the firewall filter configuration is modified after users are authenticated using the 802.1X authentication, then the established 802.1X authentication session must be terminated and re-established for the firewall filter configuration changes to take effect.

This example describes using FreeRADIUS software to configure VSAs. For specifics on configuring your server, consult the AAA documentation that was included with your server.

This topic includes the following tasks:

1. [Configuring Match Statements on the RADIUS Server on page 1911](#)
2. [Applying a Port Firewall Filter from the RADIUS Server on page 1913](#)

### Configuring Match Statements on the RADIUS Server

You can configure simple filter conditions using the **Juniper-Switching-Filter** attribute in the Juniper dictionary on the RADIUS server. These filters are then sent to a switch whenever a new user is authenticated successfully. The filters are created and applied on all EX Series switches that authenticate users through that RADIUS server without the need to configure anything on each individual switch.

To configure the **Juniper-Switching-Filter** attribute, enter one or more match conditions and a resulting action using the CLI for the RADIUS server. Enter the match statement plus an action statement enclosed within quotes (" ") using the following syntax:

```
match <destination-mac mac-address> <source-vlan vlan-name> <source-dot1q-tag
tag> <destination-ip ip-address> <ip-protocol protocol-id> <source-port port>
<destination-port port>
}
action [allow | deny] <forwarding-class class-of-service> <loss-priority (low | medium |
high)>
}
```

See “[VSA Match Conditions and Actions](#)” on [page 1918](#) for definitions of match statement options.

To configure match conditions on the RADIUS server:

1. Verify that the Juniper dictionary is loaded on your RADIUS server and includes the filtering attribute **Juniper-Switching-Filter**, attribute ID 48:

```
[root@freeradius]# cat /usr/local/share/freeradius/dictionary.juniper

# dictionary.juniper
#
```

```
# Version:      $Id: dictionary.juniper,v 1.2.6.1 2005/11/30 22:17:25 a1and
Exp
$
#  VENDOR      Juniper      2636
BEGIN-VENDOR   Juniper
ATTRIBUTE      Juniper-Local-User-Name      1      string
ATTRIBUTE      Juniper-Allow-Commands       2      string
ATTRIBUTE      Juniper-Deny-Commands        3      string
ATTRIBUTE      Juniper-Allow-Configuration   4      string
ATTRIBUTE      Juniper-Deny-Configuration   5      string
ATTRIBUTE      Juniper-Switching-Filter      48     string
<-
```

2. Enter the match conditions and actions. For example:

- To deny authentication based on the 802.1Q tag (here, the 802.1Q tag is 10):

```
[root@freeradius]#
cd /usr/local/etc/raddb
vi users
```

For each relevant user, add the **Juniper-Switching-Filter** attribute:

**Juniper-Switching-Filter = "Match Source-dot1q-tag 10 Action deny"**

- To deny access based on a destination IP address:

```
[root@freeradius]# cd /usr/local/etc/raddb
vi users
```

For each relevant user, add the **Juniper-Switching-Filter** attribute:

**Juniper-Switching-Filter = "Match Destination-ip 192.168.1.0/31 Action deny"**

- To set the packet loss priority (PLP) to **high** based on a destination MAC address and the IP protocol:

```
[root@freeradius]# cd /usr/local/etc/raddb
vi users
```

For each relevant user, add the **Juniper-Switching-Filter** attribute:

**Juniper-Switching-Filter = "Match Destination-mac 00:04:0f:fd:ac:fe, Ip-protocol 2, forwarding-class high, Action loss-priority high"**



**NOTE:** For the forwarding-class option to be applied, the forwarding class must be configured on the switch. If it is not configured on the switch, this option is ignored. You must specify both the forwarding class and the packet loss priority.

3. Stop and restart the RADIUS process to activate the configuration.



### Applying a Port Firewall Filter from the RADIUS Server

You can apply a firewall filter to user policies on the RADIUS server. The RADIUS server can then specify the firewall filters that are to be applied to each user that requests to authenticate. Use this method when the firewall filter has more extensive conditions or you want to use different conditions for the same filter on different switches. The firewall filters must be configured on each switch.

For more information about firewall filters, see [“Firewall Filters for EX Series Switches Overview” on page 4696](#).

To apply a port firewall filter centrally from the RADIUS server:



**NOTE:** If port firewall filters are also configured locally for the interface, then VSAs take precedence if they conflict with the filters. If the VSAs and the local port firewall filters do not conflict, they are merged.

1. Create the firewall filter on the local switch. In this example, the filter is called **filter1**.
2. Open the users file on the RADIUS server:

```
[root@freeradius]#  
cd /usr/local/pool/raddb  
vi users
```

3. For each relevant user, add the filter (here, the filter ID is **filter1**):

```
Filter-Id = "filter1"
```



**NOTE:** Multiple filters are not supported on a single interface. However, you can support multiple filters for multiple users that are connected to the switch on the same interface by configuring a single filter with policies for each of those users.

4. Stop and restart the RADIUS process to activate the configuration.

#### Related Documentation

- [Example: Applying a Firewall Filter to 802.1X-Authenticated Supplicants Using RADIUS Server Attributes on an EX Series Switch on page 1887](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)
- [Understanding 802.1X and VSAs on EX Series Switches on page 1837](#)

### Configuring LLDP (CLI Procedure)

Devices use Link Layer Discovery Protocol (LLDP) and Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) to learn and distribute device

information on network links. The information enables the device to quickly identify a variety of other devices, resulting in a LAN that interoperates smoothly and efficiently.

This topic describes:

- [Enabling LLDP on Interfaces on page 1914](#)
- [Adjusting LLDP Advertisement Settings on page 1914](#)
- [Adjusting SNMP Notification Settings of LLDP Changes on page 1915](#)
- [Specifying a Management Address for the LLDP Management TLV on page 1916](#)
- [Configuring LLDP Power Negotiation on page 1916](#)

---

### Enabling LLDP on Interfaces

LLDP is enabled on all interfaces by default. If it is disabled, you can enable LLDP by configuring it on all interfaces or on specific interfaces.

- To configure LLDP on all interfaces:

```
[edit protocols lldp]
user@switch# set interface all
```

- To configure LLDP on a specific interface:

```
[edit protocols lldp]
user@switch# set interface interface-name
```



**NOTE:** On EX4300 switches, LLDP cannot be configured on the me0 or vme interface. Issuing the command `set protocols lldp interface me0` generates the following error message:

```
error: name: 'me0': Invalid interface
error: statement creation failed: interface
```

Issuing the command `set protocols lldp interface vme` generates the following error message:

```
error: name: 'vme': Invalid interface
error: statement creation failed: interface
```

---

### Adjusting LLDP Advertisement Settings

You can adjust the following settings for LLDP advertisements for troubleshooting or verification purposes. The default values are applied when LLDP is enabled. For normal operations, we recommend that you do not change the default values.

- To specify the frequency at which LLDP advertisements are sent (in seconds):

```
[edit protocols lldp]
user@switch# set advertisement-interval seconds
```

For example, using the default value:

```
[edit protocols lldp]
user@switch# set advertisement-interval 45
```

- To specify the number of seconds that LLDP information is held before it is discarded (the multiplier value is used in combination with the **advertisement-interval** value):

```
[edit protocols lldp]
user@switch# set hold-multiplier seconds
```

For example, using the default value:

```
[edit protocols lldp]
user@switch# set hold-multiplier 5
```

- To specify the number of seconds the device delays before sending advertisements to neighbors after a change is made in a TLV (type, length, or value) element in LLDP or in the state of the local system, such as a change in hostname or management address, set the transmit delay. The transmit delay is enabled by default on switches to reduce the delay in notifying neighbors of a change in the local system. The default value is 2 seconds (if the **advertisement-interval** value is set to 8 seconds or more) or 1 second (if the **advertisement-interval** value is set to less than 8 seconds).

```
[edit protocols lldp]
user@switch# set transmit-delay seconds
```

For example:

```
[edit protocols lldp]
user@switch# set transmit-delay 2
```



**NOTE:** The **advertisement-interval** value must be greater than or equal to four times the **transmit-delay** value; otherwise, an error is returned when you attempt to commit the configuration.

### Adjusting SNMP Notification Settings of LLDP Changes

You can adjust the following settings for SNMP notifications of LLDP changes. If the values are not specified or if the interval values are set to 0, the notifications are disabled.

- To specify the frequency at which LLDP database changes are sent (in seconds):

```
[edit protocols lldp]
user@switch# set lldp-configuration-notification-interval seconds
```

For example:

```
[edit protocols lldp]
user@switch# set lldp-configuration-notification-interval 600
```

- To configure the time interval for SNMP trap notifications to wait for topology changes (in seconds):

```
[edit protocols lldp]
user@switch# set ptopo-configuration-trap-interval seconds
```

For example:

```
[edit protocols lldp]
user@switch# set ptopo-configuration-trap-interval 600
```

- To specify the holding time (used in combination with the **ptopo-configuration-trap-interval** value) to maintain dynamic topology entries (in seconds):

```
[edit protocols lldp]
user@switch# set ptopo-configuration-maximum-hold-time seconds
```

For example:

```
[edit protocols lldp]
user@switch# set ptopo-configuration-maximum-hold-time 2147483647
```

---

### Specifying a Management Address for the LLDP Management TLV

You can configure an IPv4 or IPv6 management address to be used in the LLDP Management Address type, length, and value (TLV) messages. Only out-of-band management addresses must be used as the value for the **management-address** statement.

To configure the management address:

```
[edit protocols lldp]
user@switch# set management-address ip-address
```



**NOTE:** Ensure that the interface with the configured management address has LLDP enabled using the `set protocols lldp interface` command. If you configure a customized management address for LLDP on an interface that has LLDP disabled, the `show lldp local-information` command output will not display the correct interface information.

---

### Configuring LLDP Power Negotiation

LLDP power negotiation enables the switch's Power over Ethernet (PoE) controller to dynamically allocate PoE power to PoE interfaces, based on the needs of the powered device, by negotiating with LLDP-enabled powered devices.



**NOTE:** LLDP power negotiation is not supported on EX3200 and EX4200 (except EX4200-24P and EX4200-48P models) switches.

LLDP power negotiation is supported on switches running PoE controller software version 4.04 or later. For information about upgrading the PoE controller software, see [“Upgrading the PoE Controller Software” on page 4467](#).

LLDP power negotiation is automatically enabled when the PoE management mode is set to **class**.

- To disable LLDP power negotiation on switch interfaces:

```
[edit protocols lldp interface all power-negotiation]
user@switch# disable
```

- To disable LLDP power negotiation on a specific switch interface:

```
[edit protocols lldp interface interface-name power-negotiation]
user@switch# disable
```

**Related  
Documentation**

- [Configuring LLDP \(J-Web Procedure\)](#)

- [Configuring LLDP-MED \(CLI Procedure\) on page 1917](#)
- [Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831](#)

## Configuring LLDP-MED (CLI Procedure)

Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) is an extension of LLDP. The EX Series switch uses LLDP-MED to support device discovery of VoIP telephones and to create location databases for these telephone locations.

LLDP-MED is turned on by default on EX Series switches.

This topic describes:

- [Enabling LLDP-MED on Interfaces on page 1917](#)
- [Configuring Location Information Advertised by the Switch on page 1917](#)
- [Configuring for Fast Start on page 1918](#)

### Enabling LLDP-MED on Interfaces

LLDP-MED is enabled on all interfaces by default. If it is disabled, you can enable LLDP-MED by configuring it on all interfaces or on specific interfaces.



**NOTE:** On switches running Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, configure LLDP-MED on the physical interface—for example, on ge-0/0/2. For more about ELS, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

To configure LLDP-MED on all interfaces or on a specific interface:

```
[edit protocols lldp-med]
user@switch# set interface (LLDP-MED) ge-0/0/2.0
```

### Configuring Location Information Advertised by the Switch

You can configure the location information that is advertised from the switch to the LLDP-MED device. You can specify a civic-based location (geographic location) or a location based on an ELIN (Emergency Location Identification Number):

- To specify a location by geography:

```
[edit protocols lldp-med]
user@switch# set interface ge-0/0/2.0 location civic-based country-code US
user@switch# set interface ge-0/0/2.0 location civic-based ca-type 1 ca-value "El Dorado County"
user@switch# set interface ge-0/0/2.0 location civic-based ca-type 2 ca-value CA
user@switch# set interface ge-0/0/2.0 location civic-based ca-type 3 ca-value Somerset
user@switch# set interface ge-0/0/2.0 location civic-based ca-type 6 ca-value "Mount Aukum Road"
user@switch# set interface ge-0/0/2.0 location civic-based ca-type 19 ca-value 6450
user@switch# set interface ge-0/0/2.0 location civic-based ca-type 21 ca-value "Holiday Market"
```

- To specify a location using an elin string:

```
[edit protocols lldp-med]
user@switch# set interface ge-0/0/2.0 location elin 4085551212
```

### Configuring for Fast Start

---

You can specify the number of LLDP-MED advertisements sent from the switch in the first second after it has detected an LLDP-MED device. The default is 3; to set it to another value:

```
[edit protocols lldp-med]
user@switch# set fast-start 6
```



**NOTE:** If an interface is configured as a VoIP interface, then the switch does not wait for an attached phone to identify itself as an LLDP-MED device before it performs an LLDP-MED fast start after a graceful Routing Engine switchover (GRES) or a reboot. Instead, it immediately performs an LLDP-MED fast start after a GRES or reboot. This behavior prevents certain models of IP phones from resetting after a GRES.

---

#### Related Documentation

- [Configuring LLDP \(J-Web Procedure\)](#)
- [Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch](#)
- [Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch on page 1858](#)
- [Configuring LLDP \(CLI Procedure\) on page 1913](#)
- [Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831](#)

## VSA Match Conditions and Actions

Devices support the configuration of RADIUS server attributes specific to Juniper Networks. These attributes are known as vendor-specific attributes (VSAs). They are configured on RADIUS servers and work in combination with 802.1X authentication. Using VSAs, you can apply port firewall filter attributes as a subset of match conditions and actions sent from the RADIUS server to the switch as a result of successful 802.1X authentication.

Each term in a VSA configured through the RADIUS server consists of *match conditions* and an *action*. Match conditions are the values or fields that the packet must contain. You can define single, multiple, or no match conditions. If no match conditions are specified for the term, the packet is accepted by default. The action is the action that the switch takes if a packet matches the match conditions for the specific term. Allowed actions are to accept a packet or to discard a packet.

The following guidelines apply when you specify match conditions and actions for VSAs:

- Both **match** and **action** statements are mandatory.
- Any or all options (separated by commas) may be included in each **match** and **action** statement.

- Fields separated by commas will be ANDed if they are of a different type. The same types cannot be repeated.
- For OR cases (for example, match **10.1.1.0/24 OR 11.1.1.0/24**), apply multiple VSAs to the 802.1X supplicant.
- In order for the **forwarding-class** option to be applied, the forwarding class must be configured on the switch. If it is not configured on the switch, this option is ignored.

Table 175 on page 1919 describes the match conditions you can specify when configuring a VSA using the **match** command on the RADIUS server. The string that defines a match condition is called a *match statement*.

Table 175: Match Conditions

Option	Description
<b>destination-mac</b> <i>mac-address</i>	Destination media access control (MAC) address of the packet.
<b>source-vlan</b> <i>source-vlan</i>	Name of the source VLAN.
<b>source-dot1q-tag</b> <i>tag</i>	Tag value in the 802.1Q header, in the range 0 through 4095.
<b>destination-ip</b> <i>ip-address</i>	Address of the final destination node.
<b>ip-protocol</b> <i>protocol-id</i>	IPv4 protocol value. In place of the numeric value, you can specify one of the following text synonyms:  <b>ah</b> , <b>egp</b> (8), <b>esp</b> (50), <b>gre</b> (47), <b>icmp</b> (1), <b>igmp</b> (2), <b>ipip</b> (4), <b>ipv6</b> (41), <b>ospf</b> (89), <b>pim</b> (103), <b>rsvp</b> (46), <b>tcp</b> (6), or <b>udp</b> (17)
<b>source-port</b> <i>port</i>	TCP or User Datagram Protocol (UDP) source port field. Normally, you specify this match statement in conjunction with the <b>ip-protocol</b> match statement to determine which protocol is being used on the port. In place of the numeric field, you can specify one of the text options listed under <b>destination-port</b> .
<b>destination-port</b> <i>port</i>	TCP or UDP destination port field. Normally, you specify this match in conjunction with the <b>ip-protocol</b> match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed):  <b>afs</b> (1483), <b>bgp</b> (179), <b>biff</b> (512), <b>bootpc</b> (68), <b>bootps</b> (67), <b>cvspserver</b> (2401), <b>cmd</b> (514), <b>dhcp</b> (67), <b>domain</b> (53), <b>eklogin</b> (2105), <b>ekshell</b> (2106), <b>exec</b> (512), <b>finger</b> (79), <b>ftp</b> (21), <b>ftp-data</b> (20), <b>http</b> (80), <b>https</b> (443), <b>ident</b> (113), <b>imap</b> (143), <b>kerberos-sec</b> (88), <b>klogin</b> (543), <b>kpasswd</b> (761), <b>krb-prop</b> (754), <b>krbupdate</b> (760), <b>kshell</b> (544), <b>ldap</b> (389), <b>login</b> (513), <b>mobileip-agent</b> (434), <b>mobileip-mn</b> (435), <b>msdp</b> (639), <b>netbios-dgm</b> (138), <b>netbios-ns</b> (137), <b>netbios-ssn</b> (139), <b>nfsd</b> (2049), <b>nntp</b> (119), <b>ntalk</b> (518), <b>ntp</b> (123), <b>pop3</b> (110), <b>pptp</b> (1723), <b>printer</b> (515), <b>radacct</b> (1813), <b>radius</b> (1812), <b>rip</b> (520), <b>rkinit</b> (2108), <b>smtp</b> (25), <b>snmp</b> (161), <b>snmptrap</b> (162), <b>snpp</b> (444), <b>socks</b> (1080), <b>ssh</b> (22), <b>sunrpc</b> (111), <b>syslog</b> (514), <b>telnet</b> (23), <b>tacacs-ds</b> (65), <b>talk</b> (517), <b>tftp</b> (69), <b>timed</b> (525), <b>who</b> (513), <b>xdmcp</b> (177), <b>zephyr-clt</b> (2103), <b>zephyr-hm</b> (2104)

When you define one or more terms that specify the filtering criteria, you also define the action to take if the packet matches all criteria. [Table 176 on page 1920](#) shows the actions that you can specify in a term.

**Table 176: Actions for VSAs**

Option	Description
(allow   deny)	Accept a packet or discard a packet silently without sending an Internet Control Message Protocol (ICMP) message.
forwarding-class <i>class-of-service</i>	(Optional) Classify the packet in one of the following forwarding classes: <ul style="list-style-type: none"> <li>assured-forwarding</li> <li>best-effort</li> <li>expedited-forwarding</li> <li>network-control</li> </ul>
loss-priority (low   medium   high)	(Optional) Set the packet loss priority (PLP) to <b>low</b> , <b>medium</b> , or <b>high</b> . Specify both the forwarding class and loss priority.

- Related Documentation**
- [Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 1910](#)
  - [Understanding 802.1X and VSAs on EX Series Switches on page 1837](#)
  - [Understanding VSAs](#)



## Configuring Server Fail Fallback (CLI Procedure)

Server fail fallback allows you to specify how end devices connected to the switch are supported if the RADIUS authentication server becomes unavailable or sends a RADIUS access-reject message.

802.1X and MAC RADIUS authentication work by using an *authenticator port access entity* (the EX Series switch) to block all traffic to and from an end device at the interface until the end device's credentials are presented and matched on the *authentication server* (a RADIUS server). When the end device has been authenticated, the switch stops blocking and opens the interface to the end device.

When you set up 802.1X or MAC RADIUS authentication on the switch, you specify a primary authentication server and one or more backup authentication servers. If the primary authentication server cannot be reached by the switch and the secondary authentication servers are also unreachable, a RADIUS server timeout occurs. Because the authentication server grants or denies access to the end devices awaiting authentication, the switch does not receive access instructions for end devices attempting access to the LAN and normal authentication cannot be completed. Server fail fallback allows you to configure authentication alternatives that permit the switch to take appropriate actions toward end devices awaiting authentication or reauthentication.



**NOTE:** The authentication fallback method called *server-reject VLAN* provides limited access to a LAN, typically just to the Internet, for responsive end devices that are 802.1X-enabled but that have sent the wrong credentials. If the end device that is authenticated using the server-reject VLAN is an IP phone, voice traffic is not allowed.

To configure basic server fail fallback options using the CLI:

- Configure an interface to allow traffic to flow from a supplicant to the LAN if a RADIUS server timeout occurs (as if the end device had been successfully authenticated by a RADIUS server):

```
[edit protocols dot1x authenticator]
user@switch# set interface ge-0/0/1 server-fail permit
```

- Configure an interface to prevent traffic flow from an end device to the LAN (as if the end device had failed authentication and had been rejected by the RADIUS server):

```
[edit protocols dot1x authenticator]
user@switch# set interface ge-0/0/1 server-fail deny
```

- Configure an interface to move an end device to a specified VLAN if a RADIUS server timeout occurs (in this case, the VLAN name is `vlan1`):

```
[edit protocols dot1x authenticator]
user@switch# set interface ge-0/0/1 server-fail vlan-name vlan1
```

- Configure an interface to recognize already connected end devices as reauthenticated if there is a RADIUS timeout during reauthentication (new users will be denied access):

```
[edit protocols dot1x authenticator]
user@switch# set interface ge-0/0/1 server-fail use-cache
```

- Configure an interface that receives a RADIUS access-reject message from the authentication server to move end devices attempting LAN access on the interface to a specified VLAN already configured on the switch (in this case, the VLAN name is `vlan-sf`):

```
[edit protocols dot1x authenticator]
user@switch# set interface ge-0/0/1 server-reject-vlan vlan-sf
```



**NOTE:** If an IP phone is authenticated in the server-reject VLAN, voice traffic is not allowed.

---

#### Related Documentation

- [Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to an EX Series Switch on page 1875](#)
- [Configuring 802.1X Authentication \(J-Web Procedure\)](#)
- [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)
- [Monitoring 802.1X Authentication on page 1993](#)
- [Understanding Server Fail Fallback and Authentication on EX Series Switches on page 1838](#)

## Configuring MAC RADIUS Authentication (CLI Procedure)

You can permit devices that are not 802.1X-enabled LAN access by configuring MAC RADIUS authentication on the EX Series switch interfaces to which the hosts are connected.



**NOTE:** You can also allow non-802.1X-enabled devices to access the LAN by configuring their MAC address for static MAC bypass of authentication.

You can configure MAC RADIUS authentication on an interface that also allows 802.1X authentication, or you can configure either authentication method alone.

If both MAC RADIUS and 802.1X authentication are enabled on the interface, the switch first sends the host three EAPOL requests to the host. If there is no response from the host, the switch sends the host's MAC address to the RADIUS server to check whether it is a permitted MAC address. If the MAC address is configured as permitted on the RADIUS server, the RADIUS server sends a message to the switch that the MAC address is a permitted address, and the switch opens LAN access to the nonresponsive host on the interface to which it is connected.

If MAC RADIUS authentication is configured on the interface but 802.1X authentication is not (by using the **mac-radius restrict** option), the switch attempts to authenticate the MAC address with the RADIUS server without delaying by attempting 802.1X authentication first.

Before you configure MAC RADIUS authentication, be sure you have:

- Configured basic access between the EX Series switch and the RADIUS server. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).

To configure MAC RADIUS authentication using the CLI:

- On the switch, configure the interfaces to which the nonresponsive hosts are attached for MAC RADIUS authentication, and add the **restrict** qualifier for interface **ge-0/0/20** to have it use only MAC RADIUS authentication:

```
[edit]
user@switch# set protocols dot1x authenticator interface ge-0/0/19 mac-radius
user@switch# set protocols dot1x authenticator interface ge-0/0/20 mac-radius restrict
```

- On a RADIUS authentication server, create user profiles for each nonresponsive host using the MAC address (without colons) of the nonresponsive host as the username and password (here, the MAC addresses are **00:04:0f:fd:ac:fe** and **00:04:ae:cd:23:5f**):

```
[root@freeradius]#
edit /etc/raddb
vi users
00040ffdacfe Auth-type:=Local, User-Password = "00040ffdacfe"
0004aecdc235f Auth-type:=Local, User-Password = "0004aecdc235f"
```

### Related Documentation

- [Example: Configuring MAC RADIUS Authentication on an EX Series Switch on page 1881](#)

- [Verifying 802.1X Authentication on page 1994](#)
- [Understanding Authentication on EX Series Switches on page 1824](#)

## Configuring Static MAC Bypass of Authentication (CLI Procedure)

You can configure a static MAC bypass list (sometimes called the exclusion list) on the switch to specify MAC addresses of devices allowed access to the LAN without 802.1X or MAC RADIUS authentication requests to the RADIUS server.

To configure the static MAC bypass list:

- Specify a MAC address to bypass authentication:

```
[edit protocols dot1x]  
user@switch# set authenticator static 00:04:0f:fd:ac:fe
```

- Configure a supplicant to bypass authentication if connected through a particular interface:

```
[edit protocols dot1x]  
user@switch# set authenticator static 00:04:0f:fd:ac:fe interface ge-0/0/5
```

- You can configure a supplicant to be moved to a specific VLAN after it is authenticated:

```
[edit protocols dot1x]  
user@switch# set authenticator static 00:04:0f:fd:ac:fe interface ge-0/0/5 vlan-assignment  
default-vlan
```

### Related Documentation

- [Example: Configuring Static MAC Bypass of Authentication on an EX Series Switch on page 1872](#)
- [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)
- [Configuring 802.1X Authentication \(J-Web Procedure\)](#)

## Specifying RADIUS Server Connections on an EX Series Switch (CLI Procedure)

IEEE 802.1X and MAC RADIUS authentication both provide network edge security, protecting Ethernet LANs from unauthorized user access by blocking all traffic to and from devices at the interface until the supplicant's credentials or MAC address are presented and matched on the *authentication server* (a RADIUS server). When the supplicant is authenticated, the switch stops blocking access and opens the interface to the supplicant.

To use 802.1X or MAC RADIUS authentication, you must specify the connections on the switch for each RADIUS server to which you will connect.

To configure a RADIUS server on the switch:

1. Define the IP address of the RADIUS server, the RADIUS server authentication port number, and the secret password. You can define more than one RADIUS server. The secret password on the switch must match the secret password on the server:

```
[edit access]
user@switch# set radius-server 10.0.0.100 port 1812 secret abc
```



**NOTE:** Specifying the authentication port is optional, and port 1812 is the default. However, we recommend that you configure it in order to avoid confusion as some RADIUS servers might refer to an older default.

2. (Optional) Specify the IP address by which the switch is identified by the RADIUS server. If you do not specify this, the RADIUS server uses the address of the interface sending the RADIUS request. We recommend that you specify this IP address because if the request gets diverted on an alternate route to the RADIUS server, the interface relaying the request might not be an interface on the switch.

```
[edit access]
user@switch# set access radius-server source-address 10.93.14.100
```

3. Configure the authentication order, making **radius** the first method of authentication:

```
[edit access]
user@switch# set profile profile1 authentication-order radius
```

4. Create a profile and specify the list of RADIUS servers to be associated with the profile. For example, you might choose to group your RADIUS servers geographically by city. This feature enables easy modification whenever you want to change to a different set of authentication servers.

```
[edit access profile]
user@switch# set atlanta radius authentication-server 10.0.0.100 10.2.14.200
```

5. Specify the group of servers to be used for 802.1X or MAC RADIUS authentication by identifying the profile name:

```
[edit access profile]
user@switch# set protocols dot1x authenticator authentication-profile-name denver
```

6. Configure the IP address of the EX Series switch in the list of clients on the RADIUS server. For specifics on configuring the RADIUS server, consult the documentation for your server.

**Related Documentation**

- [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)
- [Configuring 802.1X Authentication \(J-Web Procedure\)](#)
- [Configuring MAC RADIUS Authentication \(CLI Procedure\) on page 1923](#)
- [Configuring 802.1X RADIUS Accounting \(CLI Procedure\) on page 1909](#)

## Configuring Captive Portal Authentication (CLI Procedure)

---



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Configuring Captive Portal Authentication \(CLI Procedure\)](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

---

Configure captive portal authentication (hereafter referred to as captive portal) on an EX Series switch so that users connected to the switch are authenticated before being allowed to access the network. When the user requests a webpage, a login page is displayed that requires the user to input a username and password. Upon successful authentication, the user is allowed to continue with the original page request and subsequent access to the network.

Before you begin, be sure you have:

- Performed basic bridging and VLAN configuration on the switch. See [“Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch” on page 2281](#).
- Generated an SSL certificate and installed it on the switch. See [“Generating SSL Certificates to Be Used for Secure Web Access” on page 532](#).
- Configured basic access between the EX Series switch and the RADIUS server. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).
- Designed your captive portal login page. See [“Designing a Captive Portal Authentication Login Page on an EX Series Switch” on page 1928](#).

This topic includes the following tasks:

- [Configuring Secure Access for Captive Portal on page 1926](#)
- [Enabling an Interface for Captive Portal on page 1927](#)
- [Configuring Bypass of Captive Portal Authentication on page 1927](#)

### Configuring Secure Access for Captive Portal

---

To configure secure access for captive portal:

1. Associate the security certificate with the Web server and enable HTTPS on the switch:

[edit]

```
user@switch# set system services web-management https local-certificate my-signed-cert
```



**NOTE:** You can enable HTTP instead of HTTPS, but we recommend HTTPS for security purposes.

2. Configure captive portal to use HTTPS:

```
[edit]
user@switch# set services captive-portal secure-authentication https
```

### Enabling an Interface for Captive Portal

To enable an interface for use with captive portal authentication:

```
[edit]
user@switch# set services captive-portal interface ge-0/0/10
```

### Configuring Bypass of Captive Portal Authentication

You can allow specific clients to bypass captive portal authentication:

```
[edit]
user@switch# set switch-options authentication-whitelist 00:10:12:e0:28:22
```



**NOTE:** Optionally, you can use `set switch-options authentication-whitelist 00:10:12:e0:28:22 interface ge-0/0/10.0` to limit the scope to the interface.



**NOTE:** If the client is already attached to the switch, you must clear its MAC address from the captive portal authentication by using the `clear captive-portal mac-address session-mac-addr` command after adding its MAC address to the whitelist. Otherwise the new entry for the MAC address will not be added to the Ethernet switching table and the authentication bypass will not be allowed.

#### Related Documentation

- [Example: Setting Up Captive Portal Authentication on an EX Series Switch on page 1898](#)
- [Understanding Authentication on EX Series Switches on page 1824](#)

## Designing a Captive Portal Authentication Login Page on an EX Series Switch

You can set up captive portal authentication on your switch to redirect all Web browser requests to a login page that requires the user to input a username and password before they are allowed access. Upon successful authentication, the user is allowed access to the network and redirected to the original page requested.

Junos OS provides a customizable template for the captive portal window that allows you to easily design and modify the look of the captive portal login page. You can modify the design elements of the template to change the look of your captive portal login page and to add instructions or information to the page. You can also modify any of the design elements of a captive portal login page.

The first screen displayed before the captive login page requires the user to read the “Terms and Conditions of Use”. By clicking the Agree button, the user can access the captive portal login page.

Figure 21 on page 1928 shows an example of a captive portal login page:

Figure 21: Example of a Captive Portal Login Page

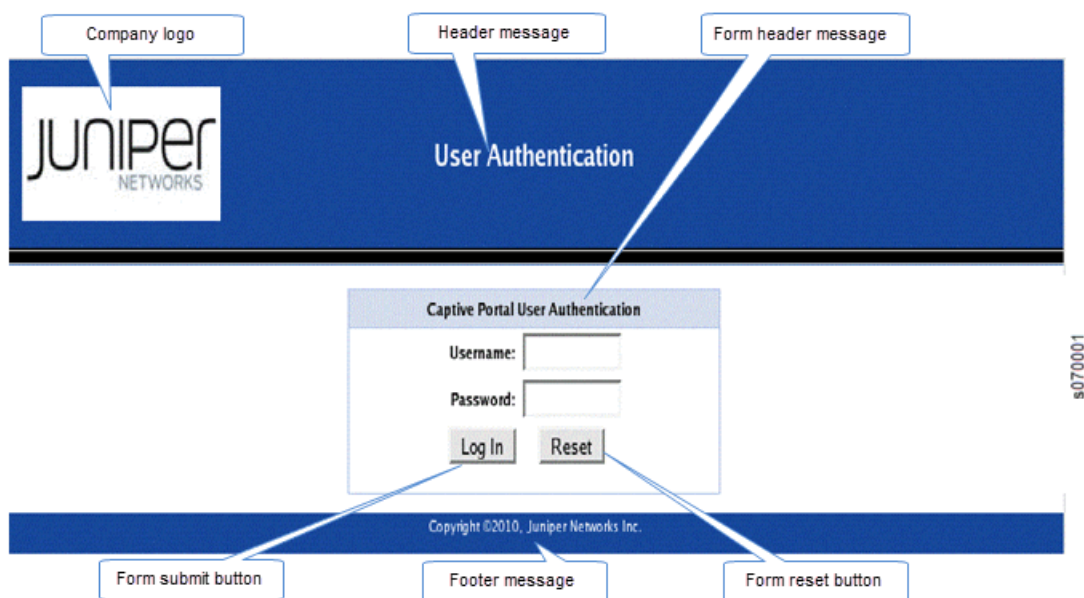


Table 177 on page 1928 summarizes the configurable elements of a captive portal login page.

Table 177: Configurable Elements of a Captive Portal Login Page

Element	CLI Statement	Description
Footer background color	<b>footer-bgcolor</b> <i>hex-color</i>	The HTML hexadecimal code for the background color of the captive portal login page footer.



Table 177: Configurable Elements of a Captive Portal Login Page (*continued*)

Element	CLI Statement	Description
Footer message	<b>footer-message</b> <i>text-string</i>	Text displayed in the footer of the captive portal login page. You can include copyright information, links, and additional information such as help instructions, legal notices, or a privacy policy  The default text shown in the footer is <b>Copyright ©2010, Juniper Networks Inc.</b>
Footer text color	<b>footer- text-color</b> <i>color</i>	Color of the text in the footer. The default color is white.
Form header background color	<b>form-header-bgcolor</b> <i>hex-color</i>	The HTML hexadecimal code for the background color of the header bar across the top of the form area of the captive portal login page.
Form header message	<b>form-header-message</b> <i>text-string</i>	Text displayed in the header of the captive portal login page. The default text is <b>Captive Portal User Authentication</b>
Form header text color	<b>form-header- text- color</b> <i>color</i>	Color of the text in the form header. The default color is black.
Form reset button label	<b>form-reset-label</b> <i>label-name</i>	Using the <b>Reset</b> button, the user can clear the username and password fields on the form.
Form submit button label	<b>form-submit-label</b> <i>label-name</i>	Using the <b>Login</b> button, the user can submit the login information.
Header background color	<b>header-bgcolor</b> <i>hex-color</i>	The HTML hexadecimal code for the background color of the captive portal login page header.
Header logo	<b>header-logo</b> <i>filename</i>	Filename of the file containing the image of the logo that you want to appear in the header of the captive portal login page. The image file can be in GIF, JPEG, or PNG format  You can upload a logo image file to the switch. Copy the logo to the /var/tmp directory on the switch (during commit, the files are saved to persistent locations).  If you do not specify a logo image, the Juniper Networks logo is displayed.
Header message	<b>header-message</b> <i>text-string</i>	Text displayed in the page header. The default text is <b>User Authentication</b> .
Header text color	<b>header-text- color</b> <i>color</i>	Color of the text in the header. The default color is white.
Post-authentication URL	<b>post-authentication-url</b> <i>url</i>	URL to which the users are directed on successful authentication. By default, users are directed to the page they had originally requested.

To design the captive portal login page:

- (Optional) Upload your logo image file to the switch:  

```
user@switch> file copy ftp://username:prompt@ftp.hostname.net/var/tmp/my-logo.jpeg
```
- Configure the custom options to specify the background colors and text displayed in the captive portal page:

```
[edit system services captive-portal]
user@switch# set custom-options header-bgcolor #006600
set custom-options header-message "Welcome to Our Network"
set custom-options banner-message "Please enter your username and password".The banner
displays the message "XXXXXXX" by default. The user can modify this message.
set custom-options footer-message "Copyright ©2010, Our Network"
```

Now you can commit the configuration.



**NOTE:** For the custom options that you do not specify, the default value is used.

#### Related Documentation

- *Example: Setting Up Captive Portal Authentication on an EX Series Switch*
- [Understanding Authentication on EX Series Switches on page 1824](#)
- *captive-portal*

## Controlling Authentication Session Timeouts (CLI Procedure)

For 802.1X and MAC RADIUS authentication sessions, you can specify authentication session timeout values using the **reauthentication** statement.

The session might also end when the MAC table aging time expires, because the session is removed from the authentication session table when the MAC address is removed from the Ethernet switching table. In order to prevent the session from being removed from the authentication session table, you must disassociate the authentication table from the Ethernet switching table using the **no-mac-table-binding** statement.

Before you begin:

- Specify the RADIUS server or servers to be used as the authentication server. See [“Specifying RADIUS Server Connections on an EX Series Switch \(CLI Procedure\)” on page 1925](#).
- Configure 802.1X authentication on the switch. See [“Configuring 802.1X Interface Settings \(CLI Procedure\)” on page 1908](#).

To configure the authentication session time on all interfaces:

```
[edit]
user@switch# set protocols dot1x authenticator interface all reauthentication seconds;
```

To configure the authentication session time on a single interface:

```
[edit]
user@switch# set protocols dot1x authenticator interface interface-name reauthentication
seconds;
```

To disable removal of authentication sessions from the authentication session table when a MAC address ages out of the Ethernet switching table, remove the binding of the authentication table to the Ethernet switching table.

To remove the binding on all interfaces:

```
[edit]
user@switch# set protocols dot1x authenticator no-mac-table-binding interface all;
```

To remove the binding on a single interface:

```
[edit]
user@switch# set protocols dot1x authenticator no-mac-table-binding interface interface-name;
```

#### Related Documentation

- [Configuring MAC Table Aging \(CLI Procedure\)](#)
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Understanding Authentication on EX Series Switches on page 1824](#)
- [Understanding Authentication Session Timeout on page 1841](#)

## Configuration Statements

- [\[edit access\] Configuration Statement Hierarchy on EX Series Switches on page 1933](#)
- [\[edit protocols dot1x\] Configuration Statement Hierarchy on EX Series Switches on page 1939](#)
- [accounting on page 1942](#)
- [accounting \(Access Profile\) on page 1943](#)
- [accounting-order on page 1944](#)
- [accounting-port on page 1944](#)
- [address-assignment \(Address-Assignment Pools\) on page 1945](#)
- [address-protection on page 1946](#)
- [authorization-order on page 1947](#)
- [authentication-order on page 1948](#)
- [authentication-whitelist on page 1949](#)
- [authenticator on page 1950](#)
- [client-accounting-algorithm on page 1951](#)
- [client-authentication-algorithm on page 1952](#)
- [coa-dynamic-variable-validation on page 1952](#)

- [destination \(Accounting\)](#) on page 1953
- [destination-host \(Gx-Plus\)](#) on page 1954
- [destination-realm \(Gx-Plus\)](#) on page 1954
- [diameter-instance \(Gx-Plus\)](#) on page 1955
- [domain \(Domain Map\)](#) on page 1956
- [domain-name-server \(Routing Instances and Access Profiles\)](#) on page 1957
- [domain-name-server-inet \(Routing Instances and Access Profiles\)](#) on page 1958
- [domain-name-server-inet6 \(Routing Instances and Access Profiles\)](#) on page 1959
- [ethernet-port-type-virtual](#) on page 1959
- [global \(Gx-Plus\)](#) on page 1960
- [gx-plus \(Gx-Plus\)](#) on page 1960
- [ignore](#) on page 1961
- [include-ipv6 \(Gx-Plus\)](#) on page 1962
- [interface \(Static MAC Bypass\)](#) on page 1963
- [interface \(VoIP\)](#) on page 1964
- [interface-description-format](#) on page 1965
- [juniper-dsl-attributes](#) on page 1966
- [max-outstanding-requests \(Gx-Plus\)](#) on page 1967
- [nas-identifier](#) on page 1967
- [nas-port-extended-format \(Access Profile\)](#) on page 1968
- [nas-port-id-delimiter \(Subscriber Management\)](#) on page 1969
- [nas-port-id-format \(Subscriber Management\)](#) on page 1970
- [nas-port-type \(Subscriber Management\)](#) on page 1971
- [options \(Access Profile\)](#) on page 1973
- [partition \(Gx-Plus\)](#) on page 1974
- [port](#) on page 1975
- [provisioning-order](#) on page 1976
- [radius \(Access Profile\)](#) on page 1977
- [radius \(System\)](#) on page 1979
- [radius-options \(Protocols 802.1X\)](#) on page 1980
- [radius-options \(Access\)](#) on page 1981
- [radius-server \(System\)](#) on page 1981
- [retry](#) on page 1982
- [revert-interval](#) on page 1983
- [routing-instance](#) on page 1983
- [secret](#) on page 1984
- [send-acct-status-on-config-change \(Access Profile\)](#) on page 1984

- [server \(RADIUS Accounting\) on page 1985](#)
- [service \(Service Accounting\) on page 1985](#)
- [source-address on page 1986](#)
- [timeout \(RADIUS\) on page 1987](#)
- [vlan \(VoIP\) on page 1988](#)
- [vlan-assignment on page 1989](#)
- [vlan-nas-port-stacked-format on page 1990](#)
- [voip on page 1990](#)
- [wait-for-acct-on-ack \(Access Profile\) on page 1991](#)

## [edit access] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit access]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit access\] Hierarchy Level on page 1933](#)
- [Unsupported Statements in the \[edit access\] Hierarchy Level on page 1938](#)

### Supported Statements in the [edit access] Hierarchy Level

The following hierarchy shows the **[edit access]** configuration statements supported on EX Series switches:

```
access {
  address-assignment {
    abated-utilization;
    abated-utilization-v6;
    high-utilization;
    high-utilization-v6;
    neighbor-discovery-router-advertisement;
    pool name {
      family {
        inet {
          dhcp-attributes {
            boot-file filename;
            boot-server server-address;
            domain-name domain-name;
            grace-period seconds;
            maximum-lease-time (length | infinite);
```

```
name-server ip-address;
netbios-node-type (b-node | h-node | m-node | p-node);
option option-identifier-code;
option-match {
  option-82 {
    circuit-id match-value;
    remote-id match-value;
  }
}
router ip-address;
server-identifier;
tftp-server;
wins-server ip-address;
}
host;
network;
range;
xauth-attributes;
}
inet6 {
  dhcp-attributes;
  prefix;
  range;
}
}
link name {
  family {
    inet;
    inet6;
  }
}
}
}
address-pool pool-name {
  address address-or-prefix;
  address-range <low lower-limit> <high upper-limit>;
}
address-protection;
domain {
  delimiter characters;
  map name {
    aaa-logical-system name {
      aaa-routing-instance;
    }
    aaa-routing-instance aaa-routing-instance;
    access-profile;
    address-pool;
    dynamic-profile;
    padn destination-ip-address;
    strip-domain;
    target-logical-system;
    target-routing-instance;
  }
  parse-direction (left-to-right | right-to-left);
}
domain-name-server address;
```

```

domain-name-server-inet address;
domain-name-server-inet6 address;
group-profile;
gx-plus {
  global {
    include-ipv6;
    max-outstanding-requests;
  }
  partition {
    destination-host;
    destination-realm;
    diameter-instance;
  }
}
ldap-options {
  assemble {
    common-name name;
  }
  base-distinguished-name name;
  revert-interval seconds;
  search {
    admin-search {
      distinguished-name name;
      password password;
    }
    search-filter filter;
  }
}
ldap-server address {
  port number;
  retry number;
  routing-instance routing-instance;
  source-address address;
  timeout seconds;
}
ppp-options;
profile profile-name {
  accounting (Access Profile) {
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    coa-immediate-update;
    duplication;
    immediate-update;
    order (radius | none);
    statistics (time | volume-time);
    wait-for-acct-on-ack;
  }
  accounting-order (radius | [accounting-order-data-list]);
  address-assignment {
    pool;
  }
  authentication-order [(ldap | none | password | radius | secureid)];
  authorization-order (jsrc | [authorization-order-data-list]);
  client client-name {
    chap-secret chap-secret;
    client-group;
  }
}

```

```
    firewall-user {
        password password;
    }
    ike;
    no-rfc2486;
    pap-password password;
}
client-name-filter {
    count number;
    domain-name name;
    separator character;
}
domain-name-server;
domain-name-server-inet;
domain-name-server-inet6;
ldap-options {
    assemble {
        common-name name;
    }
    base-distinguished-name name;
    revert-interval seconds;
    search {
        admin-search {
            distinguished-name name;
            password password;
        }
        search-filter filter;
    }
}
}
ldap-server address {
    port number;
    retry number;
    routing-instance routing-instance;
    source-address address;
    timeout seconds;
}
provisioning-order {
    gx-plus;
    jsr;
}
radius {
    accounting-server [server-addresses];
    attributes {
        exclude [exclude-options];
        ignore [ignore-options];
    }
    authentication-server [server-addresses];
    options {
        accounting-session-id-format (decimal | description);
        client-accounting-algorithm (direct | round-robin);
        client-authentication-algorithm (direct | round-robin);
        coa-dynamic-variable-validation;
        ethernet-port-type-virtual;
        interface-description-format {
            exclude-adapter;
            exclude-sub-interface;
```



```

    }
    juniper-dsl-attributes;
    nas-identifier nas-identifier;
    nas-port-extended-format {
        adapter-width adapter-width;
        ae-width ae-width;
        port-width port-width;
        slot-width slot-width;
        stacked-vlan-width stacked-vlan-width;
        vlan-width vlan-width;
    }
    nas-port-id-delimiter nas-port-id-delimiter;
    nas-port-id-format {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        nas-identifier;
    }
    nas-port-type {
        ethernet;
    }
    revert-interval seconds;
    vlans-nas-port-stacked-format;
}

radius-server address {
    max-outstanding-requests max-outstanding-requests;
    port port-number;
    retry retry;
    routing-instance instance-name;
    secret secret;
    source-address address;
    timeout seconds;
}

service {
    accounting-order {
        activation-protocol;
        radius;
    }
}

session-options {
    client-group [group-names];
    client-idle-timeout minutes;
    client-session-timeout minutes;
}

}

radius-options {
    interim-rate number;
    interim-update-tolerance interim-update-tolerance;
    request-rate number;
    revert-interval interval;
}

radius-server server-address {
    accounting-port port-number;
    max-outstanding-requests number;
    port port-number;

```

```

    retry attempts;
    routing-instance instance-name;
    secret password;
    source-address address;
    timeout seconds;
  }
  securid-server server-name {
    configuration-file file-path;
  }
  terminate-code {
  }
}

```

### Unsupported Statements in the [edit access] Hierarchy Level

All statements in the [edit access] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 178: Unsupported [edit access] Configuration Statements on EX Series Switches**

Statement	Hierarchy Level
<b>NOTE:</b> Variables, such as <i>filename</i> , are not shown in the statements or hierarchies listed below.	
aaa	[edit access terminate-code]
administrative-reset	[edit access terminate-code aaa shutdown]
authentication-denied	[edit access terminate-code aaa deny]
client-request	[edit access terminate-code aaa dhcp]
compliance	[edit access ppp-options]
deny	[edit access terminate-code aaa]
dhcp	[edit access terminate-code]
group-profile	[edit access]
ike	[edit access profile client]
initiate-dead-peer-detection	[edit access profile client ike]
lost-carrier	[edit access terminate-code dhcp]
nak	[edit access terminate-code dhcp]
nas-logout	[edit access terminate-code dhcp]
no-offers	[edit access terminate-code dhcp]

Table 178: Unsupported [edit access] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy Level
no-resources	[edit access terminate-code aaa deny]
ppp-options	[edit access]
preference	[edit access profile client ike reverse-route]
remote-reset	[edit access terminate-code aaa shutdown]
rfc	[edit access ppp-options compliance]
reverse-route	[edit access profile client ike]
server-request-timeout	[edit access terminate-code aaa deny]
shutdown	[edit access terminate-code aaa]
terminate-code	[edit access]

- Related Documentation**
- [Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch on page 1843](#)
  - [Configuring 802.1X RADIUS Accounting \(CLI Procedure\) on page 1909](#)
  - [Security Features for EX Series Switches Overview on page 4693](#)

## [edit protocols dot1x] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols dot1x] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols dot1x\] Hierarchy Level on page 1940](#)
- [Unsupported Statements in the \[edit protocols dot1x\] Hierarchy Level on page 1940](#)

### Supported Statements in the [edit protocols dot1x] Hierarchy Level

---

The following hierarchy shows the [edit protocols dot1x] configuration statements supported on EX Series switches:

```
protocols {
  dot1x {
    authenticator {
      authentication-profile-name access-profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan (vlan-id | vlan-name);
        mac-radius {
          flap-on-disconnect;
          restrict;
        }
        maximum-requests number;
        no-reauthentication;
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
        retries number;
        server-fail (deny | permit | use-cache | vlan-id | vlan-name);
        server-reject-vlan (vlan-id | vlan-name) {
          eapol-block;
          block-interval block-interval;
        }
        server-timeout seconds;
        supplicant (single | single-secure | multiple);
        supplicant-timeout seconds;
        transmit-period seconds;
      }
      no-mac-table-binding {
        interface interface-names
        static mac-address
      }
      static mac-address {
        interface interface-names;
        vlan-assignment (vlan-id | vlan-name);
      }
    }
  }
}
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regex>;
  flag flag ;
}
```

### Unsupported Statements in the [edit protocols dot1x] Hierarchy Level

---

All statements in the [edit protocols dot1x] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

**Related  
Documentation**

- [Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an EX Series Switch on page 1847](#)
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to an EX Series Switch on page 1875](#)
- [Example: Configuring Fallback Options on EX Series Switches for EAP-TTLS Authentication and Odyssey Access Clients on page 1903](#)
- [Example: Configuring Static MAC Bypass of Authentication on an EX Series Switch on page 1872](#)
- [802.1X for EX Series Switches Overview on page 1821](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## accounting

---

**Syntax**

```
accounting {
  events [ login change-log interactive-commands ];
  destination {
    radius {
      server {
        server-address {
          accounting-port port-number;
          secret password;
          source-address address;
          retry number;
          timeout seconds;
        }
      }
    }
    tacplus {
      server {
        server-address {
          port port-number;
          secret password;
          single-connection;
          timeout seconds;
        }
      }
    }
  }
  enhanced-avs-max <number>;
}
```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
**enhanced-avs-max** statement introduced in Junos OS Release 14.1.

**Description** Configure audit of TACACS+ or RADIUS authentication events, configuration changes, and interactive commands.

The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring RADIUS System Accounting*
- *Configuring TACACS+ System Accounting*
- *enhanced-avs-max*

## accounting (Access Profile)

**Syntax**

```

accounting {
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    address-change-immediate-update;
    coa-immediate-update;
    coa-no-override service-class-attribute;
    duplication;
    duplication-vrf {
        access-profile-name profile-name;
        vrf-name vrf-name;
    }
    immediate-update;
    order [ accounting-method ];
    send-acct-status-on-config-change
    statistics (time | volume-time);
    update-interval minutes;
    wait-for-acct-on-ack;
}

```

**Hierarchy Level** [edit access profile *profile-name*]

**Release Information** Statement introduced in Junos OS Release 9.1.  
Statement introduced in Junos OS Release 9.1 for EX Series switches.

**Description** Configure RADIUS accounting parameters and enable RADIUS accounting for an access profile.

The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring Authentication and Accounting Parameters for Subscriber Access*
- *Configuring Per-Subscriber Session Accounting*
- *Understanding RADIUS Accounting Duplicate Reporting*

## accounting-order

---

<b>Syntax</b>	accounting-order (radius   [ <i>accounting-order-data-list</i> ]);
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Enable RADIUS accounting for an L2TP profile.
<b>Options</b>	<b>radius</b> —Use the RADIUS accounting method.  <b>[<i>accounting-order-data-list</i>]</b> —Set of data listing the accounting order to be used, enclosed by brackets. This can be any combination of accounting methods, up to and including and entire list of the entire accounting order.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Access Profiles for L2TP or PPP Parameters</i></li></ul>

## accounting-port

---

<b>Syntax</b>	accounting-port <i>port-number</i> ;
<b>Hierarchy Level</b>	[edit access radius-server <i>server-address</i> ], [edit access profile <i>profile-name</i> radius-server <i>server-address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure the port number on which to contact the accounting server.
<b>Options</b>	<b><i>port-number</i></b> —Port number on which to contact the accounting server. Most RADIUS servers use port number 1813 (as specified in RFC 2866).
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Router or Switch Interaction with RADIUS Servers</i></li><li>• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i></li><li>• <i>Configuring RADIUS Authentication for L2TP</i></li></ul>



## address-assignment (Address-Assignment Pools)

```
Syntax  address-assignment {
        abated-utilization percentage;
        abated-utilization-v6 percentage;
        high-utilization percentage;
        high-utilization-v6 percentage;
        neighbor-discovery-router-advertisement ndra-pool-name;
        pool pool-name {
            family family {
                dhcp-attributes {
                    protocol-specific attributes;
                }
                host hostname {
                    hardware-address mac-address;
                    ip-address ip-address;
                }
                network ip-prefix / <prefix-length>;
                prefix ipv6-prefix;
                range range-name {
                    high upper-limit;
                    low lower-limit;
                    prefix-length prefix-length;
                }
            }
            link pool-name;
        }
    }
```

**Hierarchy Level** [edit access]

**Release Information** Statement introduced in Junos OS Release 9.0.  
Statement introduced in Junos OS Release 12.1 for EX Series switches.

**Description** Configure address-assignment pools that can be used by different client applications.



**NOTE:** Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

**Options** *pool-name*—Name assigned to an address-assignment pool.

The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- [Address-Assignment Pools Overview on page 1370](#)
- [Configuring Address-Assignment Pools](#)

- *Configuring an Address-Assignment Pool for L2TP LNS with Inline Services*
- *Configuring a DHCP Server on Switches (CLI Procedure)*

## address-protection

---

<b>Syntax</b>	address-protection;
<b>Hierarchy Level</b>	[edit access], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> access]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	<p>Prevents IPv6 prefixes from being used more than once when AAA is used to supply IPv6 prefixes for router advertisement.</p> <p>If enabled, the router checks the following attributes received from external servers:</p> <ul style="list-style-type: none"><li>• <i>Framed-IPv6-Prefix</i></li><li>• <i>Framed-IPv6-Pool</i></li></ul> <p>The router then takes one of the following actions:</p> <ul style="list-style-type: none"><li>• If a prefix overlaps with a prefix in an address pool, the prefix is taken from the pool if it is available.</li><li>• If the prefix is already in use, it is rejected as unavailable.</li><li>• If the prefix length requested from the external server does not match the pool's prefix length exactly, the authentication request is denied. If configured, the Acct-Stop message includes a termination cause.</li></ul>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Duplicate Prefix Protection for Router Advertisement</i></li><li>• <i>Duplicate Prefix Protection for NDRA</i></li></ul>

## authorization-order

---

<b>Syntax</b>	authorization-order (jsrc   [ <i>authorization-order-data-list</i> ]);
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure AAA to use JSRC in an SRC environment to request authorization from the SAE when verifying that a DHCP subscriber can access the router. When you include this statement, AAA ignores any configured authentication order settings. This statement is ignored for non-DHCP subscribers.
<b>Options</b>	<p>jsrc—Use JSRC application to communicate with the SAE for subscriber authorization. JSRC is the only application that is currently available.</p> <p>[<i>authorization-order-data-list</i>]<i>—Set of data listing the authorization order to be used, enclosed by brackets. This can be any combination of authorization methods, up to and including and entire list of the entire authorization order.</i></p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring JSRC</i></li> <li>• <i>Authorizing Subscribers with JSRC</i></li> </ul>

## authentication-order

---

<b>Syntax</b>	authentication-order [(none   ldap   password   radius   secureid)];
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(EX and QFX Series only) Configure the order of authentication, authorization, and accounting (AAA) servers to use while sending authentication messages.
<b>Default</b>	Not enabled
<b>Options</b>	<b>none</b> —No authentication for specified subscribers.  <b>ldap</b> —Lightweight Directory Access Protocol.  <b>password</b> —Locally configured password in access profile.  <b>radius</b> —RADIUS authentication.  <b>secureid</b> —RSA SecurID authentication.
<b>Required Privilege Level</b>	<b>admin</b> —To view this statement in the configuration. <b>admin-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch on page 1843</a></li><li>• <a href="#">Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 1909</a></li></ul>

## authentication-whitelist

---

<b>Syntax</b>	<pre>authentication-whitelist {   mac-address {     interface <i>interface-name</i>;     vlan-assignment ( <i>vlan-id</i>   <i>vlan-name</i> );   } }</pre>
<b>Hierarchy Level</b>	[edit ethernet-switching-options], [edit switch-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1 for EX Series switches. Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 for EX Series switches (ELS).
<b>Description</b>	Configure MAC addresses for which RADIUS authentication is to be bypassed.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Setting Up Captive Portal Authentication on an EX Series Switch</i></li> <li>• <a href="#">Example: Setting Up Captive Portal Authentication on an EX Series Switch on page 1898</a></li> <li>• <a href="#">Configuring Captive Portal Authentication (CLI Procedure)</a></li> <li>• <a href="#">Configuring Captive Portal Authentication (CLI Procedure) on page 1926</a></li> </ul>

## authenticator

---

**Syntax**    authenticator {  
              authentication-profile-name *access-profile-name*;  
              interface (all | [ *interface-names* ]) {  
                  disable;  
                  guest-vlan ( *vlan-id* | *vlan-name* );  
                  lldp-med-bypass;  
                  mac-radius <restrict>;  
                  maximum-requests *number*;  
                  no-reauthentication;  
                  quiet-period *seconds*;  
                  reauthentication *interval*;  
                  retries *number*;  
                  server-fail (deny | permit | use-cache | *vlan-id* | *vlan-name*);  
                  server-reject-vlan ( *vlan-id* | *vlan-name* ) {  
                      eapol-block;  
                      block-interval *block-interval*;  
                  }  
                  server-timeout *seconds*;  
                  supplicant (single | single-secure | multiple);  
                  supplicant-timeout *seconds*;  
                  transmit-period *seconds*;  
                  }  
                  no-mac-table-binding {  
                      interface *interface-names*;  
                      static mac-address;  
                  }  
                  radius-options {  
                      use-vlan-id;  
                      use-vlan-name;  
                  }  
                  static mac-address {  
                      *vlan-assignment* *vlan-identifier*;  
                  }  
              }

**Hierarchy Level**    [edit protocols dot1x]

**Release Information**    Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description**    Configure an authenticator for 802.1X authentication.

The statements are explained separately.



**NOTE:** You cannot configure 802.1X user authentication on interfaces that have been enabled for Q-in-Q tunneling.

---

**Default**    No static MAC address or VLAN is configured.

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring 802.1X Interface Settings (CLI Procedure) on page 1908</a></li> <li>• <a href="#">Specifying RADIUS Server Connections on an EX Series Switch (CLI Procedure) on page 1925</a></li> <li>• <a href="#">Example: Configuring Static MAC Bypass of Authentication on an EX Series Switch on page 1872</a></li> </ul>

## client-accounting-algorithm

<b>Syntax</b>	client-accounting-algorithm (direct   round-robin);
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius <b>options</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0. Statement introduced in Junos OS for EX Series switches Release 13.2X50-D10.
<b>Description</b>	Configure the access method the router uses to access RADIUS accounting servers.
<b>Default</b>	direct
<b>Options</b>	<p><b>direct</b>—Use the direct method.</p> <p><b>round-robin</b>—Use the round-robin method.</p>
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring RADIUS Server Parameters for Subscriber Access</a></li> <li>• <a href="#">Configuring RADIUS Server Options for Subscriber Access</a></li> </ul>

## client-authentication-algorithm

---

<b>Syntax</b>	client-authentication-algorithm (direct   round-robin);
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius <a href="#">options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure the access method the router uses to access RADIUS authentication servers.
<b>Default</b>	direct
<b>Options</b>	<b>direct</b> —Use the direct method.  <b>round-robin</b> —Use the round-robin method.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li><li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li></ul>

## coa-dynamic-variable-validation

---

<b>Syntax</b>	coa-dynamic-variable-validation;
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius <a href="#">options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Specify that when a CoA operation includes a change to a client profile dynamic variable that cannot be applied (such as an update to a non-existent filter), the router does not apply any changes to client profile dynamic variables in the request, and responds with a NACK message.
<b>Default</b>	If you do not configure this statement, the router does not apply any incorrect variable updates but does make any other changes to the client profile dynamic variables, and then responds with an ACK message.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li><li>• <i>RADIUS Server Options for Subscriber Access</i></li><li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li></ul>



## destination (Accounting)

```
Syntax destination {
    radius {
        server {
            server-address {
                accounting-port port-number;
                secret password;
                source-address address;
                retry number;
                timeout seconds;
            }
        }
    }
    tacplus {
        server {
            server-address {
                port port-number;
                secret password;
                single-connection;
                timeout seconds;
            }
        }
    }
}
```

**Hierarchy Level** [edit system [accounting](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
**radius** statement added in Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the authentication server.

**Options** The remaining statements are explained separately.

**Required Privilege Level** system—To view this statement in the configuration.  
 system-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring RADIUS System Accounting*
- *Configuring TACACS+ System Accounting*

## destination-host (Gx-Plus)

---

<b>Syntax</b>	<code>destination-host <i>hostname</i>;</code>
<b>Hierarchy Level</b>	[edit access gx-plus <a href="#">partition</a> <i>partition-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure the host on which the PCRF application resides.
<b>Options</b>	<i>hostname</i> —Host on which the PCRF is installed.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Gx-Plus</i></li><li>• <i>Configuring the Gx-Plus Partition</i></li></ul>

## destination-realm (Gx-Plus)

---

<b>Syntax</b>	<code>destination-realm <i>realm</i>;</code>
<b>Hierarchy Level</b>	[edit access gx-plus <a href="#">partition</a> <i>partition-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure the realm in which the PCRF host resides.
<b>Options</b>	<i>realm</i> —Realm in which the PCRF host resides.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Gx-Plus</i></li><li>• <i>Configuring the Gx-Plus Partition</i></li></ul>

---

## diameter-instance (Gx-Plus)

---

<b>Syntax</b>	<code>diameter-instance <i>instance-name</i>;</code>
<b>Hierarchy Level</b>	[edit access gx-plus <a href="#">partition</a> <i>partition-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Specify the Diameter instance associated with the Gx-Plus partition.
<b>Options</b>	<i>instance-name</i> —Name of the Diameter instance. Currently, only <b>master</b> is supported.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Gx-Plus</i></li><li>• <i>Configuring the Gx-Plus Partition</i></li></ul>

## domain (Domain Map)

---

**Syntax**    domain {  
              delimiter [*delimiter-character*];  
              map *domain-map-name* {  
                  aaa-logical-system *logical-system-name* {  
                      aaa-routing-instance *routing-instance-name*;  
                      }  
                  aaa-routing-instance *routing-instance-name*;  
                  access-profile *profile-name*;  
                  address-pool *pool-name*;  
                  dynamic-profile *profile-name*;  
                  padn *destination-address* {  
                      mask *destination-mask*;  
                      metric *route-metric*;  
                      }  
                  strip-domain;  
                  target-logical-system *logical-system-name* {  
                      target-routing-instance *routing-instance-name*;  
                      }  
                  target-routing-instance *routing-instance-name*;  
                  tunnel-profile *profile-name*;  
                  }  
              parse-direction (left-to-right | right-to-left);  
              parse-order (domain-first | realm-first);  
              realm-delimiter [*delimiter-character*];  
              realm-parse-direction (left-to-right | right-to-left);  
              }  
              }

**Hierarchy Level**    [edit access]

**Release Information**    Statement introduced in Junos OS Release 10.4.

**Description**    Configure a domain map, which is used to map access options and session parameters for subscriber sessions.

The remaining statements are explained separately.

**Required Privilege Level**    admin—To view this statement in the configuration.  
                                  admin-control—To add this statement to the configuration.

**Related Documentation**    • *Configuring a Domain Map*

## domain-name-server (Routing Instances and Access Profiles)

<b>Syntax</b>	<code>domain-name-server <i>dns-address</i>;</code>
<b>Hierarchy Level</b>	[edit access], [edit access profile]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure an IPv4 address for a DNS name server. You can configure an address globally for a routing instance at the <b>[edit access]</b> hierarchy level or for an access profile at the <b>[edit access profile <i>profile-name</i>]</b> hierarchy level. You can configure more than one address by including the statement multiple times.



**NOTE:** A DNS name server address configured with this statement is lower in preference than one configured with the [domain-name-server-inet](#) statement.

<b>Options</b>	<i>dns-address</i> —IPv4 address of the DNS name server.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring DNS Name Server Addresses for Subscriber Management</i></li> <li><i>DNS Name Server Address Overview</i></li> </ul>

## domain-name-server-inet (Routing Instances and Access Profiles)

---

<b>Syntax</b>	<code>domain-name-server-inet <i>dns-address</i>;</code>
<b>Hierarchy Level</b>	[edit access], [edit access profile]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure an IPv4 address for a DNS name server. You can configure an address globally for a routing instance at the [edit access] hierarchy level or for an access profile at the [edit access profile <i>profile-name</i> ] hierarchy level. You can configure more than one address by including the statement multiple times.



**NOTE:** A DNS name server address configured with this statement is higher in preference than one configured with the `domain-name-server` statement.

---

<b>Options</b>	<i>dns-address</i> —IPv4 address of the DNS name server.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring DNS Name Server Addresses for Subscriber Management</i></li><li>• <i>DNS Name Server Address Overview</i></li></ul>

## domain-name-server-inet6 (Routing Instances and Access Profiles)

<b>Syntax</b>	<code>domain-name-server-inet6 <i>dns-address</i>;</code>
<b>Hierarchy Level</b>	<code>[edit access],</code> <code>[edit access profile]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure an IPv6 address for a DNS name server. You can configure an address globally for a routing instance at the <b>[edit access]</b> hierarchy level or for an access profile at the <b>[edit access profile <i>profile-name</i>]</b> hierarchy level. You can configure more than one address by including the statement multiple times.
<b>Options</b>	<i>dns-address</i> —IPv6 address of the DNS name server.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring DNS Name Server Addresses for Subscriber Management</i></li> <li>• <i>DNS Name Server Address Overview</i></li> </ul>

## ethernet-port-type-virtual

<b>Syntax</b>	<code>ethernet-port-type-virtual;</code>
<b>Hierarchy Level</b>	<code>[edit access profile <i>profile-name</i> radius options]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches.
<b>Description</b>	Specify the physical port type the router or switch uses to authenticate clients. The router or switch passes a port type of <b>ethernet</b> in RADIUS attribute 61 (NAS-Port-Type) by default. This statement specifies a port type of <b>virtual</b> .



**NOTE:** This statement takes precedence over the **nas-port-type** statement if you include both statements in the same access profile.

<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li> <li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li> </ul>

## global (Gx-Plus)

---

<b>Syntax</b>	<pre>global {     include-ipv6;     max-outstanding-requests <i>number</i>; }</pre>
<b>Hierarchy Level</b>	[edit access <a href="#">gx-plus</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure global attributes for the Gx-Plus application.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Gx-Plus</a></li></ul>

## gx-plus (Gx-Plus)

---

<b>Syntax</b>	<pre>gx-plus {     global {         include-ipv6;         max-outstanding-requests <i>number</i>;     }     partition <i>partition-name</i> {         diameter-instance <i>instance-name</i>;         destination-host <i>hostname</i>;         destination-realm <i>realm</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit access]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure the Gx-Plus application to interact with a PCRF to authorize and provision subscribers.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Gx-Plus</a></li></ul>



## ignore

---

<b>Syntax</b>	<pre>ignore {     dynamic-iflset-name;     framed-ip-netmask;     input-filter;     logical-system-routing-instance;     output-filter; }</pre>
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius attributes]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p>
<b>Description</b>	Configure the router or switch to ignore the specified attributes in RADIUS Access-Accept messages. By default, the router or switch processes the attributes it receives from the external server.
<b>Options</b>	<p><b>dynamic-iflset-name</b>—Ignore Interface-Set/Dynamic-Ifset-Name (VSA 26-130).</p> <p><b>framed-ip-netmask</b>—Ignore Framed-IP-Netmask (RADIUS attribute 9).</p> <p><b>input-filter</b>—Ignore Ingress-Policy-Name (VSA 26-10).</p> <p><b>logical-system-routing-instance</b>—Ignore Virtual-Router (VSA 26-1).</p> <p><b>output-filter</b>—Ignore Egress-Policy-Name (VSA 26-11).</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li> </ul>

## include-ipv6 (Gx-Plus)

---

<b>Syntax</b>	include-ipv6;
<b>Hierarchy Level</b>	[edit access gx-plus <a href="#">global</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Include IPv6 subscribers in Gx-Plus provisioning requests.
<b>Default</b>	By default, IPv6 subscribers are not included.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Gx-Plus Global Attributes</i></li><li>• <i>Configuring Gx-Plus</i></li></ul>

## interface (Static MAC Bypass)

<b>Syntax</b>	<code>interface [<i>interface-names</i>];</code>
<b>Hierarchy Level</b>	[edit protocols dot1x <a href="#">authenticator</a> authentication-profile-name static <i>mac-address</i> ], [edit ethernet-switching-options <a href="#">authentication-whitelist</a> <i>mac-address</i> ], [edit <a href="#">switch-options on page 2212</a> <a href="#">authentication-whitelist</a> <i>mac-address</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement added to the <a href="#">[edit ethernet-switching-options authentication-whitelist]</a> hierarchy in Junos OS Release 10.1 for EX Series switches. Statement added to the <a href="#">[edit switch-options authentication-whitelist]</a> hierarchy in Junos OS Release 13.2X50-D10 for EX Series switches (ELS).
<b>Description</b>	Configure interfaces on which the specified MAC addresses are allowed to bypass RADIUS authentication and allowed to connect to the LAN without authentication.
<b>Options</b>	<i>interface-names</i> —List of interfaces.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show dot1x static-mac-address on page 2016</a></li> <li>• <a href="#">Example: Configuring Static MAC Bypass of Authentication on an EX Series Switch on page 1872</a></li> <li>• <a href="#">Example: Setting Up Captive Portal Authentication on an EX Series Switch</a></li> <li>• <a href="#">Example: Setting Up Captive Portal Authentication on an EX Series Switch</a></li> <li>• <a href="#">Configuring Captive Portal Authentication (CLI Procedure)</a></li> </ul>

## interface (VoIP)

---

<b>Syntax</b>	interface (all   [ <i>interface-name</i> ]   access-ports) { vlan <i>vlan-name</i> ; forwarding-class <assured-forwarding   best-effort   expedited-forwarding   network-control>; }
<b>Hierarchy Level</b>	<ul style="list-style-type: none"><li>• For platforms with ELS:     [edit <a href="#">switch-options</a> on page 2212 <b>voip</b>]</li><li>• For platforms without ELS:     [edit ethernet-switching-options <b>voip</b>],</li></ul>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Hierarchy level [edit <b>switch-options</b> ] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software”</a> on page 3 for information about ELS.)
<b>Description</b>	Enable voice over IP (VoIP) for all interfaces or specific interfaces.
<b>Options</b>	all   <i>interface-name</i>   access-ports—Enable VoIP on all interfaces, on a specific interface, or on all access ports.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch</i></li><li>• <i>Example: Configuring VoIP on an EX Series Switch Without Including 802.1X Authentication</i></li><li>• <i>Example: Configuring VoIP on an EX Series Switch Without Including LLDP-MED Support</i></li></ul>

## interface-description-format

---

<b>Syntax</b>	<pre>interface-description-format {     exclude-adapter;     exclude-sub-interface; }</pre>
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius <a href="#">options</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>Options <b>exclude-adapter</b> and <b>exclude-sub-interface</b> introduced in Junos OS Release 10.4.</p>
<b>Description</b>	Specify the information that is excluded from the interface description that the device passes to RADIUS for inclusion in the RADIUS attribute 87 (NAS-Port-Id). By default, the device includes both the subinterface and the adapter in the interface description.
<b>Options</b>	<p><b>exclude-adapter</b>—Exclude the adapter from the interface description.</p> <p><b>exclude-sub-interface</b>—Exclude the subinterface from the interface description.</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li> <li>• <i>RADIUS Server Options for Subscriber Access</i></li> </ul>

## juniper-dsl-attributes

---

<b>Syntax</b>	juniper-dsl-attributes;
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius <a href="#">options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	<p>Configure AAA to add Juniper Networks DSL VSAs to the RADIUS authentication and accounting request messages for subscribers. If the router has not received and processed the corresponding ANCP attributes from the access node, then AAA provides only the following in these RADIUS messages:</p> <ul style="list-style-type: none"><li>• Downstream-Calculated-QoS-Rate (IANA 4874, 26-141)—Default configured advisory transmit speed.</li><li>• Upstream-Calculated-QoS-Rate (IANA 4874, 26-142)—Default configured advisory receive speed.</li></ul>
<b>Default</b>	The Juniper Networks DSL VSAs are not added to the RADIUS authentication and accounting request messages. However, the DSL Forum VSA—if available—is added to RADIUS messages by default.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages</i></li><li>• <i>Configuring the ANCP Agent</i></li></ul>

## max-outstanding-requests (Gx-Plus)

<b>Syntax</b>	max-outstanding-requests <i>number</i> ;
<b>Hierarchy Level</b>	[edit access gx-plus <a href="#">global</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Limit the number of outstanding requests to the PCRF that Gx-Plus can retry when the requests are improperly answered. Too many requests risks overloading the PCRF and increases the chance of losing messages.
<b>Options</b>	<i>number</i> —Number of outstanding requests from Gx-Plus to the PCRF that can exist at any time. <b>Default:</b> 40 <b>Range:</b> 2 through 40
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Gx-Plus Global Attributes</i></li> <li>• <i>Configuring Gx-Plus</i></li> </ul>

## nas-identifier

<b>Syntax</b>	nas-identifier <i>identifier-value</i> ;
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius <a href="#">options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches.
<b>Description</b>	Configure the value for the client RADIUS attribute 32 (NAS-Identifier). This attribute is used for authentication and accounting requests.
<b>Options</b>	<i>identifier-value</i> —String to use for authentication and accounting requests. <b>Range:</b> 1 through 64 characters
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li> <li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li> </ul>

## nas-port-extended-format (Access Profile)

---

**Syntax**    `nas-port-extended-format {  
              adapter-width width;  
              ae-width width;  
              port-width width;  
              slot-width width;  
              stacked-vlan-width width;  
              vlan-width width;  
              atm {  
                  adapter-width width;  
                  port-width width;  
                  slot-width width;  
                  vci-width width;  
                  vpi-width width;  
              }  
              }`

**Hierarchy Level**    [edit access profile *profile-name* radius [options](#)]

**Release Information**    Statement introduced in Junos OS Release 9.1.  
Statement introduced in Junos OS Release 9.1 for EX Series switches.  
Option **ae-width** introduced in Junos OS Release 12.1.  
Option **stacked** introduced in Junos OS Release 12.3.  
Option **atm** introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.  
Option **atm** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

**Description**    In an access profile, configure the RADIUS client to use the extended format for RADIUS attribute 5 (NAS-Port) and specify the width of the fields in the NAS-Port attribute. You can use the same access profile to configure the NAS-Port extended format for Ethernet subscribers and ATM subscribers.

**Options**    **adapter-width *width***—Number of bits in the adapter field.

**ae-width *width***—Number of bits in the aggregated Ethernet identifier field.

**port-width *width***—Number of bits in the port field.

**slot-width *width***—Number of bits in the slot field.

**stacked**—Include stacked VLAN IDs, in addition to VLAN IDs, in the NAS-Port extended format.

**stacked-vlan-width *width***—Number of bits in the SVLAN ID field.

**vlan-width *width***—Number of bits in the VLAN ID field.

**atm**—Configure the NAS-Port extended format for ATM subscribers; options include:

- **adapter-width *width***—Number of bits in the adapter field.
- **port-width *width***—Number of bits in the port field.



- **slot-width *width***—Number of bits in the slot field.
- **vci-width *width***—Number of bits in the ATM virtual circuit identifier (VCI) field.
- **vpi-width *width***—Number of bits in the ATM virtual path identifier (VPI) field.



**NOTE:** Each field can be 0 through 32 bits wide; however, the total of the widths of all fields must not exceed 32 bits, or the configuration fails.

The router may truncate the values of individual fields depending on the bit width you specify.

<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li> <li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li> </ul>

## nas-port-id-delimiter (Subscriber Management)

<b>Syntax</b>	nas-port-id-delimiter <i>delimiter-character</i> ;
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius <a href="#">options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Specify the character that the router uses as a separator between the concatenated values in the NAS-Port-ID string. The router uses the delimiter when you configure more than one value in the <b>nas-port-id-format</b> statement.
<b>Default</b>	The hash (#) character.
<b>Options</b>	<b><i>delimiter-character</i></b> —Character used for the delimiter.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li> <li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li> <li>• <i>Configuring a NAS-Port-ID with Additional Options</i></li> </ul>

## nas-port-id-format (Subscriber Management)

---

<b>Syntax</b>	<pre>nas-port-id-format {   agent-circuit-id;   agent-remote-id;   interface-description;   nas-identifier; }</pre>
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius <a href="#">options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Specify the information that the router includes in the NAS-Port-ID (RADIUS attribute 87) that it is passed to the RADIUS server during authentication and accounting. You can include any combination of the optional values.
<b>Default</b>	The router includes the interface description.
<b>Options</b>	<p><b>agent-circuit-id</b>—Include the agent circuit ID from either DHCP option 82 or the DSL forum VSAs.</p> <p><b>agent-remote-id</b>—Include the agent remote ID from either DHCP option 82 or the DSL forum VSAs.</p> <p><b>interface-description</b>—Include the interface description.</p> <p><b>nas-identifier</b>—Include the NAS identifier value (RADIUS attribute 32).</p>
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li><li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li><li>• <i>Configuring a NAS-Port-ID with Additional Options</i></li></ul>

## nas-port-type (Subscriber Management)

**Syntax**    `nas-port-type {  
          ethernet {  
            port-type;  
          }  
}`

**Hierarchy Level**    [edit access profile *profile-name* radius [options](#)]

**Release Information**    Statement introduced in Junos OS Release 11.4.  
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

**Description**    Specify the port type used to authenticate subscribers. The router includes the port type in RADIUS attribute 61 (NAS-Port-Type attribute).



**NOTE:** This statement is ignored if the [ethernet-port-type-virtual](#) statement is included in the same access profile.

**Default**    The router uses a port type of **ethernet**.

**Options**    *port-type*—One of the following port types:

- *value*—A value from 0-65535
- **adsl-cap**—Asymmetric DSL, carrierless amplitude phase (CAP) modulation
- **adsl-dmt**—Asymmetric DSL, discrete multitone (DMT)
- **async**—Asynchronous
- **cable**—Cable
- **ethernet**—Ethernet
- **fddi**—Fiber Distributed Data Interface
- **g3-fax**—G.3 Fax
- **hdlc-clear-channel**—HDLC Clear Channel
- **iapp**—Inter-Access Point Protocol (IAPP)
- **isdsl**—ISDN DSL
- **isdn-sync**—ISDN Synchronous
- **isdn-v110**—ISDN Async V.110
- **isdn-v120**—ISDN Async V.120
- **piafs**—Personal Handyphone System (PHS) Internet Access Forum Standard
- **sdsl**—Symmetric DSL

- **sync**—Synchronous
- **token-ring**—Token Ring
- **virtual**—Virtual
- **wireless**—Other wireless
- **wireless-1x-ev**—Wireless 1xEV
- **wireless-cdma2000**—Wireless code division multiple access (CDMA) 2000
- **wireless-ieee80211**—Wireless 802.11
- **wireless-umts**—Wireless universal mobile telecommunications system (UMTS)
- **x25**—X.25
- **x75**—X.75
- **xdsl**—DSL of unknown type

<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
---------------------------------	---

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li><li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li></ul>
------------------------------	--

## options (Access Profile)

```
Syntax  options {
        accounting-session-id-format (decimal | description);
        calling-station-id-delimiter delimiter-character;
        calling-station-id-format {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            nas-identifier;
        }
        client-accounting-algorithm (direct | round-robin);
        client-authentication-algorithm (direct | round-robin);
        coa-dynamic-variable-validation;
        ethernet-port-type-virtual;
        access-loop-id-local;
        interface-description-format {
            exclude-adapter;
            exclude-sub-interface;
        }
        ip-address-change-notify message;
        juniper-dsl-attributes;
        nas-identifier identifier-value;
        nas-port-extended-format {
            adapter-width width;
            ae-width width;
            port-width width;
            slot-width width;
            stacked-vlan-width width;
            vlan-width width;
            atm {
                adapter-width width;
                port-width width;
                slot-width width;
                vci-width width;
                vpi-width width;
            }
        }
        nas-port-id-delimiter delimiter-character;
        nas-port-id-format {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            nas-identifier;
        }
        nas-port-type {
            ethernet {
                port-type;
            }
        }
        revert-interval interval;
        vlan-nas-port-stacked-format;
    }
```

<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> <b>radius</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches.
<b>Description</b>	Configure the options used by RADIUS authentication and accounting servers.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li><li>• <i>RADIUS Server Options for Subscriber Access</i></li></ul>

---

## partition (Gx-Plus)

---

<b>Syntax</b>	<pre>partition <i>partition-name</i> {     <b>diameter-instance</b> <i>instance-name</i>;     <b>destination-host</b> <i>hostname</i>;     <b>destination-realm</b> <i>realm</i>; }</pre>
<b>Hierarchy Level</b>	[edit access <b>gx-plus</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure a Gx-Plus partition.
<b>Options</b>	<b><i>partition-name</i></b> —Name of the Gx-Plus partition.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Gx-Plus</i></li><li>• <i>Configuring the Gx-Plus Partition</i></li></ul>

## port

---

<b>Syntax</b>	<code>port <i>port-number</i>;</code>
<b>Hierarchy Level</b>	[edit access radius-server <i>server-address</i> ], [edit access profile <i>profile-name</i> radius-server <i>server-address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the port number on which to contact the RADIUS server.
<b>Options</b>	<b><i>port-number</i></b> —Port number on which to contact the RADIUS server. <b>Default:</b> 1812 (as specified in RFC 2865)
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Router or Switch Interaction with RADIUS Servers</i></li> <li>• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i></li> </ul>

## provisioning-order

---

<b>Syntax</b>	provisioning-order (gx-plus   jsrc);
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Support for Gx-Plus introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure AAA to use the specified application for subscriber service provisioning.
<b>Options</b>	<p><b>gx-plus</b>—Specify Gx-Plus as the application used to communicate with a PCRF for subscriber service provisioning.</p> <p><b>jsrc</b>—Specify JSRC as the application used to communicate with the SAE for subscriber service provisioning. JSRC is used in an SRC environment to request services from the SAE for an authenticated subscriber. JSRC attempts to activate these services. If successful, JSRC returns an ACK message. If unsuccessful, the subscriber is denied access.</p>
<b>Required Privilege Level</b>	<p><b>admin</b>—To view this statement in the configuration.</p> <p><b>admin-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring JSRC</i></li><li>• <i>Provisioning Subscribers with JSRC</i></li><li>• <i>Configuring Gx-Plus</i></li><li>• <i>Provisioning Subscribers with Gx-Plus</i></li></ul>



## radius (Access Profile)

```
Syntax  radius {
        accounting-server [ ip-address ];
        attributes {
            exclude
            ...
        }
        ignore {
            framed-ip-netmask;
            input-filter;
            logical-system-routing-instance;
            output-filter;
        }
    }
    authentication-server [ ip-address ];
    options {
        accounting-session-id-format (decimal | description);
        calling-station-id-delimiter delimiter-character;
        calling-station-id-format {
            agent-circuit-id;
            agent-remote-id;
            interface-description;
            nas-identifier;
        }
        client-accounting-algorithm (direct | round-robin);
        client-authentication-algorithm (direct | round-robin);
        coa-dynamic-variable-validation;
        ethernet-port-type-virtual;
        interface-description-format {
            exclude-adapter;
            exclude-sub-interface;
        }
        ip-address-change-notify message;
        juniper-dsl-attributes;
        nas-identifier identifier-value;
        nas-port-extended-format {
            adapter-width width;
            ae-width width;
            port-width width;
            slot-width width;
            stacked-vlan-width width;
            vlan-width width;
            atm {
                adapter-width width;
                port-width width;
                slot-width width;
                vci-width width;
                vpi-width width;
            }
        }
        nas-port-id-delimiter delimiter-character;
        nas-port-id-format {
            agent-circuit-id;

```

```
    agent-remote-id;
    interface-description;
    nas-identifier;
  }
  nas-port-type {
    ethernet {
      port-type;
    }
  }
  revert-interval interval;
  vlan-nas-port-stacked-format;
}
preauthentication-server ip-address;
}
```

**Hierarchy Level** [edit access profile *profile-name*]

**Release Information** Statement introduced in Junos OS Release 9.1.  
Statement introduced in Junos OS Release 9.1 for EX Series switches.

**Description** Configure the RADIUS parameters that the router uses for AAA authentication and accounting for subscribers.

The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring RADIUS Server Parameters for Subscriber Access*
- *RADIUS Server Options for Subscriber Access*

## radius (System)

<b>Syntax</b>	<pre>radius {   server {     server-address {       accounting-port <i>port-number</i>;       secret <i>password</i>;       source-address <i>address</i>;       retry <i>number</i>;       timeout <i>seconds</i>;     }   } }</pre>
<b>Hierarchy Level</b>	[edit system accounting destination]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure the RADIUS accounting server.
<b>Options</b>	<p><b><i>server-address</i></b>—Address of the RADIUS accounting server.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring RADIUS System Accounting</i></li> </ul>

## radius-options (Protocols 802.1X)

---

<b>Syntax</b>	<pre>radius-options {     use-vlan-id;     use-vlan-name; }</pre>
<b>Hierarchy Level</b>	[edit protocols dot1x <a href="#">authenticator</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	Configure 802.1X authenticator so that the VLAN ID or VLAN name is included in the packet sent to the RADIUS server to request authentication.
<b>Options</b>	<p><b>use-vlan-id</b>—Include the VLAN ID in the packet sent to the RADIUS server to request authentication.</p> <p><b>use-vlan-name</b>—Include the VLAN name in the packet sent to the RADIUS server to request authentication. The VLAN name is sent even if the 802.1X interface is configured with the VLAN ID.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring 802.1X Interface Settings (CLI Procedure) on page 1908</a></li><li>• <a href="#">Specifying RADIUS Server Connections on an EX Series Switch (CLI Procedure) on page 1925</a></li><li>• <a href="#">authenticator on page 1950</a></li></ul>

## radius-options (Access)

<b>Syntax</b>	<code>radius-options {     revert-interval <i>seconds</i>; }</code>
<b>Hierarchy Level</b>	[edit access], [edit access profile <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Release 8.5 of Junos OS.
<b>Description</b>	Configure RADIUS options.
<b>Options</b>	The remaining statement is explained separately.
<b>Required Privilege Level</b>	access—To view this statement in the configuration. access-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Junos OS Security Configuration Guide</i></li> </ul>

## radius-server (System)

<b>Syntax</b>	<code>radius-server <i>server-address</i> {     accounting-port <i>port-number</i>;     port <i>number</i>;     retry <i>number</i>;     secret <i>password</i>;     source-address <i>source-address</i>;     timeout <i>seconds</i>; }</code>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>Configure a RADIUS server for Point-to-Point Protocol (PPP).</p> <p>To configure multiple RADIUS servers, include multiple <b>radius-server</b> statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.</p>
<b>Options</b>	<p><b><i>server-address</i></b>—Address of the RADIUS authentication server.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring RADIUS Authentication</i></li> </ul>

## retry

---

<b>Syntax</b>	<code>retry <i>attempts</i>;</code>
<b>Hierarchy Level</b>	[edit access radius-server <i>server-address</i> ], [edit access profile <i>profile-name</i> radius-server <i>server-address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the number of times that the router or switch is allowed to attempt to contact a RADIUS authentication or accounting server.
<b>Options</b>	<b><i>attempts</i></b> —Number of times that the router is allowed to attempt to contact a RADIUS server. <b>Range:</b> 1 through 30 <b>Default:</b> 3
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i></li><li>• <i>Configuring Router or Switch Interaction with RADIUS Servers</i></li><li>• <i>Example: Configuring CHAP Authentication with RADIUS</i></li><li>• <i>Configuring RADIUS Authentication for L2TP</i></li><li>• <a href="#">timeout on page 1987</a></li></ul>

## revert-interval

<b>Syntax</b>	<code>revert-interval <i>interval</i>;</code>
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius <b>options</b> ], [edit access radius-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches.
<b>Description</b>	Configure the amount of time the router or switch waits after a server has become unreachable. The router or switch rechecks the connection to the server when the specified interval expires. If the server is then reachable, it is used in accordance with the order of the server list.
<b>Options</b>	<i>interval</i> —Amount of time to wait. <b>Range:</b> 0 through 604800 seconds <b>Default:</b> 60 seconds
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li> <li>• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i></li> </ul>

## routing-instance

<b>Syntax</b>	<code>routing-instance <i>routing-instance-name</i>;</code>
<b>Hierarchy Level</b>	[edit access radius-server <i>server-address</i> ], [edit access profile <i>profile-name</i> radius-server <i>server-address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the routing instance used to send RADIUS packets to the RADIUS server.
<b>Options</b>	<i>routing-instance-name</i> —Routing instance name.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the PPP Authentication Protocol</i></li> <li>• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i></li> </ul>

## secret

---

<b>Syntax</b>	<code>secret password;</code>
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius-server <i>server-address</i> ], [edit access radius-disconnect <i>client-address</i> ], [edit access radius-server <i>server-address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the password to use with the RADIUS server. The secret password used by the local router or switch must match that used by the server.
<b>Options</b>	<i>password</i> —Password to use; it can include spaces if the character string is enclosed in quotation marks.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i></li><li>• <i>Configuring Router or Switch Interaction with RADIUS Servers</i></li><li>• <i>Example: Configuring CHAP Authentication with RADIUS</i></li><li>• <i>Configuring RADIUS Authentication for L2TP</i></li><li>• <i>Configuring the RADIUS Disconnect Server for L2TP</i></li><li>• <i>Configuring an EX Series Switch to Use Junos Pulse Access Control Service for Network Access Control (CLI Procedure)</i></li></ul>

## send-acct-status-on-config-change (Access Profile)

---

<b>Syntax</b>	<code>send-acct-status-on-config-change;</code>
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> <b>accounting</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.1. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure the router's authd process to send an Acct-On message when the first RADIUS server is added to an access profile, and to send an Acct-Off message when the last RADIUS server is removed from an access profile.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li><li>• <i>Configuring Per-Subscriber Session Accounting</i></li></ul>



## server (RADIUS Accounting)

<b>Syntax</b>	<pre>server {   server-address {     accounting-port <i>port-number</i>;     retry <i>number</i>     secret <i>password</i>;     source-address <i>address</i>;     timeout <i>seconds</i>;   } }</pre>
<b>Hierarchy Level</b>	[edit system accounting destination radius]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure RADIUS logging.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring RADIUS System Accounting</i></li> </ul>

## service (Service Accounting)

<b>Syntax</b>	<pre>service {   accounting-order (activation-protocol   radius); }</pre>
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Define the subscriber service accounting configuration.  The remaining statement is explained separately.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Service Accounting with JSRC</i></li> <li>• <i>Service Accounting with JSRC</i></li> </ul>

## source-address

---

<b>Syntax</b>	<code>source-address <i>source-address</i>;</code>
<b>Hierarchy Level</b>	<code>[edit access radius-server <i>server-address</i>],</code> <code>[edit access profile <i>profile-name</i> radius-server <i>server-address</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address.
<b>Options</b>	<b><i>source-address</i></b> —Valid IPv4 address configured on one of the router or switch interfaces. On M Series routers only, the source address can be an IPv6 address and the UDP source port is 514.
<b>Required Privilege Level</b>	<code>admin</code> —To view this statement in the configuration. <code>admin-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Router or Switch Interaction with RADIUS Servers</i></li><li>• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i></li><li>• <i>Example: Configuring CHAP Authentication with RADIUS</i></li><li>• <i>Configuring RADIUS Authentication for L2TP</i></li></ul>

## timeout (RADIUS)

---

<b>Syntax</b>	<code>timeout seconds;</code>
<b>Hierarchy Level</b>	[edit access radius-server <i>server-address</i> ], [edit access profile <i>profile-name</i> radius-server <i>server-address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure the amount of time that the local router or switch waits to receive a response from a RADIUS server.
<b>Options</b>	<b>seconds</b> —Amount of time to wait. <b>Range:</b> 1 through 90 seconds <b>Default:</b> 3 seconds
<b>Required Privilege Level</b>	<b>system</b> —To view this statement in the configuration. <b>system-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Router or Switch Interaction with RADIUS Servers</i></li> <li>• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i></li> <li>• <i>Example: Configuring CHAP Authentication with RADIUS</i></li> <li>• <i>Configuring RADIUS Authentication for L2TP</i></li> </ul>

## vlan (VoIP)

---

<b>Syntax</b>	<code>vlan (vlan-id   vlan-name   untagged);</code>
<b>Hierarchy Level</b>	[edit ethernet-switching-options <b>voip interface (VoIP)</b> (all   [ <i>interface-name</i>   access-ports])
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	For EX Series switches, specify the VLAN name or VLAN tag identifier associated with the VLAN to be sent from the authenticating server to the IP phone.
<b>Options</b>	<p><i>vlan-name</i>—Name of a VLAN.</p> <p><i>vlan-id</i>—The VLAN tag identifier.</p> <p><b>Range:</b> 0 through 4095. Tags 0 and 4095 are reserved by Junos OS, and you should not configure them.</p> <p><i>untagged</i>—Allow untagged VLAN traffic.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch</i></li><li>• <i>Example: Configuring VoIP on an EX Series Switch Without Including 802.1X Authentication</i></li><li>• <i>Example: Configuring VoIP on an EX Series Switch Without Including LLDP-MED Support</i></li></ul>

## vlan-assignment

<b>Syntax</b>	<code>vlan-assignment (vlan-id   vlan-name);</code>
<b>Hierarchy Level</b>	[edit protocols dot1x <a href="#">authenticator</a> authentication-profile-name static (Protocols 802.1X) <i>mac-address</i> ], [edit ethernet-switching-options <a href="#">authentication-whitelist</a> ], [edit <a href="#">switch-options</a> on page 2212 <a href="#">authentication-whitelist</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement added to the <a href="#">[edit ethernet-switching-options authentication-whitelist]</a> hierarchy in Junos OS Release 10.1 for EX Series switches. Statement added to the <a href="#">[edit switch-options authentication-whitelist]</a> hierarchy in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure the VLAN that is associated with the list of MAC addresses that are excluded from RADIUS authentication.
<b>Options</b>	<i>vlan-id   vlan-name</i> —The name of the VLAN or the VLAN tag identifier to associate with the device. The VLAN already exists on the switch.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show dot1x static-mac-address on page 2016</a></li> <li>• <a href="#">Example: Configuring Static MAC Bypass of Authentication on an EX Series Switch on page 1872</a></li> <li>• <a href="#">Example: Setting Up Captive Portal Authentication on an EX Series Switch</a></li> <li>• <a href="#">Understanding Authentication on EX Series Switches on page 1824</a></li> <li>• <a href="#">Example: Setting Up Captive Portal Authentication on an EX Series Switch</a></li> <li>• <a href="#">Configuring Captive Portal Authentication (CLI Procedure)</a></li> </ul>

## vlan-nas-port-stacked-format

---

<b>Syntax</b>	vlan-nas-port-stacked-format;
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius <a href="#">options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches.
<b>Description</b>	Configure RADIUS attribute 5 (NAS-Port) to include the S-VLAN ID, in addition to the VLAN ID, for subscribers on Ethernet interfaces.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring RADIUS Server Options for Subscriber Access</i></li><li>• <i>Configuring Authentication and Accounting Parameters for Subscriber Access</i></li></ul>

## voip

---

<b>Syntax</b>	<pre>voip {   <a href="#">interface</a> (all   [<i>interface-name</i>   access-ports]) {     <a href="#">vlan</a> <i>vlan-name</i> ;     forwarding-class &lt;assured-forwarding   best-effort   expedited-forwarding         network-control&gt;;   } }</pre>
<b>Hierarchy Level</b>	[edit ethernet-switching-options], [edit <a href="#">switch-options</a> on page 2212]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure voice over IP (VoIP) interfaces.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Setting Up VoIP with 802.1X and LLDP-MED on an EX Series Switch</i></li><li>• <i>Example: Configuring VoIP on an EX Series Switch Without Including 802.1X Authentication</i></li><li>• <i>Example: Configuring VoIP on an EX Series Switch Without Including LLDP-MED Support</i></li></ul>

---

## wait-for-acct-on-ack (Access Profile)

---

<b>Syntax</b>	wait-for-acct-on-ack;
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> <b>accounting</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Configure the router's <b>authd</b> process to wait for an Acct-On-Ack response message from RADIUS before sending new authentication and accounting updates to the RADIUS server. This configuration ensures that when a new subscriber session starts, the authentication and accounting information for the new session does not get deleted when RADIUS clears previously existing session state information.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring RADIUS Server Parameters for Subscriber Access</i></li><li>• <i>Configuring Per-Subscriber Session Accounting</i></li></ul>





## CHAPTER 29

# Administration

- [Routine Monitoring on page 1993](#)
- [Operational Commands on page 1995](#)

## Routine Monitoring

---

- [Monitoring 802.1X Authentication on page 1993](#)
- [Verifying 802.1X Authentication on page 1994](#)

## Monitoring 802.1X Authentication

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring feature to display details of authenticated users and users who have failed authentication.

### Action

To display authentication details in the J-Web interface, select **Monitoring > Security > 802.1X**.

To display authentication details in the CLI, enter the following commands:

- **show dot1x interface detail | display xml**
- **show dot1x interface detail <interface> | display xml**
- **show dot1x auth-failed-users**

### Meaning

The details displayed include:

- A list of authenticated users.
- The total number of users connected.
- A list of users who have failed authentication

You can also specify an interface for which the details must be displayed.

- Related Documentation**
- [Configuring 802.1X Authentication \(J-Web Procedure\)](#)
  - [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)
  - [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)

## Verifying 802.1X Authentication

- Purpose** Verify that supplicants are being authenticated on an interface on an EX Series switch with the interface configured for 802.1X authentication, and display the method of authentication being used.
- Action** Display detailed information about an interface configured for 802.1X (here, the interface is **ge-0/0/16**):
- ```
user@switch> show dot1x interface ge-0/0/16.0 detail
ge-0/0/16.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Strict: Disabled
  Reauthentication: Enabled Reauthentication interval: 40 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 1
  Guest VLAN member: <not configured>
  Number of connected supplicants: 1
    Supplicant: user5, 00:30:48:8C:66:BD
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: v200
      Reauthentication due in 17 seconds
```
- Meaning** The sample output from the **show dot1x interface detail** command shows that the **Number of connected supplicants** is 1. The supplicant that was authenticated and is now connected to the LAN is known as **user5** on the RADIUS server and has the MAC address **00:30:48:8C:66:BD**. The supplicant was authenticated by means of the 802.1X authentication method called **Radius** authentication. When the **Radius** authentication method is used, the supplicant is configured on the RADIUS server, the RADIUS server communicates this to the switch, and the switch opens LAN access on the interface to which the supplicant is connected. The sample output also shows that the supplicant is connected to VLAN **v200**.
- Other 802.1X authentication methods supported on EX Series switches in addition to the **RADIUS** method are:

- **Guest VLAN**—A nonresponsive host is granted Guest-VLAN access.
- **MAC Radius**—A nonresponsive host is authenticated based on its MAC address. The MAC address is configured as permitted on the RADIUS server, the RADIUS server lets the switch know that the MAC address is a permitted address, and the switch opens LAN access to the nonresponsive host on the interface to which it is connected.
- **Server-fail deny**—If the RADIUS servers time out, all supplicants are denied access to the LAN, preventing traffic from flowing from the supplicant through the interface. This is the default.
- **Server-fail permit**—When the RADIUS server is unavailable, a supplicant is still permitted access to the LAN as if the supplicant had been successfully authenticated by the RADIUS server.
- **Server-fail use-cache**—If the RADIUS servers time out during reauthentication, previously authenticated supplicants are granted access, but new supplicants are denied LAN access.
- **Server-fail VLAN**—A supplicant is configured to be moved to a specified VLAN if the RADIUS server is unavailable to reauthenticate the supplicant. (The VLAN must already exist on the switch.)

**Related  
Documentation**

- [Configuring 802.1X Interface Settings \(CLI Procedure\) on page 1908](#)
- [Configuring 802.1X Authentication \(J-Web Procedure\)](#)
- [Configuring MAC RADIUS Authentication \(CLI Procedure\) on page 1923](#)
- [Configuring Server Fail Fallback \(CLI Procedure\) on page 1921](#)

## Operational Commands

- `clear captive-portal`
- `clear dot1x`
- `clear lldp neighbors`
- `clear lldp statistics`
- `show captive-portal authentication-failed-users`
- `show captive-portal firewall`
- `show captive-portal interface`
- `show dot1x`
- `show dot1x authentication-failed-users`
- `show dot1x firewall`
- `show dot1x static-mac-address`
- `show lldp`
- `show lldp local-information`
- `show lldp neighbors`

- [show lldp remote-global-statistics](#)
- [show lldp statistics](#)
- [show network-access aaa statistics accounting](#)
- [show network-access aaa statistics authentication](#)
- [show network-access aaa statistics dynamic-requests](#)

## clear captive-portal

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>clear captive-portal</b> ( <b>firewall</b> [ <i>interface-names</i> ]   <b>interface</b> (802.1X) ( <b>all</b>   [ <i>interface-names</i> ])   <b>mac-address</b> [ <i>mac-addresses</i> ])                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Reset the authentication state of a captive portal interface or captive-portal firewall statistics on one or more interfaces.                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>firewall</b> [<i>interface-names</i>]<i>—</i>Resets captive portal statistics on all interfaces or on the specified interface.</p> <p><b>interface</b> (<b>all</b>   <i>interface-names</i>)<i>—</i>Resets the authentication state of users connected to all interfaces or the specified interfaces.</p> <p><b>mac-address</b> <i>mac-addresses</i><i>—</i>Resets the authentication state for the specified MAC addresses.</p>             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show captive-portal authentication-failed-users on page 2003</a></li> <li>• <a href="#">show captive-portal interface on page 2006</a></li> <li>• <a href="#">show captive-portal firewall on page 2004</a></li> <li>• <i>Example: Setting Up Captive Portal Authentication on an EX Series Switch</i></li> <li>• <i>Configuring Captive Portal Authentication (CLI Procedure)</i></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">clear captive-portal interface on page 1998</a><br><a href="#">clear captive-portal interface on page 1998</a><br><a href="#">clear captive-portal mac-address on page 1998</a><br><a href="#">clear captive-portal firewall on page 1998</a>                                                                                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 179 on page 1997</a> lists the output fields for the <b>clear captive-portal interface</b> command. (The <b>clear captive-portal firewall</b> and <b>clear captive-portal mac-address</b> commands have no output). Output fields are listed in the approximate order in which they appear.                                                                                                                                      |

**Table 179: clear captive-portal interface Output Fields**

| Field Name       | Field Description                                      |
|------------------|--------------------------------------------------------|
| <b>Interface</b> | Interface on which captive portal has been configured. |

Table 179: clear captive-portal interface Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>State</b>       | <p>The state of the port:</p> <ul style="list-style-type: none"> <li>• <b>Authenticated</b>—The client has been authenticated through the RADIUS server or has been permitted access through server fail fallback.</li> <li>• <b>Authenticating</b>—The client is authenticating through the RADIUS server.</li> <li>• <b>Connecting</b>—Switch is attempting to contact the RADIUS server.</li> <li>• <b>Initialize</b>—The interface link is down.</li> <li>• <b>Held</b>—An action has been triggered through server fail fallback during a RADIUS server timeout. A supplicant is denied access, permitted access through a specified VLAN, or maintains the authenticated state granted to it before the RADIUS server timeout occurred.</li> </ul> |
| <b>MAC address</b> | The MAC address of the connected client on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>User</b>        | Users connected to the captive portal interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## Sample Output

### clear captive-portal interface

```
user@switch> clear captive-portal interface
ge-0/0/3.0
```

### clear captive-portal interface

```
user@switch> clear captive-portal interface
Captive Portal Information:
Interface      State      MAC address      User
ge-0/0/3.0     Authenticated  00:03:47:e1:ba:b9  aclallow
ge-0/0/5.0     Connecting
ge-0/0/7.0     Connecting
ge-0/0/9.0     Connecting
```

### clear captive-portal mac-address

```
user@switch> clear captive-portal mac-address 00:03:47:e1:ba:b9
This command has no output.
```

### clear captive-portal firewall

```
user@switch> clear captive-portal firewall
This command has no output.
```

## clear dot1x

**Syntax** `clear dot1x (firewall <counter-name> | interface <[interface-name]> | mac-address [mac-addresses] | statistics <interface interface-name>)`

**Release Information** Command introduced in Junos OS Release 9.0 for EX Series switches.  
**firewall** option added in Junos OS Release 9.5 for EX Series switches.

**Description** Reset the authentication state of an interface or delete 802.1X statistics from the switch. When you reset an interface using the **interface** or **mac-address** options, reauthentication on the interface is also triggered. The switch sends out a multicast message on the interface to restart the authentication of all connected supplicants. If a MAC address is reset, then the switch sends out a unicast message to that specific MAC address to restart authentication.

If a supplicant is sending traffic when the **clear dot1x interface** command is issued, the authenticator immediately initiates reauthentication. This process happens quickly, and it might seem that reauthentication did not occur. To verify that reauthentication has happened, issue the **show dot1x interface detail** command. The values for **Reauthentication due** and **Reauthentication interval** will be about the same.



**CAUTION:** When you clear the learned MAC addresses from an interface using the **clear dot1x interface** command, all MAC addresses are cleared, including those in static MAC bypass list.

If you have enabled Media Access Control Security (MACsec) using static secure association key (SAK) security mode on an EX Series switch, the SAKs are rotated when the **clear dot1x** command is entered. The **clear dot1x** command has no impact on MACsec when MACsec is enabled using static connectivity association keys (CAK) or any other security mode.

**Options** **firewall <counter-name>**—Clear 802.1X firewall counter statistics. If the *counter-name* option is specified, clear 802.1X firewall statistics for that counter.

**interface <[interface-name]>**—Reset the authentication state of all the supplicants (also, clears all the authentication bypassed clients) connected to the specified interface (when the interface is an authenticator) or reset the authentication state for the interface itself (when the interface is a supplicant).

**mac-address [mac-addresses]**—Reset the authentication state of the specified MAC addresses.

**statistics <interface interface-name>**—Clear 802.1X statistics on all 802.1X-enabled interfaces. If the **interface** option is specified, clear 802.1X firewall statistics for that interface or interfaces.

**Required Privilege Level** view

**Related Documentation**

- [show dot1x on page 2009](#)
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 1910](#)

**List of Sample Output**

- [clear dot1x firewall c1 on page 2000](#)
- [clear dot1x interface ge-1/0/0 ge-2/0/0 ge-2/0/0 ge5/0/0 on page 2000](#)
- [clear dot1x mac-address 00:04:ae:cd:23:5f on page 2000](#)
- [clear dot1x statistics interface ge-1/0/1 on page 2000](#)

## Sample Output

[clear dot1x firewall c1](#)

```
user@switch> clear dot1x firewall c1
```

[clear dot1x interface ge-1/0/0 ge-2/0/0 ge-2/0/0 ge5/0/0](#)

```
user@switch> clear dot1x interface ge-1/0/0 ge-2/0/0 ge-2/0/0 ge5/0/0
```

[clear dot1x mac-address 00:04:ae:cd:23:5f](#)

```
user@switch> clear dot1x mac-address 00:04:ae:cd:23:5f
```

[clear dot1x statistics interface ge-1/0/1](#)

```
user@switch> clear dot1x statistics interface ge-1/0/1
```



## clear lldp neighbors

---

|                                 |                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear lldp neighbors<br><interface <i>interface</i> >                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                           |
| <b>Description</b>              | Clear the learned remote neighbor information on all or selected interfaces.                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>none</b>—Clear the remote neighbor information on all interfaces.</p> <p><b>interface <i>interface</i></b>—(Optional) Clear the remote neighbor information from one or more selected interfaces.</p>                                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show lldp on page 2018</a></li> <li>• <a href="#">Configuring LLDP (CLI Procedure) on page 1913</a></li> <li>• <a href="#">Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">clear lldp neighbors on page 2001</a><br><a href="#">clear lldp neighbors interface ge-0/1/1.0 on page 2001</a>                                                                                                                                                  |

### Sample Output

#### clear lldp neighbors

```
user@switch> clear lldp neighbors
```

#### clear lldp neighbors interface ge-0/1/1.0

```
user@switch> clear lldp neighbors interface ge-0/1/1.0
```

## clear lldp statistics

---

|                                 |                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear lldp statistics</code><br><code>&lt;interface <i>interface</i>&gt;</code>                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                      |
| <b>Description</b>              | Clear LLDP statistics on one or more interfaces.                                                                                                                                                                        |
| <b>Options</b>                  | <b>none</b> —Clears LLDP statistics on all interfaces.<br><br><b>interface <i>interface-names</i></b> —(Optional) Clear LLDP statistics on one or more interfaces.                                                      |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring LLDP (CLI Procedure) on page 1913</a></li><li>• <a href="#">Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">clear lldp statistics on page 2002</a><br><a href="#">clear lldp statistics interface ge-0/1/1.0 on page 2002</a>                                                                                           |

### Sample Output

#### clear lldp statistics

```
user@switch> clear lldp statistics
```

#### clear lldp statistics interface ge-0/1/1.0

```
user@switch> clear lldp statistics interface ge-0/1/1.0
```

## show captive-portal authentication-failed-users

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show captive-portal authentication-failed-users</b>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Display the users that have failed captive portal authentication.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show captive-portal interface on page 2006</a></li> <li>• <a href="#">show captive-portal firewall on page 2004</a></li> <li>• <a href="#">clear captive-portal on page 1997</a></li> <li>• <i>Example: Setting Up Captive Portal Authentication on an EX Series Switch</i></li> <li>• <i>Configuring Captive Portal Authentication (CLI Procedure)</i></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show captive-portal authentication-failed-users on page 2003</a>                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | <a href="#">Table 180 on page 2003</a> lists the output fields for the <b>show captive-portal authentication-failed-users</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                          |

Table 180: show captive-portal authentication-failed-users Output Fields

| Field Name           | Field Description                                                           | Level of Output |
|----------------------|-----------------------------------------------------------------------------|-----------------|
| <b>Interface</b>     | The MAC address configured to bypass captive portal authentication.         | all             |
| <b>MAC address</b>   | The MAC address configured statically on the interface.                     | all             |
| <b>User</b>          | Name of the user that has failed captive portal authentication.             | all             |
| <b>Failure Count</b> | The number of times that 802.1X authentication has failed on the interface. | all             |

### Sample Output

#### show captive-portal authentication-failed-users

```

user@switch> show captive-portal authentication-failed-users

Interface    MAC address    User           Failure Count
-----
ge-0/0/17.0  00:37:00:00:00:00  003700000000    28
ge-0/0/20.0  00:04:10:00:00:00  000410000000    32
ge-0/0/18.0  00:00:03:00:0a:00  000003000a00     4
ge-0/0/19.0  00:00:03:00:0b:00  000003000b00    18

```

## show captive-portal firewall

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show captive-portal firewall</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;interface-name&gt;</code><br><code>&lt;interface-name detail&gt;</code>                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display information about the firewall filters for each user that is authenticated on each captive portal interface.                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>none</b>—Display all the firewall filters on all captive portal interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface-name</b>—(Optional) Display all the terms of the firewall filters for the specified interface.</p> <p><b>interface-name detail</b>—(Optional) Display all of the terms of the firewall filters for the specified interface.</p>                 |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show captive-portal authentication-failed-users on page 2003</a></li><li>• <a href="#">show captive-portal interface on page 2006</a></li><li>• <a href="#">clear captive-portal on page 1997</a></li><li>• <i>Example: Setting Up Captive Portal Authentication on an EX Series Switch</i></li><li>• <i>Configuring Captive Portal Authentication (CLI Procedure)</i></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">show captive-portal firewall brief on page 2004</a><br><a href="#">show captive-portal firewall ge-0/0/10.0 on page 2005</a><br><a href="#">show captive-portal firewall on page 2005</a>                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | Output fields for the <b>show captive-portal firewall</b> command include any action modifier specified in firewall filters except policers. Policers are not supported in the terms of the internally generated dynamic firewall filters that are created when multiple supplicants authenticate on 802.1X-enabled interfaces.                                                                                                      |

## Sample Output

### show captive-portal firewall brief

```
user@switch> show captive-portal firewall brief
Captive Portal Information:
Interface      State      MAC address      User
ge-0/0/1.0     Connecting
ge-0/0/10.0    Connecting    00:30:48:8c:66:bd    No User
```

**show captive-portal firewall ge-0/0/10.0**

```

user@switch> show captive-portal firewall ge-0/0/10.0
Filter name: dot1x_ge-0/0/10
Counters:
Name                               Bytes      Packets
dot1x_ge-0/0/10_CP_arp             7616       119
dot1x_ge-0/0/10_CP_dhcp             0           0
dot1x_ge-0/0/10_CP_http             0           0
dot1x_ge-0/0/10_CP_https            0           0
dot1x_ge-0/0/10_CP_t_dns            0           0
dot1x_ge-0/0/10_CP_u_dns            0           0

```

**show captive-portal firewall**

```

user@switch> show captive-portal firewall
Filter name: dot1x_ge-0/0/0
Counters:
Name                               Bytes      Packets
dot1x_ge-0/0/0_CP_arp              0           0
dot1x_ge-0/0/0_CP_dhcp              0           0
dot1x_ge-0/0/0_CP_http              0           0
dot1x_ge-0/0/0_CP_https             0           0
dot1x_ge-0/0/0_CP_t_dns             0           0
dot1x_ge-0/0/0_CP_u_dns             0           0
Filter name: dot1x_ge-0/0/1
Counters:
Name                               Bytes      Packets
dot1x_ge-0/0/1_CP_arp              0           0
dot1x_ge-0/0/1_CP_dhcp              0           0
dot1x_ge-0/0/1_CP_http              0           0
dot1x_ge-0/0/1_CP_https             0           0
dot1x_ge-0/0/1_CP_t_dns             0           0
dot1x_ge-0/0/1_CP_u_dns             0           0
Filter name: dot1x_ge-0/0/10
Counters:
Name                               Bytes      Packets
dot1x_ge-0/0/10_CP_arp             7616       119
dot1x_ge-0/0/10_CP_dhcp             0           0
dot1x_ge-0/0/10_CP_http             0           0
dot1x_ge-0/0/10_CP_https            0           0
dot1x_ge-0/0/10_CP_t_dns            0           0
dot1x_ge-0/0/10_CP_u_dns            0           0
Filter name: dot1x_ge-0/0/11

```

## show captive-portal interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show captive-portal interface</b><br><b>&lt;interface-name&gt;</b><br><b>detail</b>                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Display the current operational state of all captive portal interfaces with the list of connected users and the configured values of captive portal attributes on the interfaces.                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Display all captive portal interfaces.</p> <p><b>interface-name</b>—(Optional) Display the state for the specified captive portal interface and lists the MAC address and user names of any clients authenticated on the interface.</p> <p><b>interface-name detail</b>—(Optional) Displays the configured values of captive portal attributes on the specified captive portal interface.</p>                                                                                          |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show captive-portal authentication-failed-users on page 2003</a></li> <li>• <a href="#">show captive-portal firewall on page 2004</a></li> <li>• <a href="#">captive-portal</a></li> <li>• <a href="#">clear captive-portal on page 1997</a></li> <li>• <a href="#">Example: Setting Up Captive Portal Authentication on an EX Series Switch</a></li> <li>• <a href="#">Configuring Captive Portal Authentication (CLI Procedure)</a></li> </ul> |
| <b>List of Sample Output</b>    | <p><a href="#">show captive-portal interface (Only captive portal is enabled) on page 2008</a></p> <p><a href="#">show captive-portal interface (Both 802.1X authentication and captive portal are enabled) on page 2008</a></p> <p><a href="#">show captive-portal interface detail (Only captive portal is enabled) on page 2008</a></p> <p><a href="#">show captive-portal interface detail (Both 802.1X authentication and captive portal are enabled) on page 2008</a></p>                       |
| <b>Output Fields</b>            | <a href="#">Table 181 on page 2006</a> lists the output fields for the <b>show captive-portal interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                          |

**Table 181: show captive-portal interface Output Fields**

| Field Name       | Field Description                                      | Level of Output |
|------------------|--------------------------------------------------------|-----------------|
| <b>Interface</b> | Interface on which captive portal has been configured. | All levels      |

Table 181: show captive-portal interface Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>State</b>                           | <p>The state of the interface:</p> <ul style="list-style-type: none"> <li>• <b>Authenticated</b>—The client has been authenticated through the RADIUS server or has been permitted access through server fail fallback.</li> <li>• <b>Authenticating</b>—The client is authenticating through the RADIUS server.</li> <li>• <b>Connecting</b>—Switch is attempting to contact the RADIUS server.</li> <li>• <b>Initialize</b>—The interface link is down.</li> <li>• <b>Held</b>—An action has been triggered through server fail fallback during a RADIUS server timeout. A supplicant is denied access, permitted access through a specified VLAN, or maintains the authenticated state granted to it before the RADIUS server timeout occurred.</li> </ul> | All levels      |
| <b>MAC address</b>                     | The MAC address of the connected client on the interface..                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | brief           |
| <b>User</b>                            | Users connected to the captive portal interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | brief           |
| <b>Fallen back</b>                     | <p>Indicates when 802.1X authentication and captive portal are both enabled on an interface:</p> <ul style="list-style-type: none"> <li>• If 802.1X authentication and captive portal are both enabled, <b>CP fallen back</b> status is <b>Yes</b>.</li> <li>• If 802.1X authentication and captive portal are not both enabled, <b>CP fallen back</b> status is <b>No</b>.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                        |                 |
| <b>Supplicant mode</b>                 | Mode used to authenticate clients—multiple, single, or single-supplicant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | detail          |
| <b>Number of retries</b>               | Number of times the user can attempt to submit authentication information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | detail          |
| <b>Quiet period</b>                    | Time, in seconds, after a user exceeds the maximum number of retries before they can attempt to authenticate.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | detail          |
| <b>Configured CP session timeout</b>   | Time, in seconds, that a client can be idle before the session expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | detail          |
| <b>Server timeout</b>                  | Time, in seconds, that an interface will wait for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | detail          |
| <b>Number of connected supplicants</b> | <p>Number of users connecting through the captive portal interface. Information for each user includes:</p> <ul style="list-style-type: none"> <li>• <b>Supplicant</b>—User name and MAC address.</li> <li>• <b>Operational state</b>—See State (above).</li> <li>• <b>Dynamic CP session timeout</b>—Timeout value dynamically downloaded from the RADIUS server for this user, if any.</li> <li>• <b>CP Session expiration due in</b>—Time remaining in session.</li> </ul>                                                                                                                                                                                                                                                                                 | detail          |

## Sample Output

### show captive-portal interface (Only captive portal is enabled)

```
user@switch> show captive-portal interface
Captive Portal Information:
Interface      State           MAC address      User             Fallen back
ge-0/0/1.0     Connecting      00:30:48:8c:66:bd No User
ge-0/0/10.0    Connecting      00:30:48:8c:66:bd No User
ge-6/0/5.0     Authenticated   00:30:48:8d:7a:9b abcdeX           No
```

### show captive-portal interface (Both 802.1X authentication and captive portal are enabled)

```
user@switch> show captive-portal interface
Captive Portal Information:
Interface      State           MAC address      User             Fallen back
ge-0/0/1.0     Connecting      00:30:48:8c:66:bd No User
ge-0/0/10.0    Connecting      00:30:48:8c:66:bd No User
ge-6/0/5.0     Authenticated   00:30:48:8d:7a:9b abcdeX           Yes
```

### show captive-portal interface detail (Only captive portal is enabled)

```
user@switch> show captive-portal interface detail ge-6/0/5.0
Supplicant mode: Multiple
Number of retries: 3
Quiet period: 60 seconds
Configured CP session timeout: 3600 seconds
Server timeout: 15 seconds
CP fallen back: No
Number of connected supplicants: 1
  Supplicant: abcdeX, 00:30:48:8d:7a:9b
    Operational state: Authenticated
    Dynamic CP Session Timeout: 3600 seconds
    CP Session Expiration due in: 3583 seconds
```

### show captive-portal interface detail (Both 802.1X authentication and captive portal are enabled)

```
user@switch> show captive-portal interface detail ge-6/0/5.0
Supplicant mode: Multiple
Number of retries: 3
Quiet period: 60 seconds
Configured CP session timeout: 3600 seconds
Server timeout: 15 seconds
CP fallen back: Yes
Number of connected supplicants: 1
  Supplicant: abcdeX, 00:30:48:8d:7a:9b
    Operational state: Authenticated
    Dynamic CP Session Timeout: 3600 seconds
    CP Session Expiration due in: 3583 seconds
```



## show dot1x

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show dot1x</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;interface <i>interface-name</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Display the current operational state of all ports with the list of connected users.</p> <p>This command displays the list of connected supplicants received from the RADIUS authentication server regardless of the session state—that is, for both authenticated supplicants and for supplicants that attempted authentication.</p>                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>none</b>—Display information for all authenticator ports.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—Display information for the specified port with a list of connected supplicants.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear dot1x on page 1999</a></li> <li>• <a href="#">Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852</a></li> <li>• <a href="#">Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to an EX Series Switch on page 1875</a></li> <li>• <a href="#">Example: Configuring Fallback Options on EX Series Switches for EAP-TTLS Authentication and Odyssey Access Clients on page 1903</a></li> <li>• <a href="#">Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 1910</a></li> <li>• <a href="#">Verifying 802.1X Authentication on page 1994</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show dot1x interface brief on page 2012</a><br><a href="#">show dot1x interface detail on page 2012</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | <a href="#">Table 182 on page 2009</a> lists the output fields for the <b>show dot1x</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

Table 182: show dot1x Output Fields

| Field Name  | Field Description                                        | Level of Output |
|-------------|----------------------------------------------------------|-----------------|
| Interface   | Name of a port.                                          | All levels      |
| MAC address | The MAC address of the connected supplicant on the port. | All levels      |

Table 182: show dot1x Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output      |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| <b>Role</b>                 | The 802.1X authentication role of the interface. When 802.1X is enabled on an interface, the role is <b>Authenticator</b> . As <b>Authenticator</b> , the interface blocks LAN access until a supplicant is authenticated through 802.1X or MAC RADIUS authentication.                                                                                                                                                                                                                                                                                                                                                                | <b>brief, detail</b> |
| <b>State</b>                | <p>The state of the port:</p> <ul style="list-style-type: none"> <li>• <b>Authenticated</b>—The supplicant has been authenticated through the RADIUS server or has been permitted access through server fail fallback.</li> <li>• <b>Authenticating</b>—The supplicant is authenticating through the RADIUS server.</li> <li>• <b>Held</b>—An action has been triggered through server fail fallback during a RADIUS server timeout. A supplicant is denied access, permitted access through a specified VLAN, or maintains the authenticated state granted to it before the RADIUS server timeout occurred.</li> </ul>               | <b>brief</b>         |
| <b>User</b>                 | The user name of the connected supplicant                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>brief</b>         |
| <b>Administrative state</b> | <p>The administrative state of the port:</p> <ul style="list-style-type: none"> <li>• <b>auto</b>—Traffic is allowed through the port based on the authentication result. (Default)</li> <li>• <b>force-authorize</b>—All traffic flows through the port irrespective of the authentication result. This state is not allowed on an interface whose VLAN membership has been set to <b>dynamic</b>.</li> <li>• <b>force-unauthorize</b>—All traffic drops on the port irrespective of the authentication result. This state is not allowed on an interface whose VLAN membership has been set to <b>dynamic</b>.</li> </ul>           | <b>detail</b>        |
| <b>Supplicant</b>           | <p>The mode for the supplicant:</p> <ul style="list-style-type: none"> <li>• <b>single</b>—Authenticates only the first supplicant. All other supplicants who connect later to the port are allowed full access without any further authentication. They effectively “piggyback” on the first supplicant’s authentication.</li> <li>• <b>single-secure</b>—Allows only one supplicant to connect to the port. No other supplicant is allowed to connect until the first supplicant logs out.</li> <li>• <b>multiple</b>—Allows multiple supplicants to connect to the port. Each supplicant is authenticated individually.</li> </ul> | <b>detail</b>        |
| <b>Quiet period</b>         | The number of seconds the port remains in the wait state following a failed authentication exchange with the supplicant before reattempting the authentication. The default value is 60 seconds. The range is 0 through 65,535 seconds.                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail</b>        |
| <b>Transmit period</b>      | The number of seconds the port waits before retransmitting the initial EAPOL PDUs to the supplicant. The default value is 30 seconds. The range is 1 through 65,535 seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail</b>        |

Table 182: show dot1x Output Fields (*continued*)

| Field Name                                                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output |
|-------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>MAC radius</b>                                           | MAC RADIUS authentication: <ul style="list-style-type: none"> <li>• <b>enabled</b>—The switch sends an EAPOL request to the connecting host to attempt 802.1X authentication and if the connecting host is unresponsive, the switch tries to authenticate using the MAC address.</li> <li>• <b>disabled</b>—The default. The switch will not attempt to authenticate the MAC address of the connecting host.</li> </ul> | <b>detail</b>   |
| <b>MAC radius restrict</b>                                  | The authentication method is restricted to MAC RADIUS only. 802.1X authentication is not enabled.                                                                                                                                                                                                                                                                                                                       | <b>detail</b>   |
| <b>Reauthentication</b>                                     | The reauthentication state: <ul style="list-style-type: none"> <li>• <b>disable</b>—Periodic reauthentication of the client is disabled.</li> <li>• <b>interval</b>—Sets the periodic reauthentication time interval. The default value is 3600 seconds. The range is 1 through 65,535 seconds.</li> </ul>                                                                                                              | <b>detail</b>   |
| <b>Supplicant timeout</b>                                   | The number of seconds the port waits for a response when relaying a request from the authentication server to the supplicant before resending the request. The default value is 30 seconds. The range is 1 through 60 seconds.                                                                                                                                                                                          | <b>detail</b>   |
| <b>Server timeout</b>                                       | The number of seconds the port waits for a reply when relaying a response from the supplicant to the authentication server before timing out. The default value is 30 seconds. The range is 1 through 60 seconds.                                                                                                                                                                                                       | <b>detail</b>   |
| <b>Maximum EAPOL requests</b>                               | The maximum number of retransmission times of an EAPOL request packet to the supplicant before the authentication session times out. The default value is 2. The range is 1 through 10.                                                                                                                                                                                                                                 | <b>detail</b>   |
| <b>Number of clients bypassed because of authentication</b> | The number of non-802.1X clients granted access to the LAN by means of static MAC bypass. The following fields are displayed: <ul style="list-style-type: none"> <li>• <b>Client</b>—MAC address of the client.</li> <li>• <b>vlan</b> —The name of the VLAN to which the client is connected.</li> </ul>                                                                                                               | <b>detail</b>   |
| <b>Guest VLAN member</b>                                    | The VLAN to which a supplicant is connected when the supplicant is authenticated using a guest VLAN. If a guest VLAN is not configured on the interface, this field displays <b>&lt;not configured&gt;</b> .                                                                                                                                                                                                            | <b>detail</b>   |
| <b>Number of connected supplicants</b>                      | The number of supplicants connected to a port.                                                                                                                                                                                                                                                                                                                                                                          | <b>detail</b>   |
| <b>Supplicant</b>                                           | The user name and MAC address of the connected supplicant.                                                                                                                                                                                                                                                                                                                                                              | <b>detail</b>   |

Table 182: show dot1x Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Authentication method   | <p>The 802.1X authentication method used for a supplicant:</p> <ul style="list-style-type: none"> <li><b>Guest VLAN</b>—A supplicant is connected to the LAN through the guest VLAN.</li> <li><b>MAC Radius</b>—A nonresponsive host is authenticated based on its MAC address. The MAC address is configured as permitted on the RADIUS server, the RADIUS server lets the switch know that the MAC address is a permitted address, and the switch opens LAN access to the nonresponsive host on the interface to which it is connected.</li> <li><b>Radius</b>—A supplicant is configured on the RADIUS server, the RADIUS server communicates this to the switch, and the switch opens LAN access on the interface to which the supplicant is connected.</li> <li><b>Server-fail deny</b>—If the RADIUS servers time out, all supplicants are denied access to the LAN, preventing traffic from flowing from the supplicant through the interface. This is the default.</li> <li><b>Server-fail permit</b>—When the RADIUS server is unavailable, a supplicant is still permitted access to the LAN as if the supplicant had been successfully authenticated by the RADIUS server.</li> <li><b>Server-fail use-cache</b>—If the RADIUS servers time out during reauthentication, previously authenticated supplicants are reauthenticated, but new supplicants are denied LAN access.</li> <li><b>Server-fail VLAN</b>—A supplicant is configured to be moved to a specified VLAN if the RADIUS server is unavailable to reauthenticate the supplicant. (The VLAN must already exist on the switch.)</li> </ul> | detail          |
| Authenticated VLAN      | The VLAN to which the supplicant is connected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | detail          |
| Dynamic filter          | User policy filter sent by the RADIUS server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | detail          |
| Session Reauth interval | The configured reauthentication interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | detail          |
| Reauthentication due in | The number of seconds in which reauthentication will occur again for the connected supplicant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | detail          |

## Sample Output

### show dot1x interface brief

```

user@switch> show dot1x interface brief
802.1X Information:
Interface    Role           State           MAC address    User
ge-0/0/1     Authenticator  Authenticated   00:a0:d2:18:1a:c8  user1
ge-0/0/2     Authenticator  Connecting      00:a6:55:f2:94:ae  user3
ge-0/0/3     Authenticator  Held

```

### show dot1x interface detail

```

user@switch> show dot1x interface ge-0/0/16.0 detail

ge-0/0/16.0
Role: Authenticator
Administrative state: Auto

```

Supplicant mode: Single  
Number of retries: 3  
Quiet period: 60 seconds  
Transmit period: 30 seconds  
Mac Radius: Enabled  
Mac Radius Strict: Disabled  
Reauthentication: Enabled  
Configured Reauthentication interval: 40 seconds  
Supplicant timeout: 30 seconds  
Server timeout: 30 seconds  
Maximum EAPOL requests: 1  
Guest VLAN member: <not configured>  
Number of connected supplicants: 1  
    Supplicant: abc, 00:30:48:8C:66:BD  
        Operational state: Authenticated  
        Authentication method: Radius  
        Authenticated VLAN: v200  
        Reauthentication due in 17 seconds

## show dot1x authentication-failed-users

|                                 |                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show dot1x authentication-failed-users                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Displays supplicants (users) that have failed 802.1X authentication.                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear dot1x on page 1999</a></li> <li>• <a href="#">Example: Configuring Static MAC Bypass of Authentication on an EX Series Switch on page 1872</a></li> <li>• <a href="#">Configuring 802.1X Interface Settings (CLI Procedure) on page 1908</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show dot1x authentication-failed-users on page 2014</a>                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | <a href="#">Table 183 on page 2014</a> lists the output fields for the <b>show dot1x authentication-failed-users</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                              |

**Table 183: show dot1x authentication-failed-users Output Fields**

| Field Name           | Field Description                                                                           | Level of Output |
|----------------------|---------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>     | The MAC address configured to bypass 802.1X authentication.                                 | all             |
| <b>MAC address</b>   | The MAC address configured statically on the interface.                                     | all             |
| <b>User</b>          | The user that is configured on the RADIUS server and that has failed 802.1X authentication. | all             |
| <b>Failure Count</b> | The number of times that 802.1X authentication has failed on the interface.                 | all             |

## Sample Output

### show dot1x authentication-failed-users

```
user@switch> show dot1x authentication-failed-users
```

| Interface   | MAC address       | User         | Failure Count |
|-------------|-------------------|--------------|---------------|
| ge-0/0/17.0 | 00:37:00:00:00:00 | 003700000000 | 28            |
| ge-0/0/20.0 | 00:04:10:00:00:00 | 000410000000 | 32            |
| ge-0/0/18.0 | 00:00:03:00:0a:00 | 000003000a00 | 4             |
| ge-0/0/19.0 | 00:00:03:00:0b:00 | 000003000b00 | 18            |

## show dot1x firewall

|                                 |                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show dot1x firewall &lt;interface <i>interface-name</i>&gt;</code>                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                                                             |
| <b>Description</b>              | Displays information about the firewall filters for each user or nonresponsive host that is authenticated on each 802.1X-enabled interface that is configured for multiple supplicants. For example, if the firewall filter is configured with a term for counters, the command shows the count for each user. |
| <b>Options</b>                  | <code>interface <i>interface-names</i></code> —(Optional) Display information for the specified interface.                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear dot1x on page 1999</a></li> <li>• <a href="#">Example: Applying Firewall Filters to Multiple Supplicants on Interfaces Enabled for 802.1X or MAC RADIUS Authentication on page 4798</a></li> </ul>                                                  |
| <b>List of Sample Output</b>    | <a href="#">show dot1x firewall on page 2015</a><br><a href="#">show dot1x firewall on page 2015</a>                                                                                                                                                                                                           |
| <b>Output Fields</b>            | Output fields include any action modifier that is specified in firewall filters.                                                                                                                                                                                                                               |

### Sample Output

#### show dot1x firewall

(Showing counter action)

```
user@switch> show dot1x firewall
Filter: dot1x-filter-ge-0/0/3
Counters
  counter1_dot1x_ge-0/0/3_user1    342
  counter1_dot1x_ge-0/0/3_user2    857
```

#### show dot1x firewall

(Showing policer action)

```
user@switch> show dot1x firewall
Filter: dot1x_ge-0/0/0
Counters
  p1-t1    494946
```

## show dot1x static-mac-address

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show dot1x static-mac-address &lt;(interface [interface-name])&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Displays all the static MAC addresses that are configured to bypass 802.1X authentication on the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <code>interface [ interface-name ]</code> —(Optional) Display static MAC addresses for a specific interface.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear dot1x on page 1999</a></li> <li>• <a href="#">Example: Configuring Static MAC Bypass of Authentication on an EX Series Switch on page 1872</a></li> <li>• <a href="#">Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure) on page 2361</a></li> <li>• <a href="#">Configuring 802.1X Interface Settings (CLI Procedure) on page 1908</a></li> <li>• <a href="#">Understanding Authentication on EX Series Switches on page 1824</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show dot1x static-mac-address on page 2016</a><br><a href="#">show dot1x static-mac-address interface ge-0/0/0.1 on page 2017</a>                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>            | <a href="#">Table 184 on page 2016</a> lists the output fields for the <code>show dot1x static-mac-address</code> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                      |

**Table 184: show dot1x static-mac-address Output Fields**

| Field Name             | Field Description                                                                      | Level of Output |
|------------------------|----------------------------------------------------------------------------------------|-----------------|
| <b>MAC address</b>     | The MAC address of the device that is configured to bypass 802.1X authentication.      | all             |
| <b>VLAN-Assignment</b> | The name of the VLAN to which the device is assigned.                                  | all             |
| <b>Interface</b>       | The name of the interface on which authentication is bypassed for a given MAC address. | all             |

## Sample Output

### show dot1x static-mac-address

```
user@switch> show dot1x static-mac-address
```

```
MAC address      VLAN-Assignment  Interface
00:00:00:11:22:33
```



|                   |            |            |
|-------------------|------------|------------|
| 00:00:00:00:12:12 |            | ge-0/0/3.0 |
| 00:00:00:02:34:56 | facilities | ge-0/0/1.0 |

#### `show dot1x static-mac-address interface ge-0/0/0.1`

```
user@switch> show dot1x static-mac-address interface ge-0/0/0.1
```

| MAC address       | VLAN-Assignment | Interface  |
|-------------------|-----------------|------------|
| 00:00:00:12:24:12 | support         | ge-0/0/1.0 |
| 00:00:00:72:30:58 | support         | ge-0/0/1.0 |

## show lldp

**Syntax** `show lldp`  
`<detail>`

**Release Information** Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display information about Link Layer Discovery Protocol (LLDP) and Link Level Discovery Protocol—Media Endpoint Discovery (LLDP-MED) configuration and capabilities on the switch. LLDP and LLDP-MED are used to learn about and to distribute device information on network links.



**NOTE:** LLDP-MED is not available on the QFX Series.

**Options** **none**—Display LLDP information for all interfaces.  
**detail**—(Optional) Display detailed LLDP information for all interfaces.

**Required Privilege Level** view

**Related Documentation**

- [Configuring LLDP \(CLI Procedure\) on page 1913](#)
- [Configuring LLDP-MED \(CLI Procedure\) on page 1917](#)
- [Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831](#)
- [Configuring LLDP](#)
- [Understanding LLDP](#)

**List of Sample Output** [show lldp \(EX3200 switches\) on page 2021](#)  
[show lldp \(EX4300 switches\) on page 2021](#)  
[show lldp detail \(EX4300 switches\) on page 2022](#)

**Output Fields** [Table 185 on page 1818](#) lists the output fields for the **show lldp** command. Output fields are listed in the approximate order in which they appear.

**Table 185: show lldp Output Fields**

| Field Name | Field Description                                                                                                                                                                                                                                                                                                 | Level of Output |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| LLDP       | LLDP operating state. The state can be <b>enabled</b> or <b>disabled</b> .<br><br><b>NOTE:</b> If a VLAN that has been configured for untagged packets on an interface also has Layer 2 protocol tunneling (L2PT) enabled for LLDP, the LLDP operating state for that interface is displayed as <b>disabled</b> . | All levels      |

Table 185: show lldp Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Advertisement interval</b> | Frequency, in seconds, at which LLDP advertisements are sent.<br><br>This value is set by the <i>advertisement-interval</i> configuration statement.                                                                                                                                                                                                                                                                                                             | All levels      |
| <b>Transmit delay</b>         | Seconds of delay before advertisements are sent to neighbors following a change to a TLV (type, length, or value) element in the LLDP protocol or to the state of the local system, such as a change in hostname or management address. You can set this value to reduce the delay in notifying neighbors of a change in the local system.<br><br>This value is set by the <i>transmit-delay</i> configuration statement.                                        | All levels      |
| <b>Hold timer</b>             | On EX4300 switches, the hold timer shows the length of time LLDP information is held before it is discarded. The hold timer value is equal to the advertisement interval multiplied by the hold multiplier.<br><br>On all other switches, the hold timer shows the value of the hold multiplier.<br><br>The hold multiplier value is set by the <i>hold-multiplier</i> configuration statement.                                                                  | All levels      |
| <b>Notification interval</b>  | How often LLDP trap notifications are generated as a result of LLDP database changes. If the interval value is 0, LLDP trap notifications on database changes are disabled.<br><br>This value is set by the <i>lldp-configuration-notification-interval</i> configuration statement.                                                                                                                                                                             | All levels      |
| <b>Config Trap Interval</b>   | How often LLDP trap notifications are generated as a result of changes in topology—for example, when an endpoint connects or disconnects. If the interval value is 0, LLDP trap notifications on topology changes are disabled.<br><br>This value is set by the <i>ptopo-configuration-trap-interval</i> configuration statement.                                                                                                                                | All levels      |
| <b>Connection Hold timer</b>  | Amount of time the system maintains dynamic topology entries.<br><br>This value is set by the <i>ptopo-configuration-maximum-hold-time</i> configuration statement.                                                                                                                                                                                                                                                                                              | All levels      |
| <b>LLDP-MED</b>               | LLDP-MED operating state. The state can be <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                   | All levels      |
| <b>MED fast start count</b>   | Number of advertisements sent from a switch to a device, such as a VoIP telephone, when the device is first detected by the switch. These increased advertisements are temporary. After a device and a switch exchange information and can communicate, advertisements are reduced to one per second.<br><br>This value is set by using the <i>fast-start</i> configuration statement.<br><br><b>NOTE:</b> <i>fast-start</i> is not available on the QFX Series. | All levels      |
| <b>Interface</b>              | Name of the interface for which LLDP configuration information is being reported.                                                                                                                                                                                                                                                                                                                                                                                | All levels      |
| <b>Parent Interface</b>       | Name of the aggregated Ethernet interface, if any, to which the interface belongs.                                                                                                                                                                                                                                                                                                                                                                               | All levels      |

Table 185: show lldp Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>LLDP</b>                      | LLDP operating state. The state can be <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels      |
| <b>Power Negotiation</b>         | LLDP power negotiation operating state. The state can be <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels      |
| <b>Neighbor count</b>            | Total number of new LLDP neighbors detected since the last switch reboot.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail</b>   |
| <b>Interface</b>                 | Name of the interface that is advertising VLAN information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>Vlan-id</b>                   | VLAN tag associated with the interface sending LLDP frames. If the interface is not a member of a VLAN, the VLAN ID is advertised as 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>   |
| <b>Vlan-name</b>                 | VLAN name associated with the VLAN ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail</b>   |
| <b>LLDP basic TLVs supported</b> | <p>Basic TLVs supported on the switch:</p> <ul style="list-style-type: none"> <li>• <b>Chassis identifier</b>—TLV that advertises the MAC address associated with the local system.</li> <li>• <b>Port identifier</b>—TLV that advertises the port identification for the specified port in the local system.</li> <li>• <b>Port description</b>—Interface name for the port.</li> <li>• <b>System name</b>—TLV that advertises the user-configured name of the local system.</li> <li>• <b>System description</b>—TLV that advertises the system description containing information about the software and current image running on the system. This information is taken from the software and is not configurable.</li> <li>• <b>System capabilities</b>—TLV that advertises the primary functions performed by the system—for example, bridge or router.</li> <li>• <b>Management address</b>—TLV that advertises the IP management address of the local system.</li> </ul> | <b>detail</b>   |
| <b>Supported LLDP 802 TLVs</b>   | <p>802.3 TLVs supported on the switch:</p> <ul style="list-style-type: none"> <li>• <b>MAC/PHY configuration status</b>—TLV that advertises information about the physical interface, such as autonegotiation status and support and MAU type. The information is based on the physical interface structure and is not configurable.</li> <li>• <b>Power via MDI</b>—TLV that advertises MDI power support, PSE power pair, and power class information.</li> <li>• <b>Link aggregation</b>—TLV that advertises if the interface is aggregated and its aggregated interface ID.</li> <li>• <b>Maximum frame size</b>—TLV that advertises the maximum transmission unit (MTU) of the interface sending LLDP frames.</li> <li>• <b>Port VLAN tag</b>—TLV that advertises the VLAN tag configured on the interface.</li> <li>• <b>Port VLAN name</b>—TLV that advertises the VLAN name configured on the interface.</li> </ul>                                                     | <b>detail</b>   |

Table 185: show lldp Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Supported LLDP MED TLVs | <p>LLDP-MED TLVs supported on the switch:</p> <ul style="list-style-type: none"> <li>• <b>LLDP MED capabilities</b>—TLV that advertises the primary function of the port. The capabilities values range from 0 through 15: <ul style="list-style-type: none"> <li>• 0—Capabilities</li> <li>• 1—Network Policy</li> <li>• 2—Location Identification</li> <li>• 3—Extended Power via MDI-PSE</li> <li>• 4—Inventory</li> <li>• 5–15—Reserved</li> </ul> </li> <li>• <b>Network policy</b>—TLV that advertises the port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier, application types—such as voice or streaming video—802.1Q VLAN tagging, and 802.1p priority bits and DiffServ code points.</li> <li>• <b>Endpoint location</b>—TLV that advertises the physical location of the endpoint.</li> <li>• <b>Extended power Via MDI</b>—TLV that advertises the power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port.</li> </ul> | detail          |

## Sample Output

### show lldp (EX3200 switches)

```

user@switch> show lldp
LLDP                               : Enabled
Advertisement interval             : 30 seconds
Transmit delay                     : 2 seconds
Hold timer                        : 4 seconds
Notification interval             : 0 Second(s)
Config Trap Interval              : 0 seconds
Connection Hold timer             : 300 seconds

LLDP MED                           : Disabled
MED fast start count              : 3 Packets

```

|           |                  |         |          |                   |
|-----------|------------------|---------|----------|-------------------|
| Interface | Parent Interface | LLDP    | LLDP-MED | Power Negotiation |
| all       | -                | Enabled | Enabled  | Enabled           |

### show lldp (EX4300 switches)

```

user@switch> show lldp
LLDP                               : Enabled
Advertisement interval             : 30 seconds
Transmit delay                     : 2 seconds
Hold timer                        : 120 seconds
Notification interval             : 0 Second(s)
Config Trap Interval              : 0 seconds
Connection Hold timer             : 300 seconds

LLDP MED                           : Disabled
MED fast start count              : 3 Packets

```

| Interface | Parent Interface | LLDP    | LLDP-MED | Power Negotiation |
|-----------|------------------|---------|----------|-------------------|
| all       | -                | Enabled | Enabled  | Enabled           |

**show lldp detail (EX4300 switches)**

```

user@switch> show lldp detail
LLDP : Enabled
Advertisement interval : 30 seconds
Transmit delay : 2 seconds
Hold timer : 120 seconds
Notification interval : 0 Second(s)
Config Trap Interval : 0 seconds
Connection Hold timer : 300 seconds

LLDP MED : Disabled
MED fast start count : 3 Packets

```

| Interface      | Parent Interface | LLDP    | LLDP-MED | Power Negotiation |
|----------------|------------------|---------|----------|-------------------|
| Neighbor count |                  |         |          |                   |
| all            | -                | Enabled | Enabled  | Enabled           |
| 8              |                  |         |          |                   |

| Interface  | Parent Interface | Vlan-id | Vlan-name |
|------------|------------------|---------|-----------|
| xe-3/0/0.0 | ae31.0           | 100     | v100      |
| xe-3/0/0.0 | ae31.0           | 101     | v101      |
| xe-3/0/0.0 | ae31.0           | 4000    | v4000     |
| xe-3/0/1.0 | ae31.0           | 100     | v100      |
| xe-3/0/1.0 | ae31.0           | 101     | v101      |
| xe-3/0/1.0 | ae31.0           | 4000    | v4000     |
| xe-3/0/2.0 | ae31.0           | 100     | v100      |
| xe-3/0/2.0 | ae31.0           | 101     | v101      |
| xe-3/0/2.0 | ae31.0           | 4000    | v4000     |

**LLDP basic TLVs supported:**

Chassis identifier, Port identifier, Port description, System name, System description, System capabilities, Management address.

**Supported LLDP 802 TLVs:**

MAC/PHY configuration/status, Power via MDI, Link aggregation, Maximum frame size, Port VLAN tag, Port VLAN name.

**Supported LLDP MED TLVs:**

LLDP MED capabilities, Network policy, Endpoint location, Extended power Via MDI.

## show lldp local-information

|                                 |                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show lldp local-information                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                   |
| <b>Description</b>              | Display the information that the switch provides in Link Layer Discovery Protocol (LLDP) advertisements to its neighbors.                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP (CLI Procedure) on page 1913</a></li> <li>• <a href="#">Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831</a></li> <li>• <i>management-address</i></li> <li>• <i>Configuring LLDP</i></li> <li>• <i>Understanding LLDP</i></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show lldp local-information (EX Series Switch) on page 2024</a>                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 186 on page 2023</a> lists the output fields for the <b>show lldp local-information</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                              |

**Table 186: show lldp local-information Output Fields**

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>LLDP Local Information details</b> | <p>Information about the local system (the switch):</p> <ul style="list-style-type: none"> <li>• <b>Chassis ID</b>—MAC address associated with the switch.</li> <li>• <b>System name</b>—User-configured name of the switch.</li> <li>• <b>System descr</b>—System description containing information about the switch model and the current software image running on the switch. This information is taken from the software and is not configurable.</li> </ul>                                                                                                                                                                                                                 |
| <b>System Capabilities</b>            | Capabilities (such as <b>bridge</b> or <b>router</b> ) that are supported or enabled on the system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Management Information</b>         | <p>Details of the management information: <b>Port Name</b>, <b>Port Address</b> (such as 10.204.34.35), <b>Address Type</b> (such as ipv4 or ipv6), <b>Port ID</b> (SNMP interface index), <b>Port ID Subtype</b>, and <b>Port Subtype</b>.</p> <p>The <b>Port Subtype</b> displays:</p> <ul style="list-style-type: none"> <li>• <b>ifindex(2)</b>—IP address of the switch's management Ethernet interface (<b>me0</b>) or virtual management Ethernet (<b>VME</b>) interface address (for a virtual chassis) is used to manage the switch.</li> <li>• <b>unknown(1)</b>—IP management address has been configured with set <b>protocols lldp management-address</b>.</li> </ul> |

Table 186: show lldp local-information Output Fields (*continued*)

| Field Name                   | Field Description                                                                        |
|------------------------------|------------------------------------------------------------------------------------------|
| <b>Interface name</b>        | Name of the local interface which is configured for either LLDP or LLDP-MED.             |
| <b>Parent Interface</b>      | Name of the aggregated Ethernet interface, if any, to which the local interface belongs. |
| <b>SNMP Index</b>            | SNMP interface index.                                                                    |
| <b>Interface description</b> | User-configured port description.                                                        |
| <b>Status</b>                | Administrative status of the interface: either <b>up</b> or <b>down</b> .                |
| <b>Tunneling</b>             | Status of tunneling on the interface: either <b>enabled</b> or <b>disabled</b> .         |

## Sample Output

### show lldp local-information (EX Series Switch)

```
user@switch> show lldp local-information
```

#### LLDP Local Information details

```
Chassis ID   : 00:1d:b5:aa:b9:f0
System name  : switch
System descr : Juniper Networks, Inc. ex8208 , version 10.4I0 [builder] Build
               date: 2010-11-17 12:38:30 UTC
```

#### System Capabilities

```
Supported   : Bridge Router
Enabled     : Bridge Router
```

#### Management Information

```
Port Name    : -
Port Address  : 10.93.54.6
Address Type  : IPv4
Port ID       : 34
Port ID Subtype : local(7)
Port Subtype  : ifIndex(2)
```

| Interface name | Parent Interface | SNMP Index | Interface description | Status | Tunneling |
|----------------|------------------|------------|-----------------------|--------|-----------|
| me0.0          | -                | 34         | -                     | Down   | Disabled  |
| xe-3/0/0.0     | ae31.0           | 769        | xe-3/0/0.0            | Up     | Disabled  |
| xe-3/0/1.0     | ae31.0           | 770        | xe-3/0/1.0            | Up     | Disabled  |
| xe-3/0/2.0     | ae31.0           | 771        | xe-3/0/2.0            | Up     | Disabled  |
| xe-3/0/3.0     | ae31.0           | 772        | xe-3/0/3.0            | Up     | Disabled  |
| xe-3/0/4.0     | ae31.0           | 577        | xe-3/0/4.0            | Up     | Disabled  |
| xe-3/0/5.0     | ae31.0           | 578        | xe-3/0/5.0            | Up     | Disabled  |
| xe-3/0/6.0     | ae31.0           | 579        | xe-3/0/6.0            | Up     | Disabled  |
| xe-3/0/7.0     | ae31.0           | 581        | xe-3/0/7.0            | Up     | Disabled  |



## show lldp neighbors

**Syntax** `show lldp neighbors`  
`<interface interface>`

**Release Information** Command introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Display the information about neighboring devices learned by the switch by using the Link Layer Discovery Protocol (LLDP).



**NOTE:** The Chassis ID TLV has a subtype for Network Address Family. The supported network address families are IPv4 and IPv6. LLDP frames are validated only if the Network Address subtype of the Chassis ID TLV has a value of 1 (IPv4) or 2 (IPv6). For any other value, the transmitting device is detected by LLDP as a neighbor and displayed in the output of the `show lldp neighbors` command, but is not assigned to the VLAN.

**Options** `interface interface`—(Optional) Display LLDP neighbor information for a selected interface.

**Required Privilege Level** view

**Related Documentation**

- [Configuring LLDP \(CLI Procedure\) on page 1913](#)
- [Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831](#)

**List of Sample Output**

[show lldp neighbors on page 2027](#)  
[show lldp neighbors interface ge-0/0/2 on page 2028](#)  
[show lldp neighbors interface ge-0/0/0.0 \(for a VoIP Avaya Telephone with LLDP-MED Support\) on page 2029](#)  
[show lldp neighbors interface ge-0/0/5.0 \(with NetBIOS Snooping Enabled on the Switch\) on page 2030](#)

**Output Fields** [Table 187 on page 2025](#) lists the output fields for the `show lldp neighbors` command. Output fields are listed in the approximate order in which they appear.

**Table 187: show lldp neighbors Output Fields**

| Field Name       | Field Description                                                                     |
|------------------|---------------------------------------------------------------------------------------|
| Local Interface  | List of local interfaces for which neighbor information is available.                 |
| Parent Interface | List of aggregated Ethernet interfaces, if any, to which the local interfaces belong. |
| Chassis ID       | List of chassis identifiers for neighbors.                                            |
| Port info        | This field displays the port information received from neighbors.                     |

Table 187: show lldp neighbors Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                 |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>System name</b>               | List of system names gathered from neighbors.                                                                                                                                                                                                                     |
| <b>LLDP Neighbor Information</b> | Information about both the local system (the switch) and a neighbor system on the interface (appears when the <b>interface</b> option is used).                                                                                                                   |
| <b>Local Information</b>         | Information about the local system (appears when the <b>interface</b> option is used).                                                                                                                                                                            |
| <b>Index</b>                     | Local interface index (appears when the <b>interface</b> option is used).                                                                                                                                                                                         |
| <b>Time to live</b>              | Number of seconds for which this information is valid (appears when the <b>interface</b> option is used).                                                                                                                                                         |
| <b>Time mark</b>                 | Date and timestamp of information (appears when the <b>interface</b> option is used).                                                                                                                                                                             |
| <b>Local Interface</b>           | Name of the local physical interface (appears when the <b>interface</b> option is used).                                                                                                                                                                          |
| <b>Parent Interface</b>          | Name of the aggregated Ethernet interface, if any, to which the interface belongs (appears when the <b>interface</b> option is used).                                                                                                                             |
| <b>Local Port ID</b>             | Local interface SNMP index (appears when the <b>interface</b> option is used).                                                                                                                                                                                    |
| <b>Ageout Count</b>              | Number of times the complete set of information advertised by the neighbor has been deleted from LLDP neighbor information maintained by the local system because the information timeliness interval expired (appears when the <b>interface</b> option is used). |
| <b>Neighbor Information</b>      | Information about a neighbor system on the interface (appears when the <b>interface</b> option is used).                                                                                                                                                          |
| <b>Chassis type</b>              | Type of chassis identifier supplied, such as <b>Mac address</b> (appears when the <b>interface</b> option is used).                                                                                                                                               |
| <b>Chassis ID</b>                | Chassis identifier of the chassis type listed (appears when the <b>interface</b> option is used).                                                                                                                                                                 |
| <b>Port type</b>                 | Type of port identifier supplied, such as <b>Locally assigned</b> (appears when the <b>interface</b> option is used).                                                                                                                                             |
| <b>Port ID</b>                   | Port identifier of the port type listed (appears when the <b>interface</b> option is used).                                                                                                                                                                       |
| <b>Port description</b>          | The port description field uses the configured port description, the port name or the SNMP ifIndex (appears when the <b>interface</b> option is used).                                                                                                            |
| <b>System name</b>               | Name supplied by the system on the interface (appears when the <b>interface</b> option is used).                                                                                                                                                                  |

Table 187: show lldp neighbors Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>System Description</b>  | Description supplied by the system on the interface (appears when the <b>interface</b> option is used).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>System capabilities</b> | Capabilities (such as <b>Bridge</b> , <b>Bridge Router</b> , and <b>Bridge Telephone</b> ) that are supported or enabled by the system on the interface (appears when the <b>interface</b> option is used).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Management Info</b>     | <p>Details of management information: <b>Type</b> (such as <b>IPv4</b> or <b>IPv6</b>), <b>Address</b> (such as <b>10.204.34.35</b>), <b>Port ID</b>, <b>Subtype</b>, <b>Interface Subtype</b>, and organization identifier (<b>OID</b>) (appears when the <b>interface</b> option is used).</p> <p>The <b>Interface Subtype</b> displays:</p> <ul style="list-style-type: none"> <li><b>ifIndex(2)</b>—IP address of the neighbor's management Ethernet interface (<b>me0</b>) or virtual management Ethernet (<b>VME</b>) interface address (for a Virtual Chassis) is used to manage the switch.</li> <li><b>unknown(1)</b>—Neighbor's IP management address has been configured with set <b>protocols lldp management-address</b>.</li> </ul> |
| <b>Media Info</b>          | Additional details about the endpoint device appear when a device that supports LLDP-MED is attached to the interface. The specific details depend upon the capabilities of the device. Details might include: <b>Media endpoint class</b> (such as Class 3 for communication devices such as IP phones), <b>MED Hardware revision</b> , <b>MED Firmware revision</b> , <b>MED Software revision</b> , <b>MED Serial number</b> , <b>MED Manufacturer name</b> , <b>MED Model name</b> .                                                                                                                                                                                                                                                          |
| <b>Organization Info</b>   | One or more entries listing remote information by organizationally unique identifier ( <b>OUI</b> ), <b>Subtype</b> , <b>Index</b> , and <b>Info</b> (appears when the <b>interface</b> option is used).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Age</b>                 | How long the neighbor has been identified (appears when the <b>interface</b> option is used and NetBIOS snooping is enabled on the switch).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Local Interface</b>     | Name of the local physical interface (appears when the <b>interface</b> option is used and NetBIOS snooping is enabled on the switch).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parent Interface</b>    | Name of the aggregated Ethernet interface, if any, to which the interface belongs (appears when the <b>interface</b> option is used and NetBIOS snooping is enabled on the switch).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Chassis ID</b>          | Chassis identifier of the chassis type listed (appears when the <b>interface</b> option is used and NetBIOS snooping is enabled on the switch).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>System name</b>         | NetBIOS name of the host (appears when the <b>interface</b> option is used and NetBIOS snooping is enabled on the switch).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## Sample Output

### show lldp neighbors

```
user@switch> show lldp neighbors
```

| Local Interface | Parent Interface | Chassis Id        | Port info  | System Name |
|-----------------|------------------|-------------------|------------|-------------|
| xe-3/0/4.0      | ae31.0           | b0:c6:9a:63:80:40 | xe-0/0/0.0 | newyork31   |
| xe-3/0/5.0      | ae31.0           | b0:c6:9a:63:80:40 | xe-0/0/1.0 | newyork31   |
| xe-3/0/6.0      | ae31.0           | b0:c6:9a:63:80:40 | xe-0/0/2.0 | newyork31   |
| xe-3/0/7.0      | ae31.0           | b0:c6:9a:63:80:40 | xe-0/0/3.0 | newyork31   |
| xe-3/0/0.0      | ae31.0           | b0:c6:9a:63:80:40 | xe-0/1/0.0 | newyork31   |
| xe-3/0/1.0      | ae31.0           | b0:c6:9a:63:80:40 | xe-0/1/1.0 | newyork31   |
| xe-3/0/2.0      | ae31.0           | b0:c6:9a:63:80:40 | xe-0/1/2.0 | newyork31   |
| xe-3/0/3.0      | ae31.0           | b0:c6:9a:63:80:40 | xe-0/1/3.0 | newyork31   |

### show lldp neighbors interface ge-0/0/2

```
user@switch> show lldp neighbors interface ge-0/0/2
```

#### LLDP Neighbor Information:

##### Local Information:

```
Index: 1 Time to live: 240 Time mark: Wed Dec 1 10:23:24 2010 Age: 29 secs
Local Interface   : ge-0/0/2.0
Parent Interface  : -
Local Port ID     : 507
Ageout Count     : 0
```

##### Neighbour Information:

```
Chassis type      : Mac address
Chassis ID       : 00:1f:12:38:7f:c0
Port type        : Locally assigned
Port ID          : 507
Port description  : ge-0/0/2.0
System name      : bng-148p5-dev
```

```
System Description : Juniper Networks, Inc. ex4200-48p , version 10.4I0 Build
date: 2010-11-30 09:32:17 UTC
```

#### System capabilities

```
Supported : Bridge Router
Enabled   : Bridge Router
```

#### Management Info

```
Type           : IPv4
Address        : 10.204.96.235
Port ID       : 34
Subtype       : 1
Interface Subtype : ifIndex(2)
OID           : 1.3.6.1.2.1.31.1.1.1.1.34
```

```
Media endpoint class: Network Connectivity
```

#### Organization Info

```
OUI       : IEEE 802.3 Private (0x00120f)
Subtype   : MAC/PHY Configuration/Status (1)
Info      : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation
Capability (0x1d00), MAU Type (0x0)
Index     : 1
```

#### Organization Info

```
OUI       : IEEE 802.3 Private (0x00120f)
Subtype   : MDI Power (2)
Info      : MDI Power Support [PSE supported ], MDI Power Pair (signal),
MDI Power Class (class0)
Index     : 2
```

### show lldp neighbors interface ge-0/0/0.0 (for a VoIP AvayaTelephone with LLDP-MED Support)

```
user@switch>show lldp neighbors interface ge-0/0/0.0
```

#### LLDP Neighbor Information:

##### Local Information:

```
Index: 20 Time to live: 120 Time mark: Thu Apr 15 22:26:22 2010 Age: 16 secs
Local Interface   : ge-0/0/0.0
Parent Interface  : -
Local Port ID     : 517
Ageout Count      : 0
```

##### Neighbour Information:

```
Chassis type      : Network address
Chassis ID        : 0.0.0.0
Port type         : Mac address
Port ID           : 00:04:0d:fc:55:48
System name       : AVAFC5548
```

##### System capabilities

```
Supported : Bridge Telephone
Enabled   : Bridge
```

##### Management Info

```
Type           : IPv4
Address         : 0.0.0.0
Port ID         : 1
Subtype         : 1
Interface Subtype : ifIndex(2)
OID             : 1.3.6.1.2.1.31.1.1.1.1.1
```

```
Media endpoint class: Class III Device
```

```
MED Hardware revision : 4610D01A
MED Firmware revision : b10d01b2_9.bin
MED Software revision : a10d01b2_9.bin
MED Serial number     : 07N510103424
MED Manufacturer name : Avaya
MED Model name        : 4610
```

##### Organization Info

```
OUI       : IEEE 802.3 Private (0x00120f)
Subtype    : MAC/PHY Configuration/Status (1)
Info       : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation
Capability (0x1d00), MAU Type (0x0)
Index      : 1
```

##### Organization Info

```
OUI       : IEEE 802.3 Private (0x00120f)
Subtype    : MDI Power (2)
Info       : MDI Power Support [PSE supported ], MDI Power Pair (signal),
MDI Power Class (class0)
Index      : 2
```

##### Organization Info

```
OUI       : IEEE 802.3 Private (0x00120f)
Subtype    : Link Aggregation (3)
Info       : Aggregation Status (supported ), Aggregation Port ID (0)
Index      : 3
```

##### Organization Info

```
OUI       : IEEE 802.3 Private (0x00120f)
```

Subtype : Maximum Frame Size (4)  
Info : MTU Size (1514)  
Index : 4

Organization Info

OUI : Ethernet Bridged (0x0080c2)  
Subtype : Port Vid (1)  
Info : VLAN ID (10),  
Index : 5

Organization Info

OUI : Juniper Specific (0x009069)  
Subtype : Chassis Serial Type (1)  
Info : Juniper Slot Serial [BQ0208211462]  
Index : 6

Organization Info

OUI : Ethernet Bridged (0x0080c2)  
Subtype : VLAN Name (3)  
Info : VLAN ID (10), VLAN Name (vtest)  
Index : 7

**show lldp neighbors interface ge-0/0/5.0 (with NetBIOS Snooping Enabled on the Switch)**

user@switch> show lldp neighbors interface ge-0/0/5

Age: 299999 secs  
Local Interface : ge-0/0/5.0  
Parent Interface : -  
Chassis ID : 00:10:94:00:00:02  
Port description : 169.254.58.17  
System name : JNPRU\

## show lldp remote-global-statistics

|                                 |                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show lldp remote-global-statistics                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                        |
| <b>Description</b>              | Display remote Link Layer Discovery Protocol (LLDP) global statistics.                                                                                                                                                     |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP (CLI Procedure) on page 1913</a></li> <li>• <a href="#">Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show lldp remote-global-statistics on page 2032</a>                                                                                                                                                            |
| <b>Output Fields</b>            | <a href="#">Table 188 on page 2031</a> describes the output fields for the <b>show lldp remote-global-statistics</b> command. Output fields are listed in the approximate order in which they appear.                      |

**Table 188: show lldp remote-global-statistics Output Fields**

| Field Name                          | Field Description                                                                                     |
|-------------------------------------|-------------------------------------------------------------------------------------------------------|
| LLDP Remote Database Table Counters | Information about remote database table counters.                                                     |
| LastchangeTime                      | Time elapsed between LLDP agent startup and the last change to the remote database table information. |
| Inserts                             | Number of insertions made in the remote database table.                                               |
| Deletes                             | Number of deletions made in the remote database table.                                                |
| Drops                               | Number of LLDP frames dropped from the remote database table because of errors.                       |
| Ageouts                             | Number of remote database table entries that have aged out of the table.                              |

## Sample Output

### show lldp remote-global-statistics

```
user@host> show lldp remote-global-statistics
user@host> show lldp remote-global-statistics
LLDP Remote Database Table Counters
LastchangeTime      Inserts    Deletes    Drops    Ageouts
00:00:76 (76 sec)   192        0           0         0
```



## show lldp statistics

|                                 |                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show lldp statistics</b><br><b>&lt;interface <i>interface</i>&gt;</b>                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                         |
| <b>Description</b>              | Display LLDP statistics for all interfaces or for the specified interface.                                                                                                                                                 |
| <b>Options</b>                  | <b>none</b> —Display LLDP statistics for all interfaces.<br><br><b>interface <i>interface</i></b> —(Optional) Display LLDP statistics for the specified interface.                                                         |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP (CLI Procedure) on page 1913</a></li> <li>• <a href="#">Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches on page 1831</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show lldp statistics on page 2034</a><br><a href="#">show lldp statistics interface xe-3/0/0.0 on page 2034</a>                                                                                                |
| <b>Output Fields</b>            | Table 189 on page 2033 lists the output fields for the <b>show lldp statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                        |

**Table 189: show lldp statistics Output Fields**

| Field Name              | Field Description                                                                                                                                                                                                                                                              |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b>        | Name of the interface.                                                                                                                                                                                                                                                         |
| <b>Parent Interface</b> | Name of the aggregated Ethernet interface, if any, to which the interface belongs.<br><br><b>NOTE:</b> Because LLDP packets are transmitted and received on member interfaces only, statistics are available only for the member interfaces, not for the aggregated interface. |
| <b>Received</b>         | Total number of LLDP frames received on an interface.                                                                                                                                                                                                                          |
| <b>Unknown TLVs</b>     | Number of unrecognized LLDP TLVs received on an interface.                                                                                                                                                                                                                     |
| <b>With Errors</b>      | Number of invalid LLDP TLVs received on an interface.                                                                                                                                                                                                                          |
| <b>Discarded</b>        | Number of LLDP TLVs received and then discarded on an interface.                                                                                                                                                                                                               |
| <b>Transmitted</b>      | Total number of LLDP frames that were transmitted on an interface.                                                                                                                                                                                                             |
| <b>Untransmitted</b>    | Total number of LLDP frames that were untransmitted on an interface.                                                                                                                                                                                                           |

## Sample Output

### show lldp statistics

```
user@switch> show lldp statistics
```

| Interface  | Parent Interface | Received | Unknown TLVs | With Errors |
|------------|------------------|----------|--------------|-------------|
| xe-3/0/0.0 | ae31.0           | 1564     | 0            | 0           |
| xe-3/0/1.0 | ae31.0           | 1564     | 0            | 0           |
| xe-3/0/2.0 | ae31.0           | 1565     | 0            | 0           |
| xe-3/0/3.0 | ae31.0           | 1566     | 0            | 0           |
| xe-3/0/4.0 | ae31.0           | 1598     | 0            | 0           |
| xe-3/0/5.0 | ae31.0           | 1598     | 0            | 0           |
| xe-3/0/6.0 | ae31.0           | 1596     | 0            | 0           |
| xe-3/0/7.0 | ae31.0           | 1597     | 0            | 0           |
| xe-5/0/6.0 | -                | 0        | 0            | 0           |
| xe-5/0/7.0 | -                | 0        | 0            | 0           |

| Discarded TLVs | Transmitted | Untransmitted |
|----------------|-------------|---------------|
| 0              | 3044        | 1             |
| 0              | 3044        | 1             |
| 0              | 3044        | 1             |
| 0              | 3044        | 1             |
| 0              | 3075        | 1             |
| 0              | 3075        | 1             |
| 0              | 3075        | 1             |
| 0              | 3075        | 1             |
| 0              | 17312       | 0             |
| 0              | 17312       | 0             |

### show lldp statistics interface xe-3/0/0.0

```
user@switch> show lldp statistics interface xe-3/0/0.0
```

| Interface  | Parent Interface | Received | Unknown TLVs | With Errors |
|------------|------------------|----------|--------------|-------------|
| xe-3/0/0.0 | ae31.0           | 1566     | 0            | 0           |

| Discarded TLVs | Transmitted | Untransmitted |
|----------------|-------------|---------------|
| 0              | 3046        | 1             |

## show network-access aaa statistics accounting

|                                 |                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show network-access aaa statistics accounting</b>                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for QFX Series switches.                                                                   |
| <b>Description</b>              | Display authentication, authorization, and accounting (AAA) accounting statistics.                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>accounting-server</i></li> <li>• <i>accounting-stop-on-access-deny</i></li> <li>• <i>Configuring RADIUS Accounting</i></li> </ul>                                |
| <b>List of Sample Output</b>    | <a href="#">show network-access aaa statistics accounting on page 2035</a>                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 190 on page 2035</a> lists the output fields for the <b>show network-access aaa statistics accounting</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 190: show network-access aaa statistics accounting Output Fields**

| Field Name                   | Field Description                                                                                       |
|------------------------------|---------------------------------------------------------------------------------------------------------|
| Requests received            | The number of accounting-request packets sent from a switch to a RADIUS accounting server.              |
| Accounting Response failures | The number of accounting-response failure packets sent from the RADIUS accounting server to the switch. |
| Accounting Response Success  | The number of accounting-response success packets sent from the RADIUS accounting server to the switch. |
| Requests timedout            | The number of requests-timedout packets sent from the RADIUS accounting server to the switch.           |

## Sample Output

### show network-access aaa statistics accounting

```

user@switch> show network-access aaa statistics accounting
Accounting module statistics
  Requests received: 1
  Accounting Response failures: 0
  Accounting Response Success: 1
  Requests timedout: 0

```

## show network-access aaa statistics authentication

|                                 |                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show network-access aaa statistics authentication</b>                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for QFX Series switches.                                                                       |
| <b>Description</b>              | Display authentication, authorization, and accounting (AAA) authentication statistics.                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>authentication-server</i></li> <li><a href="#">Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch on page 1843</a></li> </ul>                       |
| <b>List of Sample Output</b>    | <a href="#">show network-access aaa statistics authentication on page 2036</a><br><a href="#">show network-access aaa statistics authentication (in QFX Series Switches) on page 2036</a>                        |
| <b>Output Fields</b>            | <a href="#">Table 191 on page 2036</a> lists the output fields for the <b>show network-access aaa statistics authentication</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 191: show network-access aaa statistics authentication Output Fields**

| Field Name               | Field Description                                                   |
|--------------------------|---------------------------------------------------------------------|
| <b>Requests received</b> | The number of authentication requests received by the switch.       |
| <b>Accepts</b>           | The number of authentication accepts received by the RADIUS server. |
| <b>Rejects</b>           | The number authentication rejects sent by the RADIUS server.        |
| <b>Challenges</b>        | The number of authentication challenges sent by the RADIUS server.  |

## Sample Output

### show network-access aaa statistics authentication

```
user@switch> show network-access aaa statistics authentication
Authentication module statistics
Requests received: 2
Accepts: 1
Rejects: 0
Challenges: 1
```

### show network-access aaa statistics authentication (in QFX Series Switches)

```
user@lf0> show network-access aaa statistics authentication
Authentication module statistics
Requests received: 2
Accepts: 1
```

Rejects: 0  
Challenges: 1

## show network-access aaa statistics dynamic-requests

|                                 |                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show network-access aaa statistics dynamic-requests;</b>                                                                                                                                                        |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for QFX Series switches.                                                                         |
| <b>Description</b>              | Display authentication, authorization, and accounting (AAA) authentication statistics for disconnects.                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>authentication-server</i></li> <li><a href="#">Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch on page 1843</a></li> </ul>                         |
| <b>List of Sample Output</b>    | <a href="#">show network-access aaa statistics authentication on page 2038</a>                                                                                                                                     |
| <b>Output Fields</b>            | <a href="#">Table 192 on page 2038</a> lists the output fields for the <b>show network-access aaa statistics dynamic-requests</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 192: show network-access aaa statistics dynamic-requests Output Fields**

| Field Name                      | Field Description                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------|
| <b>Requests received</b>        | The number of dynamic requests received by the RADIUS server.                                  |
| <b>Processed successfully</b>   | The number of dynamic requests successfully processed by the RADIUS server.                    |
| <b>Errors during processing</b> | The number of errors that occurred while the RADIUS server was processing the dynamic request. |
| <b>Silently dropped</b>         | The number of silently dropped requests.                                                       |

## Sample Output

### show network-access aaa statistics authentication

```

user@switch> show network-access aaa statistics dynamic-requests
Dynamic-requests module statistics
  Requests received: 0
  Processed successfully: 0
  Errors during processing: 0
  Silently dropped: 0

```

## PART 10

# Class of Service

- [Overview on page 2041](#)
- [Configuration on page 2075](#)
- [Administration on page 2157](#)
- [Troubleshooting Procedures on page 2191](#)





## CHAPTER 30

# Overview

- [CoS Overview on page 2041](#)

## CoS Overview

---

- [Junos OS CoS for EX Series Switches Overview on page 2042](#)
- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)
- [Understanding CoS Code-Point Aliases on page 2046](#)
- [Understanding CoS Classifiers on page 2049](#)
- [Understanding CoS Forwarding Classes on page 2052](#)
- [Understanding CoS Congestion Management on page 2055](#)
- [Understanding CoS Schedulers on page 2060](#)
- [Understanding CoS Two-Color Marking on page 2067](#)
- [Understanding CoS Rewrite Rules on page 2067](#)
- [Understanding Port Shaping and Queue Shaping for CoS on EX Series Switches on page 2069](#)
- [Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches on page 2070](#)
- [Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports on page 2071](#)

## Junos OS CoS for EX Series Switches Overview

When a network experiences congestion and delay, some packets must be dropped. Junos operating system (Junos OS) class of service (CoS) divides traffic into classes to which you can apply different levels of throughput and packet loss when congestion occurs. This allows packet loss to happen according to rules that you configure.

For interfaces that carry IPv4, IPv6, and MPLS traffic, you can configure Junos OS CoS features to provide multiple classes of service for different applications. CoS also allows you to rewrite the Differentiated Services code point (DSCP), IP precedence, 802.1p, or EXP CoS bits of packets egressing out of an interface, thus allowing you to tailor packets for the remote peers' network requirements. See *Understanding Using CoS with MPLS Networks on EX Series Switches* for more information about CoS for MPLS networks.

CoS provides multiple classes of service for different applications. You can configure multiple forwarding classes for transmitting packets, define which packets are placed into each output queue, and schedule the transmission service level for each queue.

In designing CoS applications, you must give careful consideration to your service needs and thoroughly plan and design your CoS configuration to ensure consistency and interoperability across all platforms in a CoS domain.

Because Juniper Networks EX Series Ethernet Switches implement CoS in hardware rather than in software, you can experiment with and deploy CoS features without affecting packet-forwarding and switching performance.



**NOTE:** CoS policies can be enabled or disabled on each interface of an EX Series switch. Also, each physical and logical interface on the switch can have custom CoS rules associated with it. When CoS is used in an MPLS network, there are some additional restrictions. See *Understanding Using CoS with MPLS Networks on EX Series Switches*.

- [How Junos OS CoS Works on page 2042](#)
- [Default CoS Behavior on EX Series Switches on page 2043](#)

---

### How Junos OS CoS Works


Junos OS CoS works by examining traffic entering at the edge of your network. The switches classify traffic into defined service groups to provide the special treatment of traffic across the network. For example, voice traffic can be sent across certain links, and data traffic can use other links. In addition, the data traffic streams can be serviced differently along the network path. As the traffic leaves the network at the far edge, you can rewrite the traffic to meet the policies of the targeted peer.

To support CoS, you must configure each switch in the network. Generally, each switch examines the packets that enter it to determine their CoS settings. These settings then dictate which packets are transmitted first to the next downstream switch. Switches at the edges of the network might be required to alter the CoS settings of the packets that enter the network to classify the packets into the appropriate service groups.

Figure 22 on page 2043 represents the network scenario of an enterprise. Switch A is receiving traffic from various network nodes such as desktop computers, servers, surveillance cameras, and VoIP telephones. As each packet enters, Switch A examines the packet's CoS settings and classifies the traffic into one of the groupings defined by the enterprise. This definition allows Switch A to prioritize resources for servicing the traffic streams it receives. Switch A might alter the CoS settings of the packets to better match the enterprise's traffic groups.

When Switch B receives the packets, it examines the CoS settings, determines the appropriate traffic groups, and processes the packets according to those settings. It then transmits the packets to Switch C, which performs the same actions. Switch D also examines the packets and determines the appropriate groups. Because Switch D sits at the far end of the network, it can rewrite the CoS bits of the packets before transmitting them.

Figure 22: Packet Flow Across the Network



### Default CoS Behavior on EX Series Switches

If you do not configure any CoS settings on the switch, the software performs some CoS functions to ensure that user traffic and protocol packets are forwarded with minimum delay when the network is experiencing congestion. Some CoS settings, such as classifiers, are automatically applied to each logical interface that you configure. Other settings, such as rewrite rules, are applied only if you explicitly associate them with an interface.

#### Related Documentation

- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)
- [Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches on page 2070](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Example: Combining CoS with MPLS on EX Series Switches](#)

### Understanding Junos OS CoS Components for EX Series Switches

This topic describes the Juniper Networks Junos operating system (Junos OS) class-of-service (CoS) components for Juniper Networks EX Series Ethernet Switches:

- [Code-Point Aliases on page 2044](#)
- [Policers on page 2044](#)

- [Classifiers on page 2044](#)
- [Forwarding Classes on page 2044](#)
- [Tail Drop Profiles on page 2045](#)
- [Schedulers on page 2045](#)
- [Rewrite Rules on page 2045](#)

---

## Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

---

## Policers

Policers limit traffic of a certain class to a specified bandwidth and *burst size*. Packets exceeding the policer limits can be discarded. You define policers with filters that can be associated with input interfaces.

For more information about policers, see [“Understanding the Use of Policers in Firewall Filters” on page 4767](#).



**NOTE:** You can configure policers to discard packets that exceed the rate limits. If you want to configure CoS parameters such as *loss-priority* and *forwarding-class*, you must use firewall filters.

---

## Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. In Juniper Networks Junos operating system (Junos OS), *classifiers* associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. Junos OS supports two general types of classifiers:

- Behavior aggregate or CoS value traffic classifiers—Examines the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value, IP precedence value, and IEEE 802.1p value.
- Multifield traffic classifiers—Examines multiple fields in the packet such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

---

## Forwarding Classes

Forwarding classes group the packets for transmission. Based on forwarding classes, you assign packets to output queues. Forwarding classes affect the forwarding, scheduling, and marking policies applied to packets as they transit a switch. By default, four categories of forwarding classes are defined: best effort, assured forwarding, expedited forwarding,

and network control. For EX Series switches, 16 forwarding classes are supported, providing granular classification capability.

### Tail Drop Profiles

Drop profile is a mechanism that defines parameters that allow packets to be dropped from the network. Drop profiles define the meanings of the loss priorities. When you configure drop profiles you are essentially setting the value for queue fullness. The queue fullness represents a percentage of the queue used to store packets in relation to the total amount that has been allocated for that specific queue.

Loss priorities set the priority of dropping a packet. Loss priority affects the scheduling of a packet without affecting the packet's relative ordering. You can use the loss priority setting to identify packets that have experienced congestion. Typically you mark packets exceeding some service level with a high loss priority.

### Schedulers

Each switch interface has multiple queues assigned to store packets. The switch determines which queue to service based on a particular method of scheduling. This process often involves determining which type of packet should be transmitted before another. You can define the priority, bandwidth, delay buffer size, and tail drop profiles to be applied to a particular queue for packet transmission.

A scheduler map associates a specified forwarding class with a scheduler configuration. You can associate up to four user-defined scheduler maps with the interfaces.

### Rewrite Rules

A *rewrite rule* sets the appropriate CoS bits in the outgoing packet, thus allowing the next downstream device to classify the packet into the appropriate service group. Rewriting, or marking, outbound packets is useful when the switch is at the border of a network and must alter the CoS values to meet the policies of the targeted peer.



**NOTE:** Egress firewall filters can also assign forwarding class and loss priority so that the packets are rewritten based on forwarding class and loss priority.

#### Related Documentation

- [Understanding CoS Code-Point Aliases on page 2046](#)
- [Understanding CoS Classifiers](#)
- [Understanding CoS Forwarding Classes](#)
- [Understanding CoS Tail Drop Profiles](#)
- [Understanding CoS Schedulers on page 2060](#)
- [Understanding CoS Two-Color Marking on page 2067](#)
- [Understanding CoS Rewrite Rules on page 2067](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)

## Understanding CoS Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

Behavior aggregate classifiers use class-of-service (CoS) values such as Differentiated Services code points (DSCPs), IP precedence, and IEEE 802.1p bits to associate incoming packets with a particular CoS servicing level. On a switch, you can assign a meaningful name or alias to the CoS values and use this alias instead of bits when configuring CoS components. These aliases are not part of the specifications but are well known through usage. For example, the alias for DSCP 101110 is widely accepted as ef (expedited forwarding).

When you configure classes and define classifiers, you can refer to the markers by alias names. You can configure user-defined classifiers in terms of alias names. If the value of an alias changes, it alters the behavior of any classifier that references it.

This topic covers:

- [Default Code-Point Aliases on page 2046](#)

### Default Code-Point Aliases

[Table 193 on page 2046](#) shows the default mappings between the bit values and standard aliases.

**Table 193: Default Code-Point Aliases**

| CoS Value Types        | Mapping |
|------------------------|---------|
| <b>DSCP CoS Values</b> |         |
| ef                     | 101110  |
| af11                   | 001010  |
| af12                   | 001100  |
| af13                   | 001110  |
| af21                   | 010010  |
| af22                   | 010100  |
| af23                   | 010110  |
| af31                   | 011010  |
| af32                   | 011100  |
| af33                   | 011110  |

Table 193: Default Code-Point Aliases (*continued*)

| CoS Value Types                        | Mapping |
|----------------------------------------|---------|
| af41                                   | 100010  |
| af42                                   | 100100  |
| af43                                   | 100110  |
| be                                     | 000000  |
| cs1                                    | 001000  |
| cs2                                    | 010000  |
| cs3                                    | 011000  |
| cs4                                    | 100000  |
| cs5                                    | 101000  |
| nc1/cs6                                | 110000  |
| nc2/cs7                                | 111000  |
| <b>IEEE 802.1p CoS Values</b>          |         |
| be                                     | 000     |
| be1                                    | 001     |
| ef                                     | 100     |
| ef1                                    | 101     |
| af11                                   | 010     |
| af12                                   | 011     |
| nc1/cs6                                | 110     |
| nc2/cs7                                | 111     |
| <b>Legacy IP Precedence CoS Values</b> |         |
| be                                     | 000     |
| be1                                    | 001     |
| ef                                     | 010     |

Table 193: Default Code-Point Aliases (*continued*)

| CoS Value Types | Mapping |
|-----------------|---------|
| ef1             | 011     |
| af11            | 100     |
| af12            | 101     |
| nc1/cs6         | 110     |
| nc2/cs7         | 111     |

**Related  
Documentation**

- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Defining CoS Code-Point Aliases \(CLI Procedure\) on page 2101](#)
- [Defining CoS Code-Point Aliases \(J-Web Procedure\) on page 2101](#)



## Understanding CoS Classifiers

Packet classification associates incoming packets with a particular class-of-service (CoS) servicing level. Classifiers associate packets with a forwarding class and loss priority, and packets are associated to an output queue based on the forwarding class. You can define classifiers for IPv4 and IPv6 traffic to network interfaces, aggregated Ethernet interfaces (also known as link aggregation groups (LAGs)), integrated routing and bridging (IRB) interfaces (also known as routed VLAN interfaces (RVIs)), Layer 3 interfaces, and Layer 3 VLAN-tagged logical interfaces.

There are two general types of classifiers:

- Behavior aggregate (BA) classifiers
- Multifield (MF) classifiers

You can configure both a BA classifier and an MF classifier on an interface. If you do this, the BA classification is performed first and then the MF classification. If the two classification results conflict, the MF classification result overrides the BA classification result.

On Juniper Networks EX8200 Ethernet Switches, you can specify BA classifiers for bridged multdestination traffic and for IP multdestination traffic. A BA classifier for multicast packets is applied to all interfaces on the EX8200 switch.



**NOTE:** EX8200 switches implement the on-demand allocation of memory space for ternary content addressable memory (TCAM) so that when additional TCAM space is required for CoS classifiers, it is allocated from the free TCAM space or from the unused TCAM space. An error log message is generated when you configure CoS classifiers to use memory space that exceeds the available TCAM space that includes both the free and unused space.

This topic describes:

- [Behavior Aggregate Classifiers on page 2049](#)
- [Multifield Classifiers on page 2051](#)

### Behavior Aggregate Classifiers

The behavior aggregate classifier maps packets to a forwarding class and a loss priority. The forwarding class determines the output queue for a packet. The loss priority is used by a scheduler to control packet discard during periods of congestion.

There are three types of BA classifiers:

- Differentiated Services Code Point (DSCP) for IP DiffServ
- IP precedence bits
- IEEE 802.1p CoS bits

BA classifiers are based on fixed-length fields, which makes them computationally more efficient than MF classifiers. Therefore core devices, which handle high traffic volumes, are normally configured to perform BA classification.

#### **Default Behavior Aggregate Classification**

Juniper Networks Junos operating system (Junos OS) automatically assigns implicit default BA classifiers to logical interfaces based on the type of interface.

[Table 194 on page 2050](#) lists different types of interfaces and the corresponding implicit default BA classification.

**Table 194: Default BA Classification**

| Type of Interface                                | Default BA Classification                                                                                                       |
|--------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| Trunk and Circuit Cross-Connect (CCC) interfaces | <b>ieee8021p-default</b><br><br><b>NOTE:</b> This BA classification for a CCC interface is applicable only for EX8200 switches. |
| Layer 3 interface (IPv4)                         | <b>dscp-default</b>                                                                                                             |
| Layer 3 interface (IPv6)                         | <b>dscp-ipv6-default</b>                                                                                                        |
| Access interface                                 | Untrusted                                                                                                                       |
| Routed VLAN interface (RVI)                      | No default classification                                                                                                       |
| MPLS                                             | EXP<br><br><b>NOTE:</b> This BA classification is applicable only for EX8200 switches.                                          |

When you explicitly associate a BA classifier with a logical interface, you are overriding the implicit (default) BA classifier with an explicit BA classifier.

[Table 195 on page 2050](#) describes the BA classifier types you can configure on Layer 2 and Layer 3 interfaces.

**Table 195: Allowed BA Classification**

| Type of Interface        | Allowed BA Classification                   |
|--------------------------|---------------------------------------------|
| Layer 2 interface        | IEEE 802.1p, IP precedence, DSCP, DSCP IPv6 |
| Layer 3 interface (IPv4) | IEEE 802.1p, IP precedence, DSCP            |
| Layer 3 interface (IPv6) | IEEE 802.1p, IP precedence, DSCP IPv6       |

You can configure all the allowed classifier types on the same logical interface or on different logical interfaces. If you need to apply all classifier rules on the same logical interface, configure the classifier rules allowed for both IPv4 and IPv6 on the logical interface.

If you have not explicitly configured a classifier on a logical interface, the default classifiers are assigned and classification works as follows:

- To a logical interface configured with an IPv4 address, a DSCP classifier is assigned by default, and IPv4 and IPv6 packets are classified using the DSCP classifier.
- To logical interface configured with an IPv6 address, a DSCP IPv6 classifier is assigned by default, and IPv4 and IPv6 packets are classified using the DSCP IPv6 classifier.



**NOTE:** On EX8200 switches, you can configure either one classifier of type DSCP or IEEE802.1p, or you can configure one classifier each of type DSCP and IEEE802.1p.

You can configure integrated routing and bridging (IRB) interfaces (also known as routed VLAN interfaces (RVIs)) to classify packets. After you do this, the User Priority (UP) bits in the incoming packets are rewritten according to the default IEEE 802.1p rewrite rule, except on EX8200 switches. On EX8200 switches, you must explicitly assign the default IEEE 802.1p rewrite rule to RVIs.



**NOTE:** By default, all BA classifiers classify traffic into either the best-effort forwarding class or the network-control forwarding class.

### Multifield Classifiers

Multifield classifiers examine multiple fields in a packet such as source and destination addresses and source and destination port numbers of the packet. With MF classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

MF classification is normally performed at the network edge because of the general lack of support for DSCP or IP precedence classifiers in end-user applications. On an edge switch, an MF classifier provides the filtering functionality that scans through a variety of packet fields to determine the forwarding class for a packet. Typically, any classifier performs matching operations on the selected fields against a configured value.

#### Related Documentation

- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Defining CoS Classifiers \(CLI Procedure\) on page 2103](#)
- [Defining CoS Classifiers \(J-Web Procedure\) on page 2105](#)

## Understanding CoS Forwarding Classes

Class-of-Service (CoS) forwarding classes can be thought of as output queues. In effect, the result of classifying packets is the identification of an output queue for a particular packet. For a classifier to assign an output queue to a packet, it must associate the packet with one of the following forwarding classes:

- best-effort (be)—Provides no service profile. Loss priority is typically not carried in a CoS value.
- expedited-forwarding (ef)—Provides a low loss, low latency, low jitter, assured bandwidth, end-to-end service.
- assured-forwarding (af)—Provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with two drop probabilities: low and high.
- network-control (nc)—Supports protocol control and thus is typically high priority.
- multicast best-effort (mcast-be)—Provides no service profile for multicast packets.
- multicast expedited forwarding (mcast-ef)—Supports high-priority multicast packets.
- multicast assured-forwarding (mcast-af)—Provides two drop profiles; high, and low, for multicast packets.



**NOTE:** The forwarding classes multicast expedited-forwarding, multicast assured-forwarding, and multicast best-effort are applicable only to Juniper Networks EX8200 and Juniper Networks EX4300 Ethernet Switches.

Juniper Networks EX Series Ethernet Switches support up to 16 forwarding classes, thus allowing granular packet classification. For example, you can configure multiple classes of expedited forwarding (EF) traffic such as EF, EF1, and EF2.

EX Series switches except EX4300 switches support up to eight output queues. Therefore, if you configure more than eight forwarding classes, you must map multiple forwarding classes to single output queues. EX4300 switches support up to 12 output queues. On EX8200 Virtual Chassis, you can configure only eight forwarding classes and you can assign only one forwarding class to each output queue.



**NOTE:** On EX8200 Virtual Chassis, the queue number seven carries Virtual Chassis port (VCP) traffic and can also carry high-priority user traffic.

This topic describes:

- [Default Forwarding Classes on page 2053](#)

## Default Forwarding Classes

Table 196 on page 2053 shows the four default forwarding classes defined for unicast traffic, and Table 197 on page 2053 shows the three default forwarding classes defined for multicast traffic.



**NOTE:** The default forwarding classes for multicast traffic are applicable only to EX8200 switches.

You can rename the forwarding classes associated with the queues supported on your switch. Assigning a new class name to an output queue does not alter the default classification or scheduling that is applicable to that queue. However, because CoS configurations can be quite complicated, we recommend that you avoid altering the default class names or queue number associations.

**Table 196: Default Forwarding Classes for Unicast Traffic**

| Forwarding Class Name     | Comments                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| best-effort (be)          | The software does not apply any special CoS handling to packets with 000000 in the DiffServ field. This is a backward compatibility feature. These packets are usually dropped under congested network conditions.                                                                                                                                                                                                        |
| expedited-forwarding (ef) | The software delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end-to-end for packets in this service class. The software accepts excess traffic in this class, but in contrast to the assured forwarding class, the out-of-profile expedited-forwarding class packets can be forwarded out of sequence or dropped.                                                                       |
| assured-forwarding (af)   | <p>The software offers a high level of assurance that the packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define.</p> <p>The software accepts excess traffic, but it applies a tail drop profile to determine that excess packets are dropped, and not forwarded.</p> <p>Two drop probabilities (low and high) are defined for this service class.</p> |
| network-control (nc)      | <p>The software delivers packets in this service class with a high priority. (These packets are not delay-sensitive.)</p> <p>Typically, these packets represent routing protocol hello or keep alive messages. Because loss of these packets jeopardizes proper network operation, packet delay is preferable to packet discard for these packets.</p>                                                                    |

**Table 197: Default Forwarding Classes for Multicast Traffic**

| Forwarding Class Name            | Comments                                                                                                                                         |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| multicast best-effort (mcast-be) | The software does not apply any special CoS handling to multicast packets. These packets are usually dropped under congested network conditions. |

Table 197: Default Forwarding Classes for Multicast Traffic (*continued*)

| Forwarding Class Name                     | Comments                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| multicast expedited-forwarding (mcast-ef) | The software delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end-to-end for multicast packets in this service class. The software accepts excess traffic in this class, but in contrast to the multicast assured forwarding class, out-of-profile multicast expedited-forwarding class packets can be forwarded out of sequence or dropped.                                                        |
| multicast assured-forwarding (mcast-af)   | <p>The software offers a high level of assurance that the multicast packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define.</p> <p>The software accepts excess traffic, but it applies a tail drop profile to determine if the excess packets are dropped and not forwarded.</p> <p>Two drop probabilities (low and high) are defined for this service class.</p> |
| multicast network-control (mcast-nc)      | <p>The software delivers packets in this service class with a high priority. (These packets are not delay-sensitive.)</p> <p>Typically, these packets represent routing protocol hello or keep alive messages. Because loss of these packets jeopardizes proper network operation, packet delay is preferable to packet discard for these packets.</p>                                                                               |

The following rules govern queue assignment:

- CoS configurations that specify more queues than the switch can support are not accepted. If you commit such a configuration, the commit fails and a message displays that states the number of queues available.
- All default CoS configurations are based on queue number. The name of the forwarding class that is displayed in the default configuration for a queue number is that of the forwarding class currently associated with that queue.

#### Related Documentation

- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic](#)
- [Defining CoS Forwarding Classes \(CLI Procedure\) on page 2107](#)
- [Defining CoS Forwarding Classes \(J-Web Procedure\) on page 2107](#)

## Understanding CoS Congestion Management

A congestion in a network occurs because of various parameters and some packets must be dropped to avoid congestion and to facilitate easy flow of traffic in the network. On Juniper Networks EX Series Ethernet Switches, class of service (CoS) provides congestion management mechanisms for a switch to drop arriving packets based on certain parameters when a queue is full. Based on the EX Series switch that you are using, packets are dropped depending on the priority of a packet or on both priority and drop probability of a packet.

You can specify parameters at the **[edit class-of-service drop-profiles]** hierarchy level for dropping packets and reference the parameters in a scheduler configuration.

This topic describes:

- [Weighted Tail Drop Congestion Management on page 2055](#)
- [Weighted Random Early Detection Congestion Management on page 2056](#)

### Weighted Tail Drop Congestion Management

A weighted tail drop (WTD) is a congestion management mechanism for packets to be dropped from the tail of the queue when the queue reaches a certain buffer capacity (that is, the fill level), and hence the name weighted tail drop. The packets that are dropped are based on priority and are those marked with a packet loss priority (PLP) of *high*. You can configure a WTD profile (a WTD mechanism) usually on edge devices in a network.



**NOTE:** A WTD profile is supported only on the Juniper Networks EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, and EX6200 Ethernet Switches.

When you configure a WTD profile, you are essentially setting the value for queue fullness. The queue fullness represents a percentage of the memory, known as delay-buffer bandwidth, that is used to store packets in relation to the total amount of memory that has been allocated for that specific queue. The delay-buffer bandwidth provides packet buffer space to absorb burst traffic up to the specified duration of delay. When the specified delay buffer becomes full, packets are dropped from the tail of the buffer.

By default, if you do not configure any drop profile, WTD profile is in effect and functions as the primary mechanism for managing congestion.



**NOTE:** The default WTD profile associated with the packets whose PLP is *low* cannot be modified. You can configure custom drop profile only for those packets whose PLP is *high*.

## Weighted Random Early Detection Congestion Management

---

In a weighted random early detection (WRED) congestion management mechanism, random packets with a PLP of low or high are gradually dropped (based on drop probability) when the queue reaches a certain buffer capacity (that is, fill level).



**NOTE:** The WRED mechanism is supported only on Juniper Networks EX4300 standalone switches, EX4300 Virtual Chassis, Juniper Networks EX8200 standalone switches, and EX8200 Virtual Chassis.

Following are the different implementations of WRED:

- Segmented Drop Profile
- Interpolated Drop Profile

From a high level, segmented drop profile is a stair-step-like drop profile, whereas interpolated drop profile is a smother (curve) drop profile. [Figure 23 on page 2057](#) and [Figure 24 on page 2058](#) show a graphical representation of segmented and interpolated drop profiles. Regardless of the implementation, a drop profile represents a graph where the x-axis represents the percentage of fill level (l) and the y-axis represents the percentage of drop probability (p). The origin (0,0) represents the drop profile in which the drop probability is 0 percent when the queue fullness is 0 percent, and the point (100,100) represents that the drop probability is 100 percent when the queue fullness is 100 percent. Although the formation of graph lines in [Figure 23 on page 2057](#) and [Figure 24 on page 2058](#) is different, the application of the profile is the same. When a packet reaches the head of the queue, a random number between 0 and 100 is calculated. This random number is plotted against the drop profile graph using the current queue fullness of that particular queue. When the random number falls above the graph line, the packet is transmitted. When the number falls below the graph line, the packet is dropped from the network.

The following sections discuss the WRED drop profile implementations and parameters:


- [Segmented Drop Profile on page 2056](#)
- [Interpolated Drop Profile on page 2057](#)
- [Drop Profile Parameters on page 2059](#)

### ***Segmented Drop Profile***

In a segmented drop profile configuration, you can define multiple data points for fill level and drop probability. [Figure 23 on page 2057](#) shows a graphical representation of a segmented drop profile.



Figure 23: Graphical Representation of a Segmented Drop Profile



To create the profile's graph line, the software begins at the bottom-left corner of the graph, representing a 0 percent fill level and a 0 percent drop probability (that is the point (0,0)). The configuration draws a line directly to the right until it reaches the first defined fill level (that is, 25 percent represented in the graph on the x-axis). The software then continues the line vertically until the first drop probability is reached (that is, 25 percent represented in the graph in the y-axis). This process is repeated for all of the defined fill levels and drop probabilities until the top-right corner of the graph is reached (that is point (100,100) in the graph).

### ***Interpolated Drop Profile***

An interpolated drop profile configuration forms a smoother graph line compared to the graph in a segmented drop profile configuration. In this method of congestion management also, a switch uses multiple drop profile values to drop incoming packets to reduce congestion in the output queue.

Following are interpolated drop profile configurations on EX Series switches:

- [Interpolated Drop Profile Configuration on EX Series Switches Except EX4300 Switches on page 2057](#)
- [Interpolated Drop Profile Configuration on EX4300 Switches on page 2058](#)


### ***Interpolated Drop Profile Configuration on EX Series Switches Except EX4300 Switches***

An interpolated drop profile on all EX Series switches except EX4300 switches automatically generates 64 pairs of data points on the graph beginning at (0, 0) and

ending at (100, 100). Along the way, the graph line intersects specific data points that you define for fullness and drop probability.

Figure 24 on page 2058 shows a graphical representation of an interpolated drop profile.

**Figure 24: Graphical Representation of an Interpolated Drop Profile on EX Series Switches Except EX4300 Switches**



#### *Interpolated Drop Profile Configuration on EX4300 Switches*

On EX4300 switches, you can set two queue fill levels and two drop probabilities in each drop profile. The two fill levels and the two drop probabilities create two pairs of values. The first fill level and the first drop probability create one value pair and the second fill level and the second drop probability create the second value pair.




**NOTE:** You can configure a maximum of 64 drop profiles on EX4300 switches.

The first fill level value specifies the percentage of queue fullness at which packets begin to drop, known as the drop start point. Until the queue reaches this level of fullness, no packets are dropped. The second fill level value specifies the percentage of queue fullness at which all packets are dropped, known as the drop end point.

The first drop probability value is always 0 (zero). This pairs with the drop start point and specifies that until the queue fullness level reaches the first fill level, no packets drop. When the queue fullness exceeds the drop start point, packets begin to drop until the queue exceeds the second fill level, when all packets drop. The second drop probability value, known as the maximum drop rate, specifies the likelihood of dropping packets

when the queue fullness reaches the drop end point. As the queue fills from the drop start point to the drop end point, packets drop in a smooth, linear pattern (called an interpolated graph) as shown in [Figure 25 on page 2059](#). After the drop end point, all packets drop.

**Figure 25: Tail-Drop Profile Packet Drop on EX4300 Switches**



The thick line in [Figure 25 on page 2059](#) shows the packet drop characteristics for a sample tail drop profile. At the drop start point, the queue reaches a fill level of 30 percent. At the drop end point, the queue fill level reaches 50 percent, and the maximum drop rate is 80 percent.

No packets drop until the queue fill level reaches the drop start point of 30 percent. When the queue reaches the 30 percent fill level, packets begin to drop. As the queue fills, the percentage of packets dropped increases in a linear fashion. When the queue fills to the drop end point of 50 percent, the rate of packet drop has increased to the maximum drop rate of 80 percent. When the queue fill level exceeds the drop end point of 50 percent, all of the packets drop until the queue fill level drops below 50 percent.

#### **Drop Profile Parameters**

You can specify the following two values in drop profile configuration:

- **Fill level**—The queue fullness value, which represents a percentage of the memory used to store packets in relation to the total amount of memory allocated to the queue.
- **Drop probability**—The percentage value that corresponds to the likelihood that an individual packet is dropped.

#### **Related Documentation**

- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Configuring CoS Congestion Management \(CLI Procedure\) on page 2115](#)

## Understanding CoS Schedulers

You use class-of-service (CoS) schedulers to define the properties of output queues on Juniper Networks EX Series Ethernet Switches. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the drop profiles associated with the queue.

You associate the schedulers with forwarding classes by means of scheduler maps. You can then associate each scheduler map with an interface, thereby configuring the queues, packet schedulers, and tail drop processes that operate according to this mapping.

This topic describes:

- [Default Schedulers on page 2060](#)
- [Excess Rate on page 2061](#)
- [Transmission Rate on page 2061](#)
- [Scheduler Buffer Size on page 2061](#)
- [Priority Scheduling on page 2062](#)
- [Scheduler Drop-Profile Maps on page 2063](#)
- [Scheduler Maps on page 2063](#)

---

### Default Schedulers

Each forwarding class has an associated scheduler priority. On EX Series switches other than Juniper Networks EX8200 and Juniper Networks EX4300 Ethernet Switches, only two forwarding classes—best-effort (queue 0) and network-control (queue 7)—are used in the default configuration. On EX8200 switches three forwarding classes—best-effort (queue 0), multicast best-effort (queue 2), and network-control (queue 7)—are used in the default configuration.

On EX Series switches other than EX8200 and EX4300 switches, by default, the best-effort forwarding class (queue 0) receives 95 percent of the bandwidth and the buffer space for the output link, and the network-control forwarding class (queue 7) receives 5 percent. The default drop profile causes the buffer to fill completely and then to discard all incoming packets until it has free space. On EX8200 switches, by default, the best-effort forwarding class (queue 0) receives 75 percent of the bandwidth, the multicast best-effort forwarding class (queue 2) receives 20 percent, and the network-control forwarding class (queue 7) receives 5 percent of the bandwidth and buffer space for the output link.

On EX4300 switches, four forwarding classes—best-effort (queue 0), multicast best-effort (queue 8), network-control (queue 3), and multicast network-control (queue 11)—are used in the default configuration. By default, all the multicast traffic flows through the multicast best-effort queue. EX4300 switches support 12 queues (0–11), and the default scheduler transmission rates for queues 0 through 11 are 75, 0, 0, 5, 0, 0, 0, 0, 15, 0, 0 and 5 percent, respectively, of the total available bandwidth.

On EX Series switches other than EX4300 switches, the expedited-forwarding (queue 5) and assured-forwarding (queue 1) classes have no scheduler because no resources are

assigned to queue 5 or queue 1, by default. However, you can manually configure resources to be assigned to the expedited-forwarding and assured-forwarding classes. On EX4300 switches, the expedited-forwarding (queue 1) and assured-forwarding (queue 2) classes have no scheduler because no resources are assigned to queue 1 or queue 2, by default. However, you can manually configure resources to be assigned to the expedited-forwarding and assured-forwarding classes.

Also by default, any queue can exceed the assigned bandwidth if additional bandwidth is available from other queues. When a forwarding class does not fully use the allocated transmission bandwidth, the remaining bandwidth can be used by other forwarding classes if they have a traffic load that exceeds their allocated bandwidth.

### Excess Rate

Excess rate traffic determines the percentage of the excess bandwidth to share when a queue receives traffic in excess of its bandwidth allocation. By default, the excess bandwidth is shared in the ratio of the transmit rates. You can control this distribution by configuring the **excess-rate** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy. You can specify the excess rate sharing in percentage.

### Transmission Rate

Transmission-rate control determines the actual traffic bandwidth for each forwarding class you configure. The transmission rate is specified in bits per second. Each queue is allocated some portion of the bandwidth of the interface. This bandwidth can be a fixed value, such as 1 megabit per second (Mbps), a percentage of the total available bandwidth, or the rest of the available bandwidth. In case of congestion, the configured transmission rate is guaranteed for the queue. Transmission-rate control allows you to ensure that each queue receives the bandwidth appropriate for its level of service.

### Scheduler Buffer Size

To control congestion at the output stage, you can configure the delay-buffer bandwidth by using the **buffer-size** configuration statement. The delay-buffer bandwidth provides packet buffer space to absorb burst traffic up to the specified duration of delay. When the specified delay buffer becomes full, packets with 100 percent drop probability are dropped from the tail of the buffer.

On EX Series switches other than EX8200 and EX4300 switches, the default scheduler transmission rates for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent, respectively, of the total available bandwidth. The default buffer-size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent, respectively, of the total available buffer. On EX8200 switches, the default scheduler transmission rates for queues 0 through 7 are 75, 0, 20, 0, 0, 0, 0, and 5 percent, respectively, of the total available bandwidth, and the default buffer-size percentages for queues 0 through 7 are 75, 0, 20, 0, 0, 0, 0, and 5 percent, respectively, of the total available buffer. On EX4300 switches, the default scheduler transmission rates for queues 0 through 11 are 75, 0, 0, 5, 0, 0, 0, 0, 15, 0, 0 and 5 percent, respectively, of the total available buffer. On EX4300 switches, the default buffer-size percentages for queues 0 through 11 are 75, 0, 0, 5, 0, 0, 0, 0, 15, 0, 0 and 5 percent, respectively, of the total available buffer.

For each scheduler on EX Series switches other than EX8200 switches, you can configure the buffer size as one of the following:

- The exact buffer size.
- A percentage of the total buffer.
- The remaining buffer available. The remainder is the buffer percentage that is not assigned to other queues. For example, if you assign 40 percent of the delay buffer to queue 0, allow queue 2 to keep the default allotment of 20 percent, allow queue 7 to keep the default allotment of 5 percent, and assign the remainder to queue 3, then queue 3 uses 35 percent of the delay buffer.

On EX8200 switches, you can configure the buffer size as a temporal value (in microseconds), percentage of the total buffer, or the remaining buffer available. You can configure the buffer size as a temporal value on Juniper Networks EX4200 and EX4300 Ethernet Switches also.

When you configure buffer size as a temporal value on EX4200 switches, if sufficient buffer size is not available in the shared pool, an error message is logged in the system log (syslog) file and the default profile is applied to the interface. After the temporal buffer space is allocated successfully, if the shared buffer size is less than the current value (which was set using the **set class-of-service shared-buffer percent value** command), the new reduced value must be greater than a sum of the existing reserved temporal buffer size and the required minimum buffer size. Otherwise, the modification to the shared-buffer configuration fails and an error message is logged in the system log.

### Priority Scheduling

---

Priority scheduling determines the order in which an interface transmits traffic from queues, thus ensuring that queues containing important traffic are provided faster access.

Priority scheduling is accomplished through a procedure in which the scheduler examines the priority of the queue. Juniper Networks Junos operating system (Junos OS) supports two levels of transmission priority:

- Low—The scheduler determines whether the individual queue is within its defined bandwidth profile or not. This binary decision, which is re-evaluated on a regular time cycle, involves comparing the amount of data transmitted by the queue against the bandwidth allocated to it by the scheduler. If the transmitted amount is less than the allocated amount, the queue is considered to be in profile. A queue is out of profile when the amount of traffic that it transmits is larger than the queue's allocated limit. An out-of-profile queue is transmitted only if bandwidth is available. Otherwise, it is buffered.

On EX Series switches other than EX4300 switches, a queue from a set of queues is selected based on the shaped deficit weighted round robin (SDWRR) algorithm, which operates within the set. On EX4300 switches, the weighted deficit round-robin (WDRR) algorithm is used to select a queue from a set of queues.

- Strict-high—A strict-high priority queue receives preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue. On EX Series switches other than EX4300 switches, queues are scheduled according to

the queue number, starting with the highest queue, 7, with decreasing priority down through queue 0. Traffic in higher-numbered queues is always scheduled prior to traffic in lower-numbered queues. In other words, if there are two high-priority queues, the queue with the higher queue number is processed first. On EX4300 switches, you can configure multiple strict-high priority queues on an interface and an EX4300 switch processes these queues in a round-robin method.

Packets in low-priority queues are transmitted only when strict-high priority queues are empty.

### Scheduler Drop-Profile Maps

Drop-profile maps associate drop profiles with a scheduler. A drop-profile map sets the drop profile for a specific packet loss priority (PLP) and protocol type. The inputs for a drop-profile map are the PLP and the protocol type. The output is the drop profile.

### Scheduler Maps

A scheduler map associates a specified forwarding class with a scheduler configuration. After configuring a scheduler, you must include it in a scheduler map and then associate the scheduler map with an output interface.

On EX Series switches, if you configure more than the supported number of scheduler maps on a switch or for a port group in a line card, an error is logged in the system log. On any interface in a port group on a line card or on a switch, if you configure a scheduler map that causes the number of scheduler maps for that port group to exceed the maximum number supported, the default scheduler map is bound to that interface. We recommend that you check the system log for errors after the commit operation to verify that you have not configured more than the maximum permitted number of scheduler maps.



**NOTE:** On EX Series switches, you cannot configure a scheduler map on an individual interface that is a member of a link aggregation group (LAG). Instead, you must configure the scheduler map on the LAG itself (that is, on the aggregated Ethernet (ae) interface).

Table 198 on page 2063 shows the number of scheduler maps supported for each port group in a switch or line card.

**Table 198: Support for Scheduler Maps on Switches and Line Cards**

| Switch/Line Card                       | Number of Port Groups | Port Grouping Details                                | Number of Scheduler Maps Supported for Each Port Group |
|----------------------------------------|-----------------------|------------------------------------------------------|--------------------------------------------------------|
| EX2200-C-12T and EX2200-C-12P switches | 1                     | Port 0–11 and 2 uplink ports form a port group.      | 6                                                      |
| EX2200-24T and EX2200-24P switches     | 1                     | Ports 0–23 and 4 SFP uplink ports form a port group. | 5                                                      |

Table 198: Support for Scheduler Maps on Switches and Line Cards (*continued*)

| Switch/Line Card                   | Number of Port Groups | Port Grouping Details                                                                                                                                                                                                                                                                 | Number of Scheduler Maps Supported for Each Port Group |
|------------------------------------|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| EX2200-48T and EX2200-48P switches | 2                     | <ul style="list-style-type: none"> <li>Ports 0–23 and SFP uplink ports 0 and 1 form a port group.</li> <li>Ports 24–47 and SFP uplink ports 2 and 3 form a port group.</li> </ul>                                                                                                     | 5                                                      |
| EX3200-24T and EX3200-24P switches | 1                     | <ul style="list-style-type: none"> <li>Ports 0–23 and the uplink ports form a port group.</li> </ul> <p><b>NOTE:</b> Uplink ports include 2 SFP+ or XFP uplink ports, or 4 SFP uplink ports.</p>                                                                                      | 4                                                      |
| EX3200-24T and EX3200-24P switches | 1                     | <ul style="list-style-type: none"> <li>Ports 0–23 and the uplink ports form a port group.</li> </ul> <p><b>NOTE:</b> Uplink ports include 2 SFP+ or XFP uplink ports or 4 SFP uplink ports.</p>                                                                                       | 4                                                      |
| EX3200-48T and EX3200-48P switches | 2                     | <ul style="list-style-type: none"> <li>Ports 0–23 and 1 SFP+ or XFP uplink port or 4 SFP uplink ports form a port group.</li> <li>Ports 24–47 and 1 SFP+ or XFP uplink port form a port group.</li> </ul>                                                                             | 4                                                      |
| EX4200-48T and EX4200-48P switches | 3                     | <ul style="list-style-type: none"> <li>Ports 0–23 form a port group.</li> <li>Ports 24–47 form a port group.</li> <li>2 SFP+ or XFP uplink ports or 4 SFP uplink ports form a port group.</li> </ul>                                                                                  | 4                                                      |
| EX4200-24T and EX4200-24P switches | 2                     | <ul style="list-style-type: none"> <li>Ports 0–23 form a port group.</li> <li>2 SFP+ or XFP uplink ports or 4 SFP uplink ports form a port group.</li> </ul>                                                                                                                          | 4                                                      |
| EX4300-24T and EX4300-24P switches | 1                     | <ul style="list-style-type: none"> <li>Ports 0–23 ports, 4 uplink ports, and 4 ports on the rear panel form a port group.</li> </ul> <p><b>NOTE:</b> Uplink ports in the front panel contains SFP or SFP+ ports 0–3, and uplink ports in the rear panel contains QSFP+ ports 0–3.</p> | 64                                                     |



Table 198: Support for Scheduler Maps on Switches and Line Cards (*continued*)

| Switch/Line Card                                                    | Number of Port Groups | Port Grouping Details                                                                                                                                                                                                                                                                                                | Number of Scheduler Maps Supported for Each Port Group |
|---------------------------------------------------------------------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| EX4300-48T and EX4300-48P switches                                  | 1                     | <ul style="list-style-type: none"> <li>Ports 0–47, 4 uplink ports, and 4 ports on the real panel form a port group.</li> </ul> <p><b>NOTE:</b> Uplink ports in the front panel contains SFP or SFP+ ports 0–3, and uplink ports in the rear panel contains QSFP+ ports 0–3.</p>                                      | 64                                                     |
| EX4500-40F switch                                                   | 2                     | <ul style="list-style-type: none"> <li>SFP or SFP+ ports 0–19 and the first SFP or SFP+ port 0–4 form a port group.</li> <li>SFP or SFP+ ports 20–39 and the second SFP or SFP+ uplink port 0–4 form a port group.</li> </ul>                                                                                        | 4                                                      |
| EX4550-32F switch                                                   | 1                     | <ul style="list-style-type: none"> <li>SFP or SFP+ ports 0–31 and the uplink ports in the front and rear panels form a port group.</li> </ul> <p><b>NOTE:</b> Uplink ports in the front panel contains SFP, SFP+, or RJ-45 ports 0–7, and uplink ports in the rear panel contains SFP, SFP+, or RJ-45 ports 0–7.</p> | 5                                                      |
| EX6200-48T (48-port RJ-45) and EX6200-48P (48-port PoE+) line cards | 2                     | <ul style="list-style-type: none"> <li>Ports 0–23 form a port group.</li> <li>Ports 24–47 form a port group.</li> </ul>                                                                                                                                                                                              | 5                                                      |
| EX6200-SRE64-4XS                                                    | 1                     | SFP+ ports 0–3 form a port group.                                                                                                                                                                                                                                                                                    | 4                                                      |
| EX8200-8XS (8-port SFP+) line card                                  | 4                     | <ul style="list-style-type: none"> <li>SFP+ ports 0 and 1 form a port group.</li> <li>SFP+ ports 2 and 3 form a port group.</li> <li>SFP+ ports 4 and 5 form a port group.</li> <li>SFP+ ports 6 and 7 form a port group.</li> </ul>                                                                                 | 6                                                      |

Table 198: Support for Scheduler Maps on Switches and Line Cards (*continued*)

| Switch/Line Card                                                                      | Number of Port Groups | Port Grouping Details                                                                                                                                                                                                                                                                                                                                                                                        | Number of Scheduler Maps Supported for Each Port Group |
|---------------------------------------------------------------------------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| EX8200-40XS (40-port SFP+) line card                                                  | 8                     | <ul style="list-style-type: none"> <li>SFP+ ports 0–4 form a port group.</li> <li>SFP+ ports 5–9 form a port group.</li> <li>SFP+ ports 10–14 form a port group.</li> <li>SFP+ ports 15–19 form a port group.</li> <li>SFP+ ports 20–24 form a port group.</li> <li>SFP+ ports 25–29 form a port group.</li> <li>SFP+ ports 30–34 form a port group.</li> <li>SFP+ ports 35–39 form a port group.</li> </ul> | 6                                                      |
| EX8200-48-F (48-port SFP) and EX8200-48T (48-port RJ-45) line cards                   | 2                     | <ul style="list-style-type: none"> <li>SFP or RJ-45 ports 0–23 form a port group.</li> <li>SFP or RJ-45 ports 24–47 form a port group.</li> </ul>                                                                                                                                                                                                                                                            | 6                                                      |
| EX8200-2XS-40P (40-port PoE+ with 4-port SFP and 2-port SFP+) line card               | 3                     | <ul style="list-style-type: none"> <li>Ports 0–19 and SFP ports 0 and 1 form a port group.</li> <li>Ports 20–39 and SFP ports 2 and 3 form a port group.</li> </ul>                                                                                                                                                                                                                                          | 5                                                      |
|                                                                                       |                       | <ul style="list-style-type: none"> <li>2 SFP+ ports form a port group.</li> </ul>                                                                                                                                                                                                                                                                                                                            | 6                                                      |
| EX8200-2XS-40T (40-port RJ-45 with 4-port SFP and 2-port SFP+) line card              | 3                     | <ul style="list-style-type: none"> <li>Ports 0–19, and SFP ports 0 and 1 form a port group.</li> <li>Ports 20–39 and SFP ports 2 and 3 form a port group.</li> </ul>                                                                                                                                                                                                                                         | 5                                                      |
|                                                                                       |                       | <ul style="list-style-type: none"> <li>2 SFP+ ports form a port group.</li> </ul>                                                                                                                                                                                                                                                                                                                            | 6                                                      |
| EX8200-48PL (48-port PoE+ 20 Gbps) and EX8200-48TL (48-port RJ-45 20 Gbps) line cards | 2                     | <ul style="list-style-type: none"> <li>PoE+ or RJ-45 ports 0–23 form a port group.</li> <li>PoE+ or RJ-45 ports 24–47 form a port group.</li> </ul>                                                                                                                                                                                                                                                          | 5                                                      |

**Related Documentation**

- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Defining CoS Schedulers and Scheduler Maps \(CLI Procedure\) on page 2109](#)

- [Defining CoS Schedulers \(J-Web Procedure\) on page 2111](#)

## Understanding CoS Two-Color Marking

Networks police traffic by limiting the input or output transmission rate of a class of traffic on the basis of user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or classes of service.

Policers require you to apply limits to the traffic flow and set a consequence for packets that exceed these limits—usually a higher loss priority, so that packets exceeding the policer limits are discarded first.

Juniper Networks EX Series Ethernet Switches support a single-rate two-color marking type of policer, which is a simplified version of Single-Rate-Three-Color marking, defined in RFC 2697, *A Single Rate Three Color Marker*. This type of policer meters traffic based on the configured committed information rate (CIR) and committed burst size (CBS).

The single-rate two-color marker meters traffic and marks incoming packets depending on whether they are smaller than the committed burst size (CBS)—marked green—or exceed it—marked red.

The single-rate two-color marking policer operates in color-blind mode. In this mode, the policer's actions are not affected by any previous marking or metering of the examined packets. In other words, the policer is “blind” to any previous coloring a packet might have had.

### Related Documentation

- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)
- [Understanding the Use of Policers in Firewall Filters on page 4767](#)
- [Configuring Policers to Control Traffic Rates \(CLI Procedure\) on page 4818](#)

## Understanding CoS Rewrite Rules

As packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. This topic describes how to use rewrite rules to alter the CoS settings. It covers:

This topic covers:

- [How Rewrite Rules Work on page 2067](#)
- [Default Rewrite Rule on page 2068](#)

### How Rewrite Rules Work

Rewrite rules set the value of the CoS bits within a packet's header. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table, and writes this CoS value into the packet header. For rewrites to occur, rewrite rules must be explicitly assigned to an interface.

On EX Series switches, you can define rewrite rules for IPv4 and IPv6 traffic to network interfaces, aggregated Ethernet interfaces (also known as link aggregation groups (LAGs)), routed VLAN interfaces (RVIs), Layer 3 interfaces, and Layer 3 VLAN-tagged sub-interfaces. Multiple rewrite rules of different types can be assigned to a single interface.

On EX4300 switches, you cannot configure separate DSCPv4 and DSCPv6 rewrite rules on network interfaces, aggregated Ethernet interfaces, Layer 3 interfaces, and integrated routing and bridging (IRB) interfaces. If you configure a DSCPv4 rewrite rule on an interface to rewrite IPv4 traffic, then the same rewrite rule is applied to IPv6 traffic also on that interface, and vice versa. You can define only DSCPv4 rewrite rules on integrated routing and bridging (IRB) interfaces and Layer 3 VLAN-tagged logical interfaces.

In effect, the rewrite rule performs the reverse function of the behavior aggregate (BA) classifier, which is used when the packet enters the switch. As the packet leaves the switch, the final CoS action is generally the application of a rewrite rule.

You configure rewrite rules to alter CoS values in outgoing packets on the outbound interfaces of an edge switch to meet the policies of a targeted peer. This allows the downstream switch in a neighboring network to classify each packet into the appropriate service group.



**NOTE:** When an IP precedence rewrite rule is active, bits 3, 4, and 5 of the type-of-service (ToS) byte are always reset to zero when code points are rewritten.

---

### Default Rewrite Rule

To define a rewrite rule on an interface, you can either create your own rewrite rule and enable it on the interface or enable a default rewrite rule. See [“Defining CoS Rewrite Rules \(CLI Procedure\)” on page 2120](#).

[Table 199 on page 2069](#) shows the default rewrite-rule mappings. These are based on the default bit definitions of Differentiated Services code point (DSCP), IEEE 802.1p, and IP precedence values and the default forwarding classes. You can configure multiple CoS rewrite rules for DSCP, IP precedence and IEEE 802.1p.



**NOTE:** By default, rewrite rules are not assigned to an interface. You must explicitly assign a user-defined or system-defined rewrite rule to an interface for the rewrites to occur.

When the CoS values of a packet match the forwarding class and packet-loss-priority (PLP) values, the switch rewrites markings on the packet based on the rewrite table.

Table 199: Default Packet Header Rewrite Mappings

| Map from Forwarding Class | PLP Value | Map to DSCP/IEEE 802.1p/IP Precedence Value |
|---------------------------|-----------|---------------------------------------------|
| expedited-forwarding      | low       | ef                                          |
| expedited-forwarding      | high      | ef                                          |
| assured-forwarding        | low       | af11                                        |
| assured-forwarding        | high      | af12 (DSCP)                                 |
| best-effort               | low       | be                                          |
| best-effort               | high      | be                                          |
| network-control           | low       | nc1/cs6                                     |
| network-control           | high      | nc2/cs7                                     |

#### Related Documentation

- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Defining CoS Rewrite Rules \(CLI Procedure\) on page 2120](#)
- [Defining CoS Rewrite Rules \(J-Web Procedure\) on page 2121](#)

## Understanding Port Shaping and Queue Shaping for CoS on EX Series Switches

When the amount of traffic on a switch's network exceeds the maximum bandwidth, packets are lost because of congestion in the network. The excess traffic in the network must be handled carefully to ensure minimum or no data loss in the network. A class-of-service (CoS) configuration includes several parameters that classify traffic into different queues and also define packet loss priorities (PLPs) to ensure smooth transmission of data in the network. You can use these configuration parameters to control or shape traffic for a specific port on a switch or for a specific CoS queue. While port shaping defines the maximum bandwidth allocated to an interface, queue shaping defines a limit on excess-bandwidth usage for each queue.

This topic covers:

- [Port Shaping on page 2069](#)
- [Queue Shaping on page 2070](#)

### Port Shaping

Port shaping enables you to shape the aggregate traffic through a port or channel to a rate that is less than the line rate. You can configure interfaces to shape traffic based on the rate-limited bandwidth of the total interface bandwidth. This allows you to shape the output of the interface so that the interface transmits less traffic than it is capable

of transmitting. For port shaping, you can specify shaping rate as the peak rate at which traffic can pass through the interface. You can specify rate as a value in bits per second (bps) either as a decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000) and the value can range from 1000 through 160,000,000,000 bps.

By default, shaping is not configured on an interface. If you do not configure a shaping rate on an interface, the default shaping rate is 100 percent, which is the equivalent to no shaping configured for that interface.

On EX Series switches except EX4300 switches, when you configure a shaping rate on an aggregated Ethernet (ae) interface, all members of the ae interface are shaped at the configured shaping rate. For example, consider an interface ae0 that consists of three interfaces: ge-0/0/0, ge-0/0/1, and ge-0/0/2. If a shaping rate of X Mbps is configured on ae0, traffic at the rate of X Mbps flows through each of the three interfaces. Therefore, the total traffic flowing through ae0 would be at the rate of 3X Mbps. On EX4300 switches, when you configure a shaping rate on an ae interface, the traffic is equally divided among the members of the ae interface.

---

### Queue Shaping

Queue shaping throttles the rate at which queues transmit packets. For example, using queue shaping, you can rate-limit a strict-priority queue so that the strict-priority queue does not lock out (or starve) low-priority queues. Similarly, for any queue, you can configure queue shaping.

You can specify queue shaping as the maximum rate at which traffic can pass through the queue or as a percentage of the available bandwidth. On EX Series switches except EX4300 switches, you can specify the rate as a value between 3200 and 160,000,000,000 bps and the percentage as a value from 0 to 100 percent. On EX4300 switches, you can specify the rate as a value between 8000 and 160,000,000,000 bps and the percentage as a value from 0 to 100 percent.

#### Related Documentation

- [Understanding CoS Schedulers on page 2060](#)
- [Defining CoS Schedulers and Scheduler Maps \(CLI Procedure\) on page 2109](#)

## Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches

Junos operating system (Junos OS) EZQoS on Juniper Networks EX Series Ethernet Switches eliminates the complexities involved in configuring class of service (CoS) across the network. EZQoS offers templates for key traffic classes.

Junos OS CoS allows you to divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. You can use CoS to ensure that different types of traffic (voice, video, and data) get the bandwidth and consideration they need to meet user expectations and business objectives.

Configuring CoS requires careful consideration of your service needs and thorough planning and design to ensure consistency across all switches in a CoS domain. To configure CoS manually, you must define and fine-tune all CoS components such as classifiers, rewrite rules, forwarding classes, schedulers, and scheduler-maps and then apply these

components to the interfaces. Therefore, configuring CoS can be a fairly complex and time-consuming task.

EZQoS works by automatically assigning preconfigured values to all CoS parameters based on the typical application requirements. These preconfigured values are stored in a template with a unique name. You can change the preconfigured values of these parameters to suit your particular application needs.

For using EZQoS, you must identify which switch ports are being used for a specific application (such as VoIP, video, and data) and manually apply the corresponding application-specific EZQoS template to these switch ports.



**NOTE:** Currently, we provide an EZQoS template for configuring CoS for VoIP.



**NOTE:** We recommend that you do not use the term EZQoS for defining a classifier.

#### Related Documentation

- [Junos OS CoS for EX Series Switches Overview on page 2042](#)
- [Configuring Junos OS EZQoS for CoS \(CLI Procedure\) on page 2125](#)

## Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports

Some line cards available for Juniper Networks EX8200 Ethernet Switches include oversubscribed ports that are combined in logical port groups that share bandwidth. These oversubscribed ports handle traffic differently than ports that provide continuous line-rate bandwidth. You might need to configure CoS queues differently for oversubscribed ports than for line-rate ports.

This topic describes:

- [Oversubscribed Ports on Line Cards on page 2071](#)
- [EX8200 Line Cards That Include Oversubscribed Ports on page 2072](#)
- [Ingress Queueing on page 2072](#)
- [Egress Queues on page 2073](#)

### Oversubscribed Ports on Line Cards

Oversubscribed ports on a line card are grouped into logical port groups. A port group collectively supports a certain bandwidth.

An EX8200 switch supports different line cards that provide line-rate and oversubscribed ports. Based on your requirement, you can choose the appropriate line card for an EX8200 switch. Line cards are field-replaceable units (FRUs) that can be installed in the line card slots in an EX8200 switch. In a line-rate EX8200 line card, each port in the line card supports the same amount of bandwidth and a single port can utilize that complete bandwidth. In an oversubscribed line card, a group of ports collectively support a certain

total bandwidth and each port in that group can use either a portion or all of the available bandwidth. However, the total utilization of bandwidth by the ports in the group cannot exceed the bandwidth available for that group.

Because the port groups share bandwidth, class-of-service (CoS) ingress and egress queues are handled differently for these shared-bandwidth ports in logical port groups than they are for ports that individually support line-rate bandwidth. Some EX8200 line cards combine both port types, those that share bandwidth across port groups and those that individually support line-rate bandwidth.

### EX8200 Line Cards That Include Oversubscribed Ports

Table 200 on page 2072 lists EX8200 line cards that include oversubscribed ports in logical port groups.

**Table 200: EX8200 Line Cards That Include Oversubscribed Ports**

| Line Card Model | Name                                          | Number of Oversubscribed Ports/Port Connector                                                                                                                                                     |
|-----------------|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EX8200-40XS     | 40-port SFP+                                  | 40 oversubscribed 10-gigabit SFP+ ports                                                                                                                                                           |
| EX8200-2XS-40P  | 40-port PoE+ with 4-port SFP and 2-port SFP+  | 40 oversubscribed 10/100/1000 Gigabit Ethernet ports with RJ-45 connectors, four small form-factor pluggable (SFP) ports (in which you can install 1-gigabit SFP transceivers) and two SFP+ ports |
| EX8200-2XS-40T  | 40-port RJ-45 with 4-port SFP and 2-port SFP+ | 40 oversubscribed 10/100/1000 Gigabit Ethernet ports with RJ-45 connectors, four SFP ports (in which you can install 1-gigabit small form-factor pluggable (SFP) transceivers) and two SFP+ ports |
| EX8200-48PL     | 48-port PoE+ 20 Gbps                          | 48 oversubscribed 10/100/1000 Gigabit Ethernet ports with RJ-45 connectors                                                                                                                        |
| EX8200-48TL     | 48-port RJ-45 20 Gbps                         | 48 oversubscribed 10/100/1000 Gigabit Ethernet ports with RJ-45 connectors                                                                                                                        |

### Ingress Queueing

Classification of packets occurs in two phases for the oversubscribed ports in the port groups.

- [Preclassification of Packets and Port Ingress Queueing on page 2072](#)
- [Full Classification of Packets and Fabric Ingress Queueing on page 2073](#)

#### ***Preclassification of Packets and Port Ingress Queueing***

Packets entering ports are forwarded to one of the ingress queues. The ingress queues schedule traffic from ports into the Packet Forwarding Engine.

The ingress queues are:



- Low-priority queue—Each interface in the line card has one low-priority queue. Traffic on these queues is scheduled using the shaped deficit weighted round-robin (SDWRR) algorithm, with each interface's queue having equal weight. On EX4300 switches, traffic is queued using the weighted deficit round-robin (WDRR) algorithm.
- High-priority queue—A set of interfaces in the line card shares a single high-priority queue. Traffic on this queue is scheduled by strict-high priority. The switch always sends critical network control packets on the high-priority queue.
- Line-rate priority queue—The packets entering line-rate ports are forwarded to this queue. Traffic on this queue is scheduled by strict priority and is always given higher priority than the traffic on the high-priority queue. This queue is used only in the following oversubscribed lines cards for an EX8200 switch:
  - EX8200-2XS-40P
  - EX8200-2XS-40T

For the purpose of port ingress queuing on oversubscribed ports, packets are classified only by behavior aggregate (BA) classification. To control the ingress queue (high priority or low priority) to which packets are sent, configure a BA classifier on the physical port and specify switch fabric priorities for the forwarding classes. On EX8200 switches, fabric priority determines the priority of packets ingressing the switch fabric. For the EX8200-40XS line card, fabric priority also determines the priority of packets ingressing the port group.

By default, the fabric priority for all forwarding classes is low. To direct packets belonging to a forwarding class to the high-priority ingress queue, set the fabric priority to high for that class.

Critical network-control packets and line-rate packets are handled differently from other packets. Instead of using the BA classifier to classify critical network-control packets, the switch always sends critical network packets to the high-priority queue. The line-rate packets are always sent to the line-rate priority queue. This difference in handling of network-control packets and line-rate packets ensures that these packets are not dropped because of congestion on the shared-bandwidth ports.

### ***Full Classification of Packets and Fabric Ingress Queuing***

When packets (apart from line-rate and critical network-control packets) from an oversubscribed port reach the Packet Forwarding Engine, it performs full packet classification, along with other actions, such as multifield (MF) classification, traffic policing, and storm control. It then schedules and queues the packets for ingressing the fabric. The fabric priority associated with the forwarding class determines whether packets are sent to the low priority or high-priority ingress queues.

### **Egress Queues**

On EX Series switches except EX4300 switches, each interface supports eight egress CoS queues. You can map up to 16 forwarding classes to these queues. An EX4300 switch interface supports 12 egress CoS queues.

In the EX8200-40XS line card, all interfaces in a port group share a single set of eight egress queues at the Packet Forwarding Engine. Egress traffic is fanned out from the Packet Forwarding Engine queues to the corresponding queues for the individual ports. For this reason, the interfaces in a port group must share the same scheduler map configuration. If you configure different scheduler map configurations for the different interfaces in a port group, an error is logged in the system log and the default scheduler map is used for all ports in the port group.

**Related  
Documentation**

- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)
- [Understanding CoS Schedulers on page 2060](#)
- *Understanding CoS Forwarding Classes*
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards \(CLI Procedure\) on page 2126](#)

## CHAPTER 31

# Configuration

- [Configuration Examples on page 2075](#)
- [Configuration Tasks on page 2099](#)
- [Configuration Statements on page 2127](#)

### Configuration Examples

---

- [Example: Configuring CoS on EX Series Switches on page 2075](#)

#### Example: Configuring CoS on EX Series Switches

Configure class of service (CoS) on your switch to manage traffic so that when the network experiences congestion and delay, critical applications are protected. Using CoS, you can divide traffic on your switch into classes and provide various levels of throughput and packet loss. This is especially important for traffic that is sensitive to jitter and delay, such as voice traffic.

This example shows how to configure CoS on a single EX Series switch in the network.

- [Requirements on page 2075](#)
- [Overview and Topology on page 2075](#)
- [Configuration on page 2078](#)
- [Verification on page 2088](#)

#### Requirements

---

This example uses the following hardware and software components:


- EX3200 and EX4200 switches
- Junos OS Release 9.0 or later for EX Series switches

#### Overview and Topology

---

This example uses the topology shown in [Figure 26 on page 2076](#).

Figure 26: Topology for Configuring CoS



The topology for this configuration example consists of EX3200 and EX4200 switches at the access layer.

The EX Series access switch is configured to support VLAN membership. Interfaces **ge-0/0/0** and **ge-0/0/1** are assigned to the voice VLAN (**voice-vlan**) for two VoIP phones. Switch port **ge-0/0/2** is assigned to the camera VLAN (**camera-vlan**) for the surveillance camera. Switch ports **ge-0/0/3**, **ge-0/0/4**, **ge-0/0/5**, and **ge-0/0/6** are assigned to the server VLAN (**server-vlan**) for the servers hosting various applications such as those provided by Citrix, Microsoft, Oracle, and SAP.

Table 201 on page 2077 shows the VLAN configuration components.

Table 201: Configuration Components: VLANs

| VLAN Name   | VLAN ID | VLAN Subnet and Available IP Addresses                                                                      | VLAN Description                                      |
|-------------|---------|-------------------------------------------------------------------------------------------------------------|-------------------------------------------------------|
| voice-vlan  | 10      | 192.168.1.0/28<br>192.168.1.1 through 192.168.1.14<br><br>192.168.1.15 is the subnet's broadcast address.   | Voice VLAN used for employee VoIP communication.      |
| camera-vlan | 20      | 192.168.1.16/28<br>192.168.1.17 through 192.168.1.30<br><br>192.168.1.31 is the subnet's broadcast address. | VLAN for the surveillance cameras.                    |
| server-vlan | 30      | 192.168.1.32/28<br>192.168.1.33 through 192.168.1.46<br><br>192.168.1.47 is the subnet's broadcast address. | VLAN for the servers hosting enterprise applications. |

PoE-capable ports on EX Series switches support Power over Ethernet (PoE) to provide both network connectivity and power for VoIP telephones connecting to the ports. [Table 202 on page 2077](#) shows the switch interfaces that are assigned to the VLANs and the IP addresses for devices connected to the switch ports on a 48-port switch, all ports of which are PoE-capable.

Table 202: Configuration Components: Switch Ports on a 48-Port All-PoE Switch

| Interfaces                             | VLAN Membership | IP Addresses                            | Port Devices                                                                                    |
|----------------------------------------|-----------------|-----------------------------------------|-------------------------------------------------------------------------------------------------|
| ge-0/0/0, ge-0/0/1                     | voice-vlan      | 192.168.1.1/28 through 192.168.1.2/28   | Two VoIP telephones.                                                                            |
| ge-0/0/2                               | camera-vlan     | 192.168.1.17/28                         | Surveillance camera.                                                                            |
| ge-0/0/3, ge-0/0/4, ge-0/0/5, ge-0/0/6 | server-vlan     | 192.168.1.33/28 through 192.168.1.36/28 | Four servers hosting applications such as those provided by Citrix, Microsoft, Oracle, and SAP. |



**NOTE:** This example shows how to configure CoS on a standalone EX Series switch. This example does not consider across-the-network applications of CoS in which you might implement different configurations on ingress and egress switches to provide differentiated treatment to different classes across a set of nodes in a network.

## Configuration

**CLI Quick Configuration** To quickly configure CoS, copy the following commands and paste them into the switch terminal window:

```
[edit]
set class-of-service forwarding-classes class app queue-num 5
set class-of-service forwarding-classes class mail queue-num 1
set class-of-service forwarding-classes class db queue-num 2
set class-of-service forwarding-classes class erp queue-num 3
set class-of-service forwarding-classes class video queue-num 4
set class-of-service forwarding-classes class best-effort queue-num 0
set class-of-service forwarding-classes class voice queue-num 6
set class-of-service forwarding-classes class network-control queue-num 7
set firewall family ethernet-switching filter voip_class term voip from source-address 192.168.1.1/28
set firewall family ethernet-switching filter voip_class term voip from source-address 192.168.1.2/28
set firewall family ethernet-switching filter voip_class term voip from protocol udp
set firewall family ethernet-switching filter voip_class term voip from source-port 2698
set firewall family ethernet-switching filter voip_class term voip then forwarding-class voice
loss-priority low
set firewall family ethernet-switching filter voip_class term network_control from precedence
[net-control internet-control]
set firewall family ethernet-switching filter voip_class term network_control then forwarding-class
network-control loss-priority low
set firewall family ethernet-switching filter voip_class term best_effort_traffic then
forwarding-class best-effort loss-priority low
set interfaces ge-0/0/0 description phone1-voip-ingress-port
set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input voip_class
set class-of-service interfaces ge-0/0/0 shaping-rate 100m
set interfaces ge-0/0/1 description phone2-voip-ingress-port
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input voip_class
set firewall family ethernet-switching filter video_class term video from source-address
192.168.1.17/28
set firewall family ethernet-switching filter video_class term video from protocol udp
set firewall family ethernet-switching filter video_class term video from source-port 2979
set firewall family ethernet-switching filter video_class term video then forwarding-class video
loss-priority low
set firewall family ethernet-switching filter video_class term network_control from precedence
[net-control internet-control]
set firewall family ethernet-switching filter video_class term network_control then forwarding-class
network-control loss-priority low
set firewall family ethernet-switching filter video_class term best_effort_traffic then
forwarding-class best-effort loss-priority low
set interfaces ge-0/0/2 description video-ingress-port
set interfaces ge-0/0/2 unit 0 family ethernet-switching filter input video_class
set firewall family ethernet-switching filter app_class term app from source-address
192.168.1.33/28
set firewall family ethernet-switching filter app_class term app from protocol tcp
set firewall family ethernet-switching filter app_class term app from source-port [1494 2512 2513
2598 2897]
set firewall family ethernet-switching filter app_class term app then forwarding-class app
loss-priority low
set firewall family ethernet-switching filter app_class term mail from source-address
192.168.1.34/28
set firewall family ethernet-switching filter app_class term mail from protocol tcp
set firewall family ethernet-switching filter app_class term mail from source-port [25 143 389
691 993 3268 3269]
set firewall family ethernet-switching filter app_class term mail then forwarding-class mail
loss-priority low
```

```

set firewall family ethernet-switching filter app_class term db from source-address 192.168.1.35/28
set firewall family ethernet-switching filter app_class term db from protocol tcp
set firewall family ethernet-switching filter app_class term db from source-port [1521 1525 1527
1571 1810 2481]
set firewall family ethernet-switching filter app_class term db then forwarding-class db loss-priority
low
set firewall family ethernet-switching filter app_class term erp from source-address 192.168.1.36/28
set firewall family ethernet-switching filter app_class term erp from protocol tcp
set firewall family ethernet-switching filter app_class term erp from source-port [3200 3300
3301 3600]
set firewall family ethernet-switching filter app_class term erp then forwarding-class erp
loss-priority low
set firewall family ethernet-switching filter app_class term network_control from precedence
[net-control internet-control]
set firewall family ethernet-switching filter app_class term network_control then forwarding-class
network-control loss-priority low
set firewall family ethernet-switching filter app_class term best_effort_traffic then forwarding-class
best-effort loss-priority low
set interfaces ge-0/0/3 unit 0 family ethernet-switching filter input app_class
set interfaces ge-0/0/4 unit 0 family ethernet-switching filter input app_class
set interfaces ge-0/0/5 unit 0 family ethernet-switching filter input app_class
set interfaces ge-0/0/6 unit 0 family ethernet-switching filter input app_class
set class-of-service schedulers voice-sched buffer-size percent 10
set class-of-service schedulers voice-sched priority strict-high
set class-of-service schedulers voice-sched transmit-rate percent 10
set class-of-service schedulers video-sched buffer-size percent 15
set class-of-service schedulers video-sched priority low
set class-of-service schedulers video-sched transmit-rate percent 15
set class-of-service schedulers app-sched buffer-size percent 10
set class-of-service schedulers app-sched priority low
set class-of-service schedulers app-sched transmit-rate percent 10
set class-of-service schedulers mail-sched buffer-size percent 5
set class-of-service schedulers mail-sched priority low
set class-of-service schedulers mail-sched transmit-rate percent 5
set class-of-service schedulers db-sched buffer-size percent 10
set class-of-service schedulers db-sched priority low
set class-of-service schedulers db-sched transmit-rate percent 10
set class-of-service schedulers erp-sched buffer-size percent 10
set class-of-service schedulers erp-sched priority low
set class-of-service schedulers erp-sched transmit-rate percent 10
set class-of-service schedulers nc-sched buffer-size percent 5
set class-of-service schedulers nc-sched priority strict-high
set class-of-service schedulers nc-sched transmit-rate percent 5
set class-of-service schedulers be-sched buffer-size percent 35
set class-of-service schedulers be-sched priority low
set class-of-service schedulers be-sched transmit-rate percent 35
set class-of-service scheduler-maps ethernet-cos-map forwarding-class voice scheduler
voice-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class video scheduler
video-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class app scheduler app-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class mail scheduler mail-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class db scheduler db-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class erp scheduler erp-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class network-control
scheduler nc-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class best-effort scheduler
be-sched
set class-of-service interfaces ge-0/0/20 scheduler-map ethernet-cos-map
set class-of-service schedulers voice-sched-queue-shap shaping-rate 30m

```

```

set class-of-service scheduler-maps sched-map-be forwarding-class best-effort scheduler
voice-sched-queue-shap
set class-of-service interfaces ge-0/0/2 scheduler-map sched-map-be

```

### Step-by-Step Procedure

To configure and apply CoS:

1. Configure one-to-one mappings between eight forwarding classes and eight queues:

```

[edit class-of-service]
user@switch# set forwarding-classes class app queue-num 5
user@switch# set forwarding-classes class mail queue-num 1
user@switch# set forwarding-classes class db queue-num 2
user@switch# set forwarding-classes class erp queue-num 3
user@switch# set forwarding-classes class video queue-num 4
user@switch# set forwarding-classes class best-effort queue-num 0
user@switch# set forwarding-classes class voice queue-num 6
user@switch# set forwarding-classes class network-control queue-num 7

```

2. Define the firewall filter **voip\_class** to classify the VoIP traffic:

```

[edit firewall]
user@switch# set family ethernet-switching filter voip_class

```

3. Define the term **voip**:

```

[edit firewall]
user@switch# set family ethernet-switching filter voip_class term voip from source-address
192.168.1.1/28
user@switch# set family ethernet-switching filter voip_class term voip from source-address
192.168.1.2/28
user@switch# set family ethernet-switching filter voip_class term voip protocol udp
user@switch# set family ethernet-switching filter voip_class term voip source-port 2698
user@switch# set family ethernet-switching filter voip_class term voip then forwarding-class
voice loss-priority low

```

4. Define the term **network\_control** (for the **voip\_class** filter):

```

[edit firewall]
user@switch# set family ethernet-switching filter voip_class term network_control from
precedence [net-control internet-control]
user@switch# set family ethernet-switching filter voip_class term network_control then
forwarding-class network-control loss-priority low

```

5. Define the term **best\_effort\_traffic** with no match conditions (for the **voip\_class** filter):

```

[edit firewall]
user@switch# set family ethernet-switching filter voip_class term best_effort_traffic then
forwarding-class best-effort loss-priority low

```

6. Apply the firewall filter **voip\_class** as an input filter to the interfaces for the VoIP phones:

```

[edit interfaces]
user@switch# set ge-0/0/0 description phone1-voip-ingress-port
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input voip_class
user@switch# set ge-0/0/1 description phone2-voip-ingress-port
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input voip_class

```

7. Apply port shaping on the interface **ge-0/0/0**:

```

[edit]
user@switch# set class-of-service interfaces ge-0/0/0 shaping-rate 100m

```

8. Define the firewall filter **video\_class** to classify the video traffic:

```

[edit firewall]
user@switch# set family ethernet-switching filter video_class

```



9. Define the term **video**:

```
[edit firewall]
user@switch# set family ethernet-switching filter video_class term video from
source-address 192.168.1.17/28
user@switch# set family ethernet-switching filter video_class term video protocol udp
user@switch# set family ethernet-switching filter video_class term video source-port 2979
user@switch# set family ethernet-switching filter video_class term video then
forwarding-class video loss-priority low
```

10. Define the term **network\_control** (for the **video\_class** filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter video_class term network_control from
precedence [net-control internet-control]
user@switch# set family ethernet-switching filter video_class term network_control then
forwarding-class network-control loss-priority low
```

11. Define the term **best\_effort\_traffic** with no match conditions (for the **video\_class** filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter video_class term best_effort_traffic then
forwarding-class best-effort loss-priority low
```

12. Apply the firewall filter **video\_class** as an input filter to the interface for the surveillance camera:

```
[edit interfaces]
user@switch# set ge-0/0/2 description video-ingress-port
user@switch# set ge-0/0/2 unit 0 family ethernet-switching filter input video_class
```

13. Define the firewall filter **app\_class** to classify the application server traffic:

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class
```

14. Define the term **app** (for the **app\_class** filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term app from source-address
192.168.1.33/28
user@switch# set family ethernet-switching filter app_class term app protocol tcp
user@switch# set family ethernet-switching filter app_class term app source-port [1494
2512 2513 2598 2897]
user@switch# set family ethernet-switching filter app_class term app then forwarding-class
app loss-priority low
```

15. Define the term **mail** (for the **app\_class** filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term mail from source-address
192.168.1.34/28
user@switch# set family ethernet-switching filter app_class term mail protocol tcp
user@switch# set family ethernet-switching filter app_class term mail source-port [25 143
389 691 993 3268 3269]
user@switch# set family ethernet-switching filter app_class term mail then forwarding-class
mail loss-priority low
```

16. Define the term **db** (for the **app\_class** filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term db from source-address
192.168.1.35/28
user@switch# set family ethernet-switching filter app_class term db protocol tcp
user@switch# set family ethernet-switching filter app_class term db source-port [1521
1525 1527 1571 1810 2481]
```

- ```

user@switch# set family ethernet-switching filter app_class term db then forwarding-class
db loss-priority low

```
17. Define the term **erp** (for the **app\_class** filter):
 

```

[edit firewall]
user@switch# set family ethernet-switching filter app_class term erp from source-address
192.168.1.36/28
user@switch# set family ethernet-switching filter app_class term erp protocol tcp
user@switch# set family ethernet-switching filter app_class term erp source-port [3200
3300 3301 3600]
user@switch# set family ethernet-switching filter app_class term erp then forwarding-class
erp loss-priority low

```
  18. Define the term **network\_control** (for the **app\_class** filter):
 

```

[edit firewall]
user@switch# set family ethernet-switching filter app_class term network_control from
precedence [net-control internet-control]
user@switch# set family ethernet-switching filter app_class term network_control then
forwarding-class network-control loss-priority low

```
  19. Define the term **best\_effort\_traffic** (for the **app\_class** filter):
 

```

[edit firewall]
user@switch# set family ethernet-switching filter app_class term best_effort_traffic then
forwarding-class best-effort loss-priority low

```
  20. Apply the firewall filter **app\_class** as an input filter to the interfaces for the servers hosting applications:
 

```

[edit interfaces]
user@switch# set ge-0/0/3 unit 0 family ethernet-switching filter input app_class
user@switch# set ge-0/0/4 unit 0 family ethernet-switching filter input app_class
user@switch# set ge-0/0/5 unit 0 family ethernet-switching filter input app_class
user@switch# set ge-0/0/6 unit 0 family ethernet-switching filter input app_class

```
  21. Configure schedulers:
 

```

[edit class-of-service]
user@switch# set schedulers voice-sched buffer-size percent 10
user@switch# set schedulers voice-sched priority strict-high
user@switch# set schedulers voice-sched transmit-rate percent 10
user@switch# set schedulers video-sched buffer-size percent 15
user@switch# set schedulers video-sched priority low
user@switch# set schedulers video-sched transmit-rate percent 15
user@switch# set schedulers app-sched buffer-size percent 10
user@switch# set schedulers app-sched priority low
user@switch# set schedulers app-sched transmit-rate percent 10
user@switch# set schedulers mail-sched buffer-size percent 5
user@switch# set schedulers mail-sched priority low
user@switch# set schedulers mail-sched transmit-rate percent 5
user@switch# set schedulers db-sched buffer-size percent 10
user@switch# set schedulers db-sched priority low
user@switch# set schedulers db-sched transmit-rate percent 10
user@switch# set schedulers erp-sched buffer-size percent 10
user@switch# set schedulers erp-sched priority low
user@switch# set schedulers erp-sched transmit-rate percent 10
user@switch# set schedulers nc-sched buffer-size percent 5
user@switch# set schedulers nc-sched priority strict-high
user@switch# set schedulers nc-sched transmit-rate percent 5
user@switch# set schedulers be-sched buffer-size percent 35
user@switch# set schedulers be-sched priority low
user@switch# set schedulers be-sched transmit-rate percent 35

```

22. Assign the forwarding classes to schedulers with the scheduler map **ethernet-cos-map**:
 

```
[edit class-of-service]
user@switch# set scheduler-maps ethernet-cos-map forwarding-class voice scheduler
voice-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class video scheduler
video-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class app scheduler
app-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class mail scheduler
mail-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class db scheduler
db-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class erp scheduler
erp-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class network-control
scheduler nc-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class best-effort scheduler
be-sched
```
23. Associate the scheduler map with the outgoing interface:
 

```
[edit class-of-service interfaces]
user@switch# set ge-0/0/20 scheduler-map ethernet-cos-map
```
24. Apply queue shaping for the best-effort queue:
 

```
[edit]
user@switch# set class-of-service schedulers voice-sched-queue-shap shaping-rate 30m
user@switch# set class-of-service scheduler-maps sched-map-be forwarding-class
best-effort scheduler voice-sched-queue-shap
user@switch# set class-of-service interfaces ge-0/0/2 scheduler-map sched-map-be
```

**Results** Display the results of the configuration:

```
user@switch> show firewall

firewall family ethernet-switching {
  filter voip_class {
    term voip {
      from {
        source-address {
          192.168.1.1/28;
          192.168.1.2/28;
        }
        protocol udp;
        source-port 2698;
      }
      then {
        forwarding-class voice;
        loss-priority low;
      }
    }
  }
  term network control {
    from {
      precedence [net-control internet-control];
    }
    then {
      forwarding-class network-control;
      loss-priority low;
    }
  }
}
```

```
    }
  }
  term best_effort_traffic {
    then {
      forwarding-class best-effort;
      loss-priority low;
    }
  }
}
filter video_class {
  term video {
    from {
      source-address {
        192.168.1.17/28;
      }
      protocol udp;
      source-port 2979;
    }
    then {
      forwarding-class video;
      loss-priority low;
    }
  }
  term network_control {
    from {
      precedence [net-control internet-control];
    }
    then {
      forwarding-class network-control;
      loss-priority low;
    }
  }
  term best_effort_traffic {
    then {
      forwarding-class best-effort;
      loss-priority low;
    }
  }
}
filter app_class {
  term app {
    from {
      source-address {
        192.168.1.33/28;
      }
      protocol tcp;
      source-port [1491 2512 2513 2598 2897];
    }
    then {
      forwarding-class app;
      loss-priority low;
    }
  }
  term mail {
    from {
      source-address {
```

```

        192.168.1.34/28;
    }
    protocol tcp;
    source-port [25 143 389 691 993 3268 3269];
}
then {
    forwarding-class mail;
    loss-priority low;
}
}
term db {
    from {
        source-address {
            192.168.1.35/28;
        }
        protocol tcp;
        source-port [1521 1525 1527 1571 1810 2481];
    }
    then {
        forwarding-class db;
        loss-priority low;
    }
}
term erp {
    from {
        source-address {
            192.168.1.36/28;
        }
        protocol tcp;
        source-port [3200 3300 3301 3600];
    }
    then {
        forwarding-class erp;
        loss-priority low;
    }
}
term network control {
    from {
        precedence [net-control internet-control];
    }
    then {
        forwarding-class network-control;
        loss-priority low;
    }
}
term best_effort_traffic {
    then {
        forwarding-class best-effort;
        loss-priority low;
    }
}
}
}
}

user@switch# show class-of-service
forwarding-classes {

```

```
class app queue-num 5;
class mail queue-num 1;
class db queue-num 2;
class erp queue-num 3;
class video queue-num 4;
class best-effort queue-num 0;
class voice queue-num 6;
class network-control queue-num 7;
}
interfaces {
  ge-0/0/0 {
    shaping-rate 100m;
  }
}
interfaces {
  ge-0/0/2 {
    scheduler-map sched-map-be;
  }
}
schedulers {
  voice-sched-queue-shap {
    shaping-rate 30m;
  }
  voice-sched {
    buffer-size percent 10;
    priority strict-high;
    transmit-rate percent 10;
  }
  video-sched {
    buffer-size percent 15;
    priority low;
    transmit-rate percent 15;
  }
  app-sched {
    buffer-size percent 10;
    priority low;
    transmit-rate percent 10;
  }
  mail-sched {
    buffer-size percent 5;
    priority low;
    transmit-rate percent 5;
  }
  db-sched {
    buffer-size percent 10;
    priority low;
    transmit-rate percent 10;
  }
  erp-sched {
    buffer-size percent 10;
    priority low;
    transmit-rate percent 10;
  }
  nc-sched {
    buffer-size percent 5;
    priority strict-high;
  }
}
```

```

        transmit-rate percent 5;
    }
    be-sched {
        buffer-size percent 35;
        priority low;
        transmit-rate percent 35;
    }
}
scheduler-maps {
    ethernet-cos-map {
        forwarding-class voice scheduler voice-sched;
        forwarding-class video scheduler video-sched;
        forwarding-class app scheduler app-sched;
        forwarding-class mail scheduler mail-sched;
        forwarding-class db scheduler db-sched;
        forwarding-class erp scheduler erp-sched;
        forwarding-class network-control scheduler nc-sched;
        forwarding-class best-effort scheduler be-sched;
    }
    sched-map-be {
        forwarding-class best-effort scheduler voice-sched-queue-shap;
    }
}

user@switch# show interfaces

ge-0/0/0 {
    unit 0 {
        family ethernet {
            filter {
                input voip_class;
            }
        }
    }
}
ge-0/0/1 {
    unit 0 {
        family ethernet {
            filter {
                input voip_class;
            }
        }
    }
}
ge-0/0/2 {
    unit 0 {
        family ethernet {
            filter {
                input video_class;
            }
        }
    }
}
ge-0/0/3 {
    unit 0 {
        family ethernet {
            filter {

```

```
        input app_class;
      }
    }
  }
}
ge-0/0/4 {
  unit 0 {
    family ethernet {
      filter {
        input app_class;
      }
    }
  }
}
ge-0/0/5 {
  unit 0 {
    family ethernet {
      filter {
        input app_class;
      }
    }
  }
}
ge-0/0/6 {
  unit 0 {
    family ethernet {
      filter {
        input app_class;
      }
    }
  }
}
```

---

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Defined Forwarding Classes Exist and Are Mapped to Queues on page 2088](#)
- [Verifying That the Forwarding Classes Have Been Assigned to Schedulers on page 2089](#)
- [Verifying That the Scheduler Map Has Been Applied to the Interface on page 2091](#)
- [Verifying That Port Shaping Has Been Applied on page 2091](#)
- [Verifying That Queue Shaping Has Been Applied on page 2095](#)

### *Verifying That the Defined Forwarding Classes Exist and Are Mapped to Queues*

**Purpose** Verify that the forwarding classes **app**, **best-effort**, **db**, **erp**, **mail**, **network-control**, **video**, and **voice** have been defined and mapped to queues.



**Action** user@switch> show class-of-service forwarding-class

Forwarding class	ID	Queue
app	0	5
db	1	2
erp	2	3
best-effort	3	0
mail	4	1
voice	5	6
video	6	4
network-control	7	7

**Meaning** This output shows that the forwarding classes have been defined and mapped to appropriate queues.

***Verifying That the Forwarding Classes Have Been Assigned to Schedulers***

**Purpose** Verify that the forwarding classes have been assigned to schedulers.

**Action** user@switch> **show class-of-service scheduler-map**

Scheduler map: ethernet-cos-map, Index: 2

Scheduler: voice-sched, Forwarding class: voice, Index: 22

Transmit rate: 5 percent, Rate Limit: none, Buffer size: 15 percent,

Priority: Strict-high

Drop profiles:

Loss priority	Protocol	Index	Name
High	non-TCP	1	<default-drop-profile>
High	TCP	1	<default-drop-profile>

Scheduler: video-sched, Forwarding class: video, Index: 22

Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent,

Priority: low

Drop profiles:

Loss priority	Protocol	Index	Name
High	non-TCP	1	<default-drop-profile>
High	TCP	1	<default-drop-profile>

Scheduler: app-sched, Forwarding class: app, Index: 22

Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent,

Priority: low

Drop profiles:

Loss priority	Protocol	Index	Name
High	non-TCP	1	<default-drop-profile>
High	TCP	1	<default-drop-profile>

Scheduler: mail-sched, Forwarding class: mail, Index: 22

Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,

Priority: low

Drop profiles:

Loss priority	Protocol	Index	Name
High	non-TCP	1	<default-drop-profile>
High	TCP	1	<default-drop-profile>

Scheduler: db-sched, Forwarding class: db, Index: 22

Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent,

Priority: low

Drop profiles:

Loss priority	Protocol	Index	Name
High	non-TCP	1	<default-drop-profile>
High	TCP	1	<default-drop-profile>

Scheduler: erp-sched, Forwarding class: erp, Index: 22

Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent,

Priority: low

Drop profiles:

Loss priority	Protocol	Index	Name
High	non-TCP	1	<default-drop-profile>
High	TCP	1	<default-drop-profile>

Scheduler: be-sched, Forwarding class: best-effort, Index: 20

Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent,

Priority: low

Drop profiles:

Loss priority	Protocol	Index	Name
High	non-TCP	1	<default-drop-profile>
High	TCP	1	<default-drop-profile>

Scheduler: nc-sched, Forwarding class: network-control, Index: 22

Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,

```

Priority: Strict-high
Drop profiles:
  Loss priority  Protocol  Index  Name
  High          non-TCP    1      <default-drop-profile>
  High          TCP        1      <default-drop-profile>

```

**Meaning** This output shows that the forwarding classes have been assigned to schedulers.

### *Verifying That the Scheduler Map Has Been Applied to the Interface*

**Purpose** Verify that the scheduler map has been applied to the interface.

```

Action  user@switch> show class-of-service interface
...
Physical interface: ge-0/0/20, Index: 149
Queues supported: 8, Queues in use: 8
Scheduler map: ethernet-cos-map, Index: 43366
Input scheduler map: <default>, Index: 3
...

```

**Meaning** This output shows that the scheduler map (**ethernet-cos-map**) has been applied to the interface (**ge-0/0/20**).

### *Verifying That Port Shaping Has Been Applied*

**Purpose** Verify that the port shaping has been applied to an interface.

**Action** Following is the output before port shaping is applied to the interface **ge-0/0/0**, when there is egress traffic of 400 Mbps exiting on that interface:

```

user@switch> show interfaces ge-0/0/0 extensive
Physical interface: ge-0/0/0, Enabled, Physical link is Up
Interface index: 239, SNMP ifIndex: 548, Generation: 242
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto, BPDU Error:
None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online,
Media type: Copper
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:23:9c:0b:ae:8d, Hardware address: 00:23:9c:0b:ae:8d
Last flapped   : 2012-07-07 03:21:52 UTC (1d 18:02 ago)
Statistics last cleared: 2012-07-07 23:54:34 UTC (21:29:59 ago)
Traffic statistics:
Input bytes   :                0                0 bps
Output bytes  :          2299853696          345934816 bps
Input packets :                0                0 pps
Output packets:          17967609          337827 pps
IPv6 transit statistics:
Input bytes   :                0
Output bytes  :                0
Input packets :                0
Output packets:                0

```

```

Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort      0          18302337          0
  1 assured-forw      0          0          0
  5 expedited-fo      0          0          0
  7 network-cont      0          0          0

Queue number:      Mapped forwarding classes
  0      best-effort
  1      assured-forwarding
  5      expedited-forwarding
  7      network-control

Active alarms : None
Active defects : None
MAC statistics:
  Total octets      Receive      Transmit
  Total packets      0      2299853696
  Unicast packets      0      17967609
  Broadcast packets      0      0
  Multicast packets      0      0
  CRC/Align errors      0      0
  FIFO errors      0      0
  MAC control frames      0      0
  MAC pause frames      0      0
  Oversized frames      0
  Jabber frames      0
  Fragment frames      0
  Code violations      0

Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link
partner Speed: 1000 Mbps
  Local resolution:
    Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 1
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority  Limit

      %      bps      %      usec
  0 best-effort      95      950000000      95      NA      low      none
  7 network-control      5      50000000      5      NA      low      none

Interface transmit statistics: Disabled

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2

```

```

Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol eth-switch, Generation: 163, Route table: 0
Flags: Trunk-Mode

```

The Traffic statistics: field in this output shows that egress traffic is ~400 Mbps (345,934,816 bps). When a port shaping of 100 Mbps is applied to the ge-0/0/0 interface, you see the following outputs for the **show interfaces ge-0/0/0 statistics** and the **show class-of-service interface ge-0/0/0** commands:

```

user@switch> show interfaces ge-0/0/0 statistics
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 239, SNMP ifIndex: 548, Generation: 242
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto, BPDU Error:
None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online,
  Media type: Copper
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags : None
  CoS queues : 8 supported, 8 maximum usable queues
  Hold-times : Up 0 ms, Down 0 ms
  Current address: 00:23:9c:0b:ae:8d, Hardware address: 00:23:9c:0b:ae:8d
  Last flapped : 2012-07-07 03:21:52 UTC (1d 18:10 ago)
  Statistics last cleared: 2012-07-07 23:54:34 UTC (21:37:58 ago)
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 15779512832 100223104 bps
  Input packets: 0 0 pps
  Output packets: 123277444 97874 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

  0 best-effort 0 123350092 57012484

  1 assured-forw 0 0 0

```

```

5 expedited-fo          0          0          0

7 network-cont          0          0          0

Queue number:           Mapped forwarding classes
0                       best-effort
1                       assured-forwarding
5                       expedited-forwarding
7                       network-control

Active alarms : None
Active defects : None
MAC statistics:
Total octets             Receive      Transmit
Total packets           0      15779512832
Unicast packets         0      123277444
Broadcast packets      0      0
Multicast packets      0      0
CRC/Align errors       0      0
FIFO errors            0      0
MAC control frames     0      0
MAC pause frames       0      0
Oversized frames      0
Jabber frames         0
Fragment frames       0
Code violations        0

Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link
partner Speed: 1000 Mbps
Local resolution:
Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 1
CoS information:
Direction : Output
CoS transmit queue      Bandwidth      Buffer Priority
Limit
%      bps      %      usec
0 best-effort          95      95000000      95      NA      low
none
7 network-control      5      5000000      5      NA      low
none

Interface transmit statistics: Disabled

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0      0 bps
Output bytes : 0      0 bps

```

```

Input packets:          0          0 pps
Output packets:         0          0 pps
Protocol eth-switch, Generation: 163, Route table: 0
Flags: Trunk-Mode

```

```

user@switch> show class-of-service interface ge-0/0/0
Physical interface: ge-0/0/0, Index: 165
Queues supported: 8, Queues in use: 4
Shaping rate: 100000000 bps
...
...

```

**Meaning** In the output for the **show interfaces ge-0/0/0 statistics** command, the Traffic statistics: field shows that egress traffic is ~100 Mbps (100,223,104 bps). The output for the **show class-of-service interface ge-0/0/0** command shows that the shaping rate is 100,000,000 bps, which indicates that a port shaping of 100 Mbps is applied to the ge-0/0/0 interface.

### *Verifying That Queue Shaping Has Been Applied*

**Purpose** Verify that the queue shaping has been applied to the best-effort queue.

**Action** Following is the output before queue shaping is applied to the best-effort queue when there is egress traffic of 400 Mbps exiting on that interface:

```

user@switch> show interfaces ge-0/0/2 extensive
Physical interface: ge-0/0/2, Enabled, Physical link is Up
Interface index: 239, SNMP ifIndex: 548, Generation: 242
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto, BPDU Error:
None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online,
Media type: Copper
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags    : None
CoS queues    : 8 supported, 8 maximum usable queues
Hold-times    : Up 0 ms, Down 0 ms
Current address: 00:23:9c:0b:ae:8d, Hardware address: 00:23:9c:0b:ae:8d
Last flapped   : 2012-07-07 03:21:52 UTC (1d 18:02 ago)
Statistics last cleared: 2012-07-07 23:54:34 UTC (21:29:59 ago)
Traffic statistics:
Input bytes   :          0          0 bps
Output bytes  :    2299853696    345934816 bps
Input packets :          0          0 pps
Output packets:    17967609    337827 pps
IPv6 transit statistics:
Input bytes   :          0
Output bytes  :          0
Input packets :          0
Output packets:          0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

```

```

0 best-effort          0          18302337          0
1 assured-forw         0          0          0
5 expedited-fo        0          0          0
7 network-cont        0          0          0

Queue number:          Mapped forwarding classes
0                      best-effort
1                      assured-forwarding
5                      expedited-forwarding
7                      network-control

Active alarms : None
Active defects : None
MAC statistics:
Total octets           Receive          Transmit
Total packets          0          2299853696
Unicast packets        0          17967609
Broadcast packets      0          0
Multicast packets      0          0
CRC/Align errors       0          0
FIFO errors            0          0
MAC control frames     0          0
MAC pause frames       0          0
Oversized frames       0
Jabber frames          0
Fragment frames        0
Code violations         0

Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link
partner Speed: 1000 Mbps
Local resolution:
Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 1
CoS information:
Direction : Output
CoS transmit queue      Bandwidth          Buffer Priority
Limit
%          bps      %          usec      low
0 best-effort          95          950000000      95          NA
none
7 network-control      5           50000000       5           NA
none
Interface transmit statistics: Disabled

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes :          0
Output bytes :          0
Input packets:          0
Output packets:          0
Local statistics:
Input bytes :          0
Output bytes :          0
Input packets:          0

```



```

Output packets:                0
Transit statistics:
Input bytes :                   0                0 bps
Output bytes :                  0                0 bps
Input packets:                  0                0 pps
Output packets:                 0                0 pps
Protocol eth-switch, Generation: 163, Route table: 0
Flags: Trunk-Mode

```

The Traffic statistics: field in this output shows that the egress traffic is ~400 Mbps (345,934,816 bps). When a queue shaping of 30 Mbps is applied to the best-effort queue, you see the following output for the **show interfaces ge-0/0/2 statistics** and **show class-of-service scheduler-map sched-map-be** commands:

```

user@switch> show interfaces ge-0/0/2 statistics
Physical interface: ge-0/0/2, Enabled, Physical link is Up
  Interface index: 239, SNMP ifIndex: 548, Generation: 242
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto, BPDU Error:
None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online,
Media type: Copper
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:23:9c:0b:ae:8d, Hardware address: 00:23:9c:0b:ae:8d
Last flapped   : 2012-07-07 03:21:52 UTC (1d 18:29 ago)
Statistics last cleared: 2012-07-08 21:46:22 UTC (00:04:56 ago)
Traffic statistics:
Input bytes :                   0                0 bps
Output bytes :                5376128896          30097712 bps
Input packets:                   0                0 pps
Output packets:              42001003          29392 pps
IPv6 transit statistics:
Input bytes :                   0
Output bytes :                   0
Input packets:                   0
Output packets:                  0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort                0          41986978          57813642

1 assured-forw                0              0              0

5 expedited-fo                0              0              0

7 network-cont                0              0              0

Queue number:      Mapped forwarding classes
0                  best-effort
1                  assured-forwarding
5                  expedited-forwarding

```

```

7                               network-control
Active alarms : None
Active defects : None
MAC statistics:
Total octets          Receive      Transmit
Total packets        0          5376128896
Unicast packets      0          42001003
Broadcast packets    0              0
Multicast packets    0              0
CRC/Align errors     0              0
FIFO errors          0              0
MAC control frames   0              0
MAC pause frames     0              0
Oversized frames     0
Jabber frames        0
Fragment frames      0
Code violations       0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: Symmetric, Remote fault: OK, Link
partner Speed: 1000 Mbps
Local resolution:
Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 1
CoS information:
Direction : Output
CoS transmit queue   Bandwidth      Buffer Priority
Limit
0 best-effort        %      bps      %      usec
none                  r      r      r      NA      low
Interface transmit statistics: Disabled

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 638) (Generation 138)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0          0 bps
Output bytes : 0          0 bps
Input packets: 0          0 pps
Output packets: 0          0 pps
Protocol eth-switch, Generation: 163, Route table: 0
Flags: Trunk-Mode

```

```

user@switch> show class-of-service scheduler-map sched-map-be

```

```

Scheduler map: sched-map-be, Index: 31271

```

```

Scheduler: voice-sched-queue-shap, Forwarding class: best-effort, Index: 64106

```

```

Transmit rate: remainder, Rate Limit: none, Buffer size: remainder,
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Shaping rate: 30000000 bps
Drop profiles:
  Loss priority  Protocol  Index  Name
  High          non-TCP    1      <default-drop-profile>
  High          TCP       1      <default-drop-profile>

```

**Meaning** In the output for the **show interfaces ge-0/0/2 statistics** command, the Traffic statistics: field shows that the egress traffic is ~30 Mbps (30,097,712 bps). The output for the **show class-of-service scheduler-map sched-map-be** command, shows that a shaping rate of 30,000,000 bps (that is 30 Mbps) is applied to the best-effort queue.

- Related Documentation**
- [Defining CoS Code-Point Aliases \(CLI Procedure\) on page 2101](#)
  - [Defining CoS Classifiers \(CLI Procedure\) on page 2103](#)
  - [Defining CoS Forwarding Classes \(CLI Procedure\) on page 2107](#)
  - [Defining CoS Schedulers and Scheduler Maps \(CLI Procedure\) on page 2109](#)
  - [Configuring CoS Tail Drop Profiles \(CLI Procedure\)](#)
  - [Assigning CoS Components to Interfaces \(CLI Procedure\) on page 2123](#)
  - [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)

## Configuration Tasks

- [Configuring CoS \(J-Web Procedure\) on page 2100](#)
- [Defining CoS Code-Point Aliases \(CLI Procedure\) on page 2101](#)
- [Defining CoS Code-Point Aliases \(J-Web Procedure\) on page 2101](#)
- [Defining CoS Classifiers \(CLI Procedure\) on page 2103](#)
- [Defining CoS Classifiers \(J-Web Procedure\) on page 2105](#)
- [Defining CoS Forwarding Classes \(CLI Procedure\) on page 2107](#)
- [Defining CoS Forwarding Classes \(J-Web Procedure\) on page 2107](#)
- [Defining CoS Schedulers and Scheduler Maps \(CLI Procedure\) on page 2109](#)
- [Defining CoS Schedulers \(J-Web Procedure\) on page 2111](#)
- [Defining CoS Scheduler Maps \(J-Web Procedure\) on page 2114](#)
- [Configuring CoS Congestion Management \(CLI Procedure\) on page 2115](#)
- [Defining CoS Drop Profiles \(J-Web Procedure\) on page 2118](#)
- [Defining CoS Rewrite Rules \(CLI Procedure\) on page 2120](#)
- [Defining CoS Rewrite Rules \(J-Web Procedure\) on page 2121](#)
- [Assigning CoS Components to Interfaces \(CLI Procedure\) on page 2123](#)
- [Assigning CoS Components to Interfaces \(J-Web Procedure\) on page 2123](#)

- [Configuring Junos OS EZQoS for CoS \(CLI Procedure\) on page 2125](#)
- [Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards \(CLI Procedure\) on page 2126](#)

## Configuring CoS (J-Web Procedure)

The Class of Service Configuration pages allow you to configure the Junos CoS components. You can configure forwarding classes for transmitting packets, define which packets are placed into each output queue, and schedule the transmission service level for each queue. After defining the CoS components you must assign classifiers to the required physical and logical interfaces.

Using the Class of Service Configuration pages, you can configure various CoS components individually or in combination to define particular CoS services.

To configure CoS components :

1. In the J-Web interface, select **Configure>Class of Service**.
2. On the Class of Service Configuration page, select one of the following options depending on the CoS component that you want to define. Enter information into the pages as described in the respective table:
  - To define or edit CoS value aliases, select **CoS Value Aliases** .
  - To define or edit forwarding classes and assign queues, select **Forwarding Classes**.
  - To define or edit classifiers, select **Classifiers** .
  - To define or edit rewrite rules, select **Rewrite Rules**.
  - To define or edit schedulers, select **Schedulers**.
  - To define or edit virtual channel groups, select **Interface Associations**.
3. Click **Apply** after completing configuration on any Configuration page.

### Related Documentation

- [Defining CoS Classifiers \(J-Web Procedure\) on page 2105](#)
- [Defining CoS Code-Point Aliases \(J-Web Procedure\) on page 2101](#)
- [Defining CoS Forwarding Classes \(J-Web Procedure\) on page 2107](#)
- [Defining CoS Rewrite Rules \(J-Web Procedure\) on page 2121](#)
- [Defining CoS Schedulers \(J-Web Procedure\) on page 2111](#)
- [Assigning CoS Components to Interfaces \(J-Web Procedure\) on page 2123](#)

## Defining CoS Code-Point Aliases (CLI Procedure)

You can use code-point aliases to streamline the process of configuring CoS features on your EX Series switch. A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

You can configure code-point aliases for the following CoS marker types:

- **dscp** and **dscp-ipv6**—Handles incoming IPv4 and IPv6 packets, respectively.
- **ieee-802.1**—Handles Layer 2 CoS.
- **inet-precedence**—Handles incoming IPv4 packets. IP precedence mapping requires only the higher order three bits of the DSCP field.

To configure a code-point alias for a specified CoS marker type (**dscp**), assign an alias (**my1**) to the code-point (**110001**):

```
[edit class-of-service code-point-aliases]
user@switch# set dscp my1 110001
```

The **my1** alias will be applicable for incoming IPv4 packets.

### Related Documentation

- [Defining CoS Code-Point Aliases \(J-Web Procedure\) on page 2101](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Monitoring CoS Value Aliases on page 2164](#)
- [Understanding CoS Code-Point Aliases on page 2046](#)

## Defining CoS Code-Point Aliases (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You can use the J-Web interface to define CoS code-point aliases on an EX Series switch. By defining aliases you can assign meaningful names to a particular set of bit values and refer to them when configuring CoS components.

To define CoS code-point aliases:

1. Select **Configure > Class of Service > CoS Value Aliases**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—Adds a code-point alias. Enter information into the code point alias page as described in [Table 203 on page 2102](#).
- **Edit**—Modifies an existing code-point alias. Enter information into the code point alias page as described in [Table 203 on page 2102](#).
- **Delete**—Deletes an existing code-point alias.

[Table 203 on page 2102](#) describes the related fields.

**Table 203: CoS Value Aliases Configuration Fields**

Field	Function	Your Action
Code point name	Specifies the name for a code-point—for example, <b>af11</b> or <b>be</b> .	Enter a name.
Code point type	Specifies a code-point type. The code-point type can be DSCP or IP precedence.	Select a value.
Code point value bits	Specifies the CoS value for which an alias is defined.  Changing this value alters the behavior of all classifiers that refer to this alias.	To specify a CoS value, type it in the appropriate format: <ul style="list-style-type: none"> <li>• For DSCP CoS values, use the format <b>xxxxxx</b>, where x is 1 or 0—for example, <b>101110</b>.</li> <li>• For IP precedence CoS values, use the format <b>xxx</b>, where x is 1 or 0—for example, <b>111</b>.</li> </ul>

**Related Documentation**

- [Defining CoS Code-Point Aliases \(CLI Procedure\) on page 2101](#)
- [Monitoring CoS Value Aliases on page 2164](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)

## Defining CoS Classifiers (CLI Procedure)

Packet classification associates incoming packets with a particular CoS servicing level. Classifiers associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. Junos OS supports two general types of classifiers:

- Behavior aggregate (BA) classifier—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value, IP precedence value, or IEEE 802.1p value. EX Series switches except EX4300 switches support two types of loss priorities: **high** and **low**. EX4300 switches support four types of loss priorities: **high**, **medium-high**, **low**, and **medium-low**.

You can configure BA classifiers for the following CoS marker types:

- **dscp** and **dscp-ipv6**—Handles incoming IPv4 and IPv6 packets, respectively.
- **ieee-802.1**—Handles Layer 2 CoS.
- **inet-precedence**—Handles incoming IPv4 packets. IP precedence mapping requires only the higher order three bits of the DSCP field.
- Multifield (MF) classifier—Examine multiple fields in the packet such as source and destination addresses and source and destination port numbers of the packet. With MF classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.



**NOTE:** Juniper Networks EX8200 Ethernet Switches implement the on-demand ternary content addressable memory (TCAM) allocation of memory so that when additional TCAM space is required for CoS, the space is allocated from the free TCAM space or from the unused TCAM space. An error log message is generated when you configure CoS classifiers beyond the available TCAM space that includes both the free and unused space.

The following example describes how to configure a BA classifier (**ba-classifier**) as the default DSCP map for handling IPv4 traffic and to apply the BA classifier to either a specific Gigabit Ethernet interface or to all the Gigabit Ethernet interfaces on the switch. The BA classifier assigns loss priorities, as shown in [Table 204 on page 2103](#), to incoming packets in the four forwarding classes.

You can use the same procedure to set MF classifiers (except that you would use firewall filter rules).

**Table 204: BA-classifier Loss Priority Assignments**

Forwarding Class	For CoS Traffic Type	ba-classifier Assignment
<b>be</b>	Best-effort traffic	High-priority code point: <b>000001</b>

Table 204: BA-classifier Loss Priority Assignments (*continued*)

<b>ef</b>	Expedited-forwarding traffic	High-priority code point: <b>101110</b>
<b>af</b>	Assured-forwarding traffic	High-priority code point: <b>001100</b>
<b>nc</b>	Network-control traffic	High-priority code point: <b>110001</b>

To configure a DSCP BA classifier named **ba-classifier** as the default DSCP map:

- Associate code point **000001** with forwarding class **be** and loss priority **high**:

```
[edit class-of-service classifiers]
user@switch# set dscp ba-classifier import default forwarding-class be loss-priority high
code-points 000001
```

- Associate code point **101110** with forwarding class **ef** and loss priority **high**:

```
[edit class-of-service classifiers]
user@switch# set dscp ba-classifier forwarding-class ef loss-priority high code-points 101110
```

- Associate code point **001100** with forwarding class **af** and loss priority **high**:

```
[edit class-of-service classifiers]
user@switch# set dscp ba-classifier forwarding-class af loss-priority high code-points 001100
```

- Associate code point **110001** with forwarding class **nc** and loss priority **high**:

```
[edit class-of-service classifiers]
user@switch# set dscp ba-classifier forwarding-class nc loss-priority high code-points 110001
```

- Apply the classifier to a specific interface or to all Gigabit Ethernet interfaces on the switch.

- To apply the classifier to a specific interface:

```
[edit class-of-service interfaces]
user@switch# set ge-0/0/0 unit 0 classifiers dscp ba-classifier
```

- To apply the classifier to all Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and the logical-interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set ge-* unit * classifiers dscp ba-classifier
```



**NOTE:** On EX8200 switches, it can take a long time to install code-point classifiers on multiple interfaces (for example, approximately 25 minutes to install 64 code-point classifiers on multiple interfaces in the order of 280 or more).

#### Related Documentation

- [Defining CoS Classifiers \(J-Web Procedure\) on page 2105](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Assigning CoS Components to Interfaces \(CLI Procedure\) on page 2123](#)
- [Monitoring CoS Classifiers on page 2157](#)
- [Understanding CoS Classifiers](#)



- [Troubleshooting a CoS Classifier Configuration for a TCAM Space Error on page 2192](#)

### Defining CoS Classifiers (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You can use the J-Web interface to define CoS classifiers on an EX Series switch. Classifiers examine the CoS value or alias of an incoming packet and assign the packet a level of service by setting its forwarding class and loss priority.

To define CoS classifiers:

1. Select **Configure > Class of Service > Classifiers**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:
  - **Add**—Adds a classifier. Enter information into the classifier page as described in [Table 205 on page 2105](#).
  - **Edit**—Modifies an existing classifier. Enter information into the classifier page as described in [Table 205 on page 2105](#).
  - **Delete**—Deletes an existing classifier.

**Table 205: Classifiers Configuration Fields**

Field	Function	Your Action
Classifier Name	Specifies the name for a classifier.	To name a classifier, type the name—for example, <b>ba-classifier</b> .
Classifier Type	Specifies the type of classifier: <b>dscp</b> , <b>ieee-802.1</b> , or <b>inet-precedence</b> .	Select a value from the list.

Table 205: Classifiers Configuration Fields (*continued*)

Field	Function	Your Action
Code Point Mapping	Sets the forwarding classes and the packet loss priorities (PLPs) for specific CoS values and aliases.	<p>To add a code point mapping:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b>.</li> <li>2. Select the code point.</li> <li>3. Select a forwarding class from the following list: <ul style="list-style-type: none"> <li>• <b>expedited-forwarding</b>—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service. Packets can be forwarded out of sequence or dropped.</li> <li>• <b>best-effort</b>—Provides no special CoS handling of packets. Typically, RED drop profile is aggressive and no loss priority is defined.</li> <li>• <b>assured-forwarding</b>—Provides high assurance for packets within the specified service profile. Excess packets are dropped.</li> <li>• <b>network-control</b>—Packets can be delayed but not dropped.</li> </ul> </li> <li>4. Select the loss priority. <p>To assign a loss priority, select one:</p> <ul style="list-style-type: none"> <li>• <b>high</b>—Packet has a high loss priority.</li> <li>• <b>low</b>—Packet has a low loss priority.</li> </ul> </li> </ol>

- Related Documentation**
- [Defining CoS Classifiers \(CLI Procedure\) on page 2103](#)
  - [Example: Configuring CoS on EX Series Switches on page 2075](#)
  - [Monitoring CoS Classifiers on page 2157](#)
  - [Understanding CoS Classifiers](#)

## Defining CoS Forwarding Classes (CLI Procedure)

Forwarding classes allow you to group packets for transmission. Based on forwarding classes, you assign packets to output queues.

By default, four categories of forwarding classes are defined: best effort, assured forwarding, expedited forwarding, and network control. EX Series switches support up to 16 forwarding classes.

You can configure forwarding classes in one of the following ways:

- Using **class** statement—You can configure up to 16 forwarding classes and you can map multiple forwarding classes to single queue.
- Using **queue** statement—You can configure up to 8 forwarding classes and you can map one forwarding class to one queue.

This example uses the **class** statement to configure forwarding classes.

To configure CoS forwarding classes, map the forwarding classes to queues:

```
[edit class-of-service forwarding-classes]
user@switch# set class be queue-num 0
user@switch# set class ef queue-num 1
user@switch# set class af queue-num 2
user@switch# set class nc queue-num 3
user@switch# set class ef1 queue-num 4
user@switch# set class ef2 queue-num 5
user@switch# set class af1 queue-num 6
user@switch# set class nc1 queue-num 7
```

### Related Documentation

- [Defining CoS Forwarding Classes \(J-Web Procedure\) on page 2107](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic](#)
- [Assigning CoS Components to Interfaces \(CLI Procedure\) on page 2123](#)
- [Monitoring CoS Forwarding Classes on page 2158](#)
- [Understanding CoS Forwarding Classes](#)

## Defining CoS Forwarding Classes (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You can define CoS forwarding classes on an EX Series switch using the J-Web interface. Assigning a forwarding class to a queue number affects the scheduling and marking of a packet as it transits a switch.

To define forwarding classes:

1. Select **Configure > Class of Service > Forwarding Classes**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See *Using the Commit Options to Commit Configuration Changes* for details about all commit options.

2. Click one of the following options:
  - **Add**—Adds a forwarding class. Enter information into the forwarding class page as described in [Table 206 on page 2108](#).
  - **Edit**—Modifies an existing forwarding class. Enter information into the forwarding class page as described in [Table 206 on page 2108](#).
  - **Delete**—Deletes an existing forwarding class.

**Table 206: Forwarding Classes Configuration Fields**

Field	Function	Your Action
<b>Forwarding Class Summary</b>		
Queue #	<p>Specifies the internal queue numbers to which forwarding classes are assigned.</p> <p>By default, if a packet is not classified, it is assigned to the class associated with queue 0. You can have more than one forwarding class to a queue number.</p>	<p>To specify an internal queue number, select an integer from 0 through 7, appropriate for your platform.</p> <p><b>NOTE:</b> For EX4300 switches, to specify an internal queue number, select an integer from 0 through 11.</p>
Forwarding Class Name	<p>Specifies the forwarding class names assigned to specific internal queue numbers.</p> <p>By default, four forwarding classes are assigned to queue numbers 0 (best-effort), 1 (assured-forwarding), 5 (expedited-forwarding), and 7 (network-connect).</p> <p><b>NOTE:</b> For EX4300 switches, by default the forwarding classes are assigned to queue numbers 0 (best-effort), 1 (expedited-forwarding), 2 (assured-forwarding), 3 (network-connect), 8 (mcast-be), 9 (mcast-ef), 10 (mcast-af), and 11 (mcast-nc).</p>	Type the name—for example, <b>be-class</b> .

**Related Documentation**

- [Defining CoS Forwarding Classes \(CLI Procedure\) on page 2107](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic](#)
- [Monitoring CoS Forwarding Classes on page 2158](#)

- [Assigning CoS Components to Interfaces \(J-Web Procedure\) on page 2123](#)
- [Understanding CoS Forwarding Classes](#)

## Defining CoS Schedulers and Scheduler Maps (CLI Procedure)

You use schedulers to define the class-of-service (CoS) properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the drop profiles associated with the queue.

You associate the schedulers with forwarding classes by means of scheduler maps. You can then associate each scheduler map with an interface, thereby configuring the queues and packet schedulers that operate according to this mapping.



**NOTE:** On EX Series switches, you cannot configure a scheduler map on an individual interface that is a member of a link aggregation group (LAG). Instead, you must configure the scheduler map on the LAG itself (that is, on the aggregated Ethernet (ae) interface).

You can associate up to four user-defined scheduler maps with an interface.

This topic describes:

- [Configuring a Scheduler and a Scheduler Map on page 2109](#)
- [Assigning a Scheduler Map to Interfaces on page 2110](#)
- [Assigning Scheduler Maps to Interfaces on EX8200 Line Cards That Include Oversubscribed Ports on page 2110](#)

### Configuring a Scheduler and a Scheduler Map

You can define the properties for an output queue by configuring a scheduler. You can then define a scheduler map to associate a forwarding class with a scheduler.

To configure a scheduler and a scheduler map:

1. Create a scheduler, and assign one or more output queue properties to it:

**[edit class-of-service]**

```
user@switch# set schedulers scheduler-name output-queue-properties
```

For various properties that you can define for an output queue, see the [schedulers](#) hierarchy.

2. Configure a scheduler map that associates the scheduler with the forwarding class:

**[edit class-of-service]**

```
user@switch# set scheduler-maps map-name forwarding-class class-name scheduler scheduler-name
```

### Assigning a Scheduler Map to Interfaces

---

After defining a scheduler map, you can assign the scheduler map to one or more interfaces. You can also assign the scheduler map to multiple interfaces by using a wildcard representation of the interface or Virtual Chassis Ports (VCPs).

Following are sample syntaxes and examples for assigning a scheduler map to a single or to multiple interfaces:

- To assign the scheduler map to one interface:

```
[edit class-of-service interfaces]
user@switch# set interface-name scheduler-map map-name
```

- To assign the scheduler map to more than one interface, you can use a wildcard representation of the interface:

```
[edit class-of-service interfaces]
user@switch# set wild-card-representation-of-interface-name scheduler-map map-name
```

For example, following is the configuration to assign the **be-map** scheduler map to all Gigabit Ethernet interfaces (**ge-\***):

```
[edit class-of-service interfaces]
user@switch# set ge-* scheduler-map be-map
```

- To assign the scheduler map to all VCPs:

```
[edit class-of-service interfaces]
user@switch# set wild-card-representation-of-vcp scheduler-map map-name
```



**NOTE:** You can assign a scheduler map to a VCP only on EX4200, EX4300 or EX4500 switches that are members of Virtual Chassis composed exclusively either of EX4200 switches, EX4300 switches or of EX4500 switches, or that are members of a mixed Virtual Chassis composed of EX4200, EX4300, and EX4500 switches.

For example, following is the configuration to assign the **be-map** scheduler map to all VCPs:

```
[edit class-of-service interfaces]
user@switch# set vcp-* scheduler-map be-map
```

### Assigning Scheduler Maps to Interfaces on EX8200 Line Cards That Include Oversubscribed Ports

---

Some line cards available for Juniper Networks EX8200 Ethernet Switches include oversubscribed ports that are combined in logical port groups that share bandwidth. These oversubscribed ports handle traffic differently than ports that provide continuous line-rate bandwidth. You might need to configure CoS queues differently for oversubscribed ports than for line-rate ports. For more information about EX8200 line cards that include oversubscribed ports, see [“Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports” on page 2071](#).

For interfaces on oversubscribed EX8200 line cards, you use the same procedure to configure CoS schedulers as you do for other interfaces. However, you must assign the

same scheduler map to all the interfaces in a port group. When you assign a scheduler map to one interface in a port group, you do not need to assign the scheduler map to the remaining interfaces in the port group. The switch automatically uses that scheduler map for all the interfaces in the port group when you bring the interfaces up. Therefore, you do not need to assign the scheduler map to the remaining interfaces in that port group.

If you assign different scheduler maps to different interfaces in a port group, you do not receive an error when you commit the configuration. Instead, an error is logged in the system log. When you bring an interface in the port group up, the default scheduler map is assigned to all interfaces in the port group. If you assign a scheduler map to an interface that is down and if that scheduler map is different from the scheduler map being used by the currently operating interfaces in the port group, then the default scheduler map is used by all interfaces in the port group, even the currently operating ones, when you bring the interface up.

To assign a scheduler map to a port group, assign a scheduler map to at least one interface in the port group:

```
[edit class-of-service interfaces]
user@switch# set interface-name scheduler-map map-name
```

Considering that the xe-0/0/2 interface is part of a port group, following is the configuration to assign a scheduler map named **ef-map** to a port group that contains xe-0/0/2:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/2 scheduler-map ef-map
```

#### Related Documentation

- [Defining CoS Schedulers \(J-Web Procedure\) on page 2111](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Assigning CoS Components to Interfaces \(CLI Procedure\) on page 2123](#)
- [Monitoring CoS Scheduler Maps on page 2162](#)
- [Understanding CoS Schedulers on page 2060](#)

## Defining CoS Schedulers (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You can use the J-Web interface to define CoS schedulers on an EX Series switch. Using schedulers, you can assign attributes to queues and thereby provide congestion control for a particular class of traffic. These attributes include the amount of interface bandwidth, memory buffer size, transmit rate, and schedule priority.

To configure schedulers:

1. Select **Configure > Class of Service > Schedulers**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—Adds a scheduler. Enter information into the Schedulers page as described in [Table 207 on page 2112](#).
- **Edit**—Modifies an existing scheduler. Enter information into the Schedulers page as described in [Table 207 on page 2112](#).
- **Delete**—Deletes an existing scheduler.

**Table 207: Schedulers Configuration Page**

Field	Function	Your Action
Scheduler name	Specifies the name for a scheduler.	To name a scheduler, type the name—for example, <b>be-scheduler</b> .
Scheduling priority	<p>Sets the transmission priority of the scheduler, which determines the order in which an output interface transmits traffic from the queues.</p> <p>You can set the scheduling priority at different levels in the order of increasing priority from low to high.</p> <p>A high-priority queue with a high transmission rate might lock out lower-priority traffic.</p>	<p>To set a priority, select one:</p> <ul style="list-style-type: none"> <li>• <b>low</b>—Packets in this queue are transmitted last.</li> <li>• <b>strict-high</b>—Packets in this queue are transmitted first.</li> <li>• To specify no scheduling priority, select the blank check box.</li> </ul>



Table 207: Schedulers Configuration Page (*continued*)

Field	Function	Your Action
Buffer size	<p>Defines the size of the delay buffer.</p> <p>By default, queues 0 through 11 are allotted the following percentages of the total available buffer space:</p> <ul style="list-style-type: none"> <li>• Queue 0—75 percent</li> <li>• Queue 1—0 percent</li> <li>• Queue 2—0 percent</li> <li>• Queue 3—5 percent</li> <li>• Queue 4—0 percent</li> <li>• Queue 5—0 percent</li> <li>• Queue 6—0 percent</li> <li>• Queue 7—0 percent</li> <li>• Queue 8—15 percent</li> <li>• Queue 9—0 percent</li> <li>• Queue 10—0 percent</li> <li>• Queue 11—5 percent</li> </ul> <p><b>NOTE:</b> A large buffer size value correlates with a greater possibility of packet delays. Such a value might not be practical for sensitive traffic such as voice or video.</p>	<p>To define a delay buffer size for a scheduler, select the appropriate option:</p> <ul style="list-style-type: none"> <li>• To specify no buffer size, select the blank check box.</li> <li>• To specify buffer size as a percentage of the total buffer, select <b>Percent</b> and type an integer from 1 through 100.</li> <li>• To specify buffer size as the remaining available buffer, select <b>Remainder</b>.</li> </ul> <p><b>NOTE:</b> On EX8200 and EX4300 switches, you can specify the buffer size as a temporal value. The queuing algorithm will then drop packets after it has queued a computed number of bytes. This number is the product of the logical interface speed and the configured temporal value.</p>
Shaping rate	<p>Specifies the rate at which queues transmit packets.</p>	<ul style="list-style-type: none"> <li>• To specify shaping rate as a percentage, select <b>Percent</b> and type an integer from 1 through 100.</li> <li>• To specify shaping rate as a number, select <b>Rate</b> and enter a value.</li> <li>• To specify no shaping rate, select the blank check box.</li> </ul>

Table 207: Schedulers Configuration Page (*continued*)

Field	Function	Your Action
Transmit rate	<p>Defines the transmission rate of a scheduler.</p> <p>The transmit rate determines the traffic bandwidth from each forwarding class you configure.</p> <p>By default, queues 0 through 11 are allotted the following percentages of the transmission capacity:</p> <ul style="list-style-type: none"> <li>• Queue 0—75 percent</li> <li>• Queue 1—0 percent</li> <li>• Queue 2—0 percent</li> <li>• Queue 3—5 percent</li> <li>• Queue 4—0 percent</li> <li>• Queue 5—0 percent</li> <li>• Queue 6—0 percent</li> <li>• Queue 7—0 percent</li> <li>• Queue 8—15 percent</li> <li>• Queue 9—0 percent</li> <li>• Queue 10—0 percent</li> <li>• Queue 11—5 percent</li> </ul>	<p>To define a transmit rate, select the appropriate option:</p> <ul style="list-style-type: none"> <li>• To enforce the exact transmission rate, select <b>Rate</b> and enter a value.</li> <li>• To specify the remaining transmission capacity, select <b>Remainder Available</b>.</li> <li>• To specify a percentage of transmission capacity, select <b>Percent</b> and type an integer from 1 through 100.</li> <li>• To specify no transmit rate, select the blank check box.</li> </ul>
Excess rate  <b>NOTE:</b> This option is supported only on EX4300 switches.	<p>Defines the excess rate of a scheduler.</p>	<p>To define the excess rate, select the appropriate option:</p> <ul style="list-style-type: none"> <li>• To specify a percentage of the excess rate, select <b>Percent</b> and type an integer from 1 through 100.</li> <li>• To specify no excess rate, select the blank check box.</li> </ul>

**Related Documentation**

- [Defining CoS Schedulers and Scheduler Maps \(CLI Procedure\) on page 2109](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Monitoring CoS Scheduler Maps on page 2162](#)

## Defining CoS Scheduler Maps (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You can use the J-Web interface to configure CoS scheduler maps on an EX Series switch.



**NOTE:** On EX Series switches, you cannot configure a scheduler map on an individual interface that is a member of a link aggregation group (LAG). Instead, you must configure the scheduler map on the LAG itself (that is, on the aggregated Ethernet (ae) interface).

To configure scheduler maps:

1. Select **Configure > Class of Service > Scheduler Maps**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See *Using the Commit Options to Commit Configuration Changes* for details about all commit options.

2. Click one of the following options:
  - **Add**—Adds a scheduler map. Enter information into the scheduler map page as described in [Table 208 on page 2115](#).
  - **Edit**—Modifies an existing scheduler map. Enter information into the scheduler map page as described in [Table 208 on page 2115](#).
  - **Delete**—Deletes an existing scheduler map.

**Table 208: Scheduler Maps Configuration Fields**

Field	Function	Your Action
Scheduler Map Name	Specifies the name for a scheduler map.	To name a map, type the name—for example, <b>be-scheduler-map</b> .
Scheduler Mapping	<p>Allows you to associate a preconfigured scheduler with a forwarding class.</p> <p>After scheduler maps have been applied to an interface, they affect the hardware queues and packet schedulers.</p>	<p>To associate a scheduler with a forwarding class, locate the forwarding class and select the scheduler in the box next to it.</p> <p>For example, for the <b>best-effort</b> forwarding class, select the configured scheduler from the list.</p>

**Related Documentation**

- [Defining CoS Schedulers \(J-Web Procedure\) on page 2111](#)
- [Defining CoS Schedulers and Scheduler Maps \(CLI Procedure\) on page 2109](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Monitoring CoS Scheduler Maps on page 2162](#)

## Configuring CoS Congestion Management (CLI Procedure)

An effective congestion management mechanism is imperative to ensure smooth flow of traffic in a network and also to ensure minimum packet drops in the network. Class of

service (CoS) provides congestion management methods that allow you to define parameters based on which packets can be dropped when the output queue is full. These parameters vary depending on the EX Series switch that you are using in a network.

You can specify parameters for dropping packets at the **[edit class-of-service drop-profiles]** hierarchy level and reference them in a scheduler configuration. The parameters that you can specify are **fill-level** and **drop-probability**. The first parameter defines the delay-buffer bandwidth, which provides packet buffer space to absorb burst traffic up to the specified duration of delay. When the specified delay buffer becomes full, packets with 100 percent drop probability are dropped from the head of the buffer. The second parameter represents a percentage value that correlates to the likelihood that an individual packet is dropped from the network.

Depending on the switch on which you are configuring a drop profile, you can configure either a weighted tail drop (WTD) profile or a weighted random early detection (WRED) profile.

This topic describes:

- [Configuring a Weighted Tail Drop Profile on page 2116](#)
- [Configuring a Weighted Random Early Detection Drop Profile on page 2116](#)

---

### Configuring a Weighted Tail Drop Profile

A weighted tail drop (WTD) is a congestion management mechanism in which packets are dropped from the tail of the queue when the queue reaches a certain buffer capacity (that is, the fill level), and hence the name weighted tail drop. When that level is reached on EX2200, EX3200, or EX4200 Switches, packets marked with a packet loss priority (PLP) of high are prevented from entering the queue (that is, they are discarded).

To configure a WTD profile, create a drop profile name and assign a fill level:

```
[edit class-of-service drop-profiles]
user@switch# set profile-name fill-level percentage
```

Following is a sample WTD profile in which the fill level is set to 80 percent:

```
[edit class-of-service drop-profiles]
user@switch# set wtd-profile fill-level 80
```

---

### Configuring a Weighted Random Early Detection Drop Profile

A WRED drop profile enables you to define multiple data points for fill level and drop probability so that packets are dropped at various levels of queue fullness, and for various drop probabilities. Unlike the WTD drop profile that can be defined only for packets with a PLP of high, WRED can be defined for packets with a PLP of high and also for packets with a PLP of low.



**NOTE:** The WRED drop profile is supported only on EX4300 standalone switches, EX4300 Virtual Chassis, EX8200 standalone switches and EX8200 Virtual Chassis.

---

WRED has two implementations: segmented and interpolated. From a high level, segmented is a stair-step-like drop profile, whereas interpolated is a smoother (curve) drop profile. For a graphical representation of both these implementations, see [“Understanding CoS Congestion Management” on page 2055](#). Although the formation of graph lines is different for both these implementations, the application of the profile is the same. On EX Series switches except EX4300 switches, when a packet reaches the head of the queue, a random number between 0 and 100 is calculated. This random number is plotted against the drop profile using the current queue fullness of that particular queue. When the random number falls above the graph line, the packet is transmitted. When the number falls below the graph line, the packet is dropped from the network.

For information about congestion management on EX4300 switches, see [“Understanding CoS Congestion Management” on page 2055](#).



**NOTE:** On EX4300 switches, you cannot enable WRED on multidestination (multicast) queues. You can enable WRED only on unicast queues.

Following is the procedure to define a segmented and an interpolated drop profiles:

- To configure a segmented drop profile, specify multiple data points for fill level (l) and drop probability (p) as follows:

**[edit class-of-service drop-profiles]**

```
user@switch# set profile-name fill-level percentage-l1 drop-probability percentage-p1
user@switch# set profile-name fill-level percentage-l2 drop-probability percentage-p2
user@switch# set profile-name fill-level percentage-l3 drop-probability percentage-p3
user@switch# set profile-name fill-level percentage-l4 drop-probability percentage-p4
```

Following is a sample segmented drop profile:

**[edit class-of-service drop-profiles]**

```
user@switch# set seg-prof fill-level 20 drop-probability 25
user@switch# set seg-prof fill-level 40 drop-probability 50
user@switch# set seg-prof fill-level 60 drop-probability 75
user@switch# set seg-prof fill-level 80 drop-probability 100
```

- To configure an interpolated drop profile on EX Series switches except EX4300 switches, specify multiple data points for fill level (l) and drop probability (p) using the **interpolate** statement as follows:

**[edit class-of-service drop-profiles]**

```
user@switch# set profile-name interpolate fill-level percentage-l1 drop-probability
percentage-l1
user@switch# set profile-name interpolate fill-level percentage-l2 drop-probability
percentage-l2
user@switch# set profile-name interpolate fill-level percentage-l3 drop-probability
percentage-p3
user@switch# set profile-name interpolate fill-level percentage-l4 drop-probability
percentage-p4
```

Following is a sample interpolated drop profile:

**[edit class-of-service drop-profiles]**

```
user@switch# set inter-prof interpolate fill-level 20 drop-probability 25
user@switch# set inter-prof interpolate fill-level 40 drop-probability 50
user@switch# set inter-prof interpolate fill-level 60 drop-probability 75
```

```
user@switch# set inter-prof interpolate fill-level 80 drop-probability 100
```

- To configure an interpolated drop profile EX4300 switches, specify two data points for fill level (l) and drop probability (p) by using the **interpolate** statement as follows:

```
[edit class-of-service drop-profiles]
```

```
user@switch# set profile-name interpolate fill-level percentage-l1 fill-level percentage-l2  
drop-probability percentage-l1 percentage-l2
```

Following is a sample interpolated drop profile:

```
[edit class-of-service drop-profiles]
```

```
user@switch# set inter-prof interpolate fill-level 20 fill-level 80 drop-probability 25  
drop-probability 100
```

#### Related Documentation

- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Understanding CoS Congestion Management on page 2055](#)

## Defining CoS Drop Profiles (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You can use the J-Web interface to define CoS drop profiles on EX4500 and EX8200 switches.

To configure CoS drop profiles:

1. Select **Configure > Class of Service > Drop Profile**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—Adds a drop profile. Enter information into the drop profiles page as described in [Table 209 on page 2118](#).
- **Edit**—Modifies an existing drop file. Enter information into the drop profiles page as described in [Table 209 on page 2118](#).
- **Delete**—Deletes an existing drop profile.

**Table 209: Drop Profiles Configuration parameters**

Field	Function	Your Action
Drop Profile Name	Specifies the name for a drop profile.	Type the name.
Drop profile graph	Specifies the drop profile graph type	Select one: <b>Segmented</b> or <b>Interpolated</b> .

Table 209: Drop Profiles Configuration parameters (*continued*)

Field	Function	Your Action
Drop profile values	<p>Specifies values for the following two parameters of the drop profile: the queue fill level and the drop probability.</p> <p>The queue fill level represents a percentage of the memory used to store packets in relation to the total amount that has been allocated for that specific queue.</p> <p>The drop probability is a percentage value that correlates to the likelihood that an individual packet is dropped from the network.</p>	<p>To add new values:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b>.</li> <li>2. Enter the fill level.</li> <li>3. Enter the drop probability.</li> <li>4. Click <b>OK</b>.</li> </ol> <p>To edit an existing value, click <b>Edit</b> and modify the fill level and drop probability.</p> <p>To delete a value, select it and click <b>Delete</b>.</p>

- Related Documentation**
- [Monitoring CoS Drop Profiles on page 2164](#)
  - [Example: Configuring CoS on EX Series Switches on page 2075](#)

## Defining CoS Rewrite Rules (CLI Procedure)

You configure rewrite rules to alter CoS values in outgoing packets on the outbound interfaces of an EX Series switch to match the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure a CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and a code point, thus creating a rewrite table, and you can enable the rewrite rule on an interface. On EX Series switches except EX4300 switches, you can also enable a rewrite rule on routed VLAN interfaces (RVIs). On EX4300 switches, you can also enable rewrite rules on integrated routing and bridging (IRB) interfaces. If you need to customize a rewrite rule, you can create a customized rewrite rule using a firewall filter configuration. You can configure CoS rewrite rules for DSCP, IP precedence and IEEE 802.1p.

You can configure rewrite rules for the following CoS marker types:

- **dscp** and **dscp-ipv6**—Handles incoming IPv4 and IPv6 packets, respectively. On EX4300 switches, you cannot configure DSCP IPv4 and DSCP IPv6 rewrite rules on the same interface. If you configure a DSCP IPv4 rewrite rule on an interface to rewrite IPv4 traffic, then the same rewrite rule is applied to IPv6 traffic also on that interface, and vice versa.
- **ieee-802.1**—Handles Layer 2 CoS.
- **inet-precedence**—Handles incoming IPv4 packets. IP precedence mapping requires only the higher order three bits of the DSCP field.



**NOTE:** To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the rewrite rule and then apply the new rule.

To create IEEE 802.1p rewrite rules and enable them on Layer 2 interfaces:

- To create an IEEE 802.1p rewrite rule named `customup-rw` in the rewrite table for all Layer 2 interfaces:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low code-point 000
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority high code-point 001
user@switch# set ieee-802.1 customup-rw forwarding-class af loss-priority low code-point 010
user@switch# set ieee-802.1 customup-rw forwarding-class af loss-priority high code-point 011
user@switch# set ieee-802.1 customup-rw forwarding-class ef loss-priority low code-point 100
user@switch# set ieee-802.1 customup-rw forwarding-class ef loss-priority high code-point 101
```



```

user@switch# set ieee-802.1p customup-rw forwarding-class nc loss-priority low code-point
110
user@switch# set ieee-802.1p customup-rw forwarding-class nc loss-priority high code-point
111

```

- To enable an IEEE 802.1p rewrite rule named customup-rw on a Layer 2 interface:

```

[edit]
user@switch# set class-of-service interfaces ge-0/0/0 unit 0 rewrite-rules ieee-802.1
customup-rw

```

(On EX4300 switches) To enable an IEEE 802.1p rewrite rule named customup-rw on a Layer 2 interface:

```

[edit]
user@switch# set class-of-service interfaces ge-0/0/0 rewrite-rules ieee-802.1 customup-rw

```

- To enable an IEEE 802.1p rewrite rule named customup-rw on all Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and logical-interface (unit) number:

```

[edit]
user@switch# set class-of-service interfaces ge-* unit * rewrite-rules customup-rw

```

(On EX4300 switches) To enable an IEEE 802.1p rewrite rule named customup-rw on all Gigabit Ethernet interfaces on the switch, use wildcards for the interface name:

```

[edit]
user@switch# set class-of-service interfaces ge-* rewrite-rules customup-rw

```

#### Related Documentation

- [Defining CoS Rewrite Rules \(J-Web Procedure\) on page 2121](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Monitoring CoS Rewrite Rules on page 2161](#)
- [Understanding CoS Rewrite Rules on page 2067](#)

## Defining CoS Rewrite Rules (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You can use the J-Web interface to define CoS rewrite rules. Use the rewrite rules to alter the CoS values in outgoing packets to meet the requirements of the targeted peer. A rewrite rule examines the forwarding class and loss priority of a packet and sets its bits to a corresponding value specified in the rule.

To define rewrite rules:

1. Select **Configure > Class of Service > Rewrite Rules**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—Adds a rewrite rule. Enter information into the rewrite rule page as described in [Table 210 on page 2122](#).
- **Edit**—Modifies an existing rewrite rule. Enter information into the rewrite rule page as described in [Table 210 on page 2122](#).
- **Delete**—Deletes an existing rewrite rule.

**Table 210: Rewrite Rules Configuration Page Summary**

Field	Function	Your Action
Rewrite Rule Name	Specifies the name for the rewrite rule.	To name a rule, type the name—for example, <b>rewrite-dscps</b> .
Rewrite rule type	Specifies the type of rewrite rule: <b>dscp</b> , <b>ieee-802.1</b> , or <b>inet-precedence</b> .	Select a value from the list.
Code Point Mapping	<p>Rewrites outgoing CoS values of a packet based on the forwarding class and loss priority.</p> <p>Allows you to remove a code point mapping entry.</p>	<p>To configure a CoS value assignment, follow these steps:</p> <p>To add a code point mapping:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b>.</li> <li>2. Select the code point.</li> <li>3. Select a forwarding class from the following list: <ul style="list-style-type: none"> <li>• <b>expedited-forwarding</b>—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service. Packets can be forwarded out of sequence or dropped.</li> <li>• <b>best-effort</b>—Provides no special CoS handling of packets. Typically, RED drop profile is aggressive and no loss priority is defined.</li> <li>• <b>assured-forwarding</b>—Provides high assurance for packets within the specified service profile. Excess packets are dropped.</li> <li>• <b>network-control</b>—Packets can be delayed but not dropped.</li> </ul> </li> <li>4. Select the loss priority. <p>To assign a loss priority, select one:</p> <ul style="list-style-type: none"> <li>• <b>high</b>—Packet has a high loss priority.</li> <li>• <b>low</b>—Packet has a low loss priority.</li> </ul> <p>To edit an existing code point mapping, select it and click <b>Edit</b>.</p> <p>To remove a code point mapping entry, select it and click <b>Remove</b>.</p> </li> </ol>

- Related Documentation**
- [Defining CoS Rewrite Rules \(CLI Procedure\) on page 2120](#)
  - [Understanding CoS Rewrite Rules on page 2067](#)
  - [Monitoring CoS Rewrite Rules on page 2161](#)
  - [Example: Configuring CoS on EX Series Switches on page 2075](#)

## Assigning CoS Components to Interfaces (CLI Procedure)

After you have defined the following CoS components, you must assign them to logical or physical interfaces.

- Forwarding classes—Assign only to logical interfaces.
- Classifiers—Assign only to logical interfaces.
- Scheduler maps—Assign to either physical or logical interfaces.
- Rewrite rules—Assign to either physical or logical interfaces.

You can assign a CoS component to a single interface or to multiple interfaces using wild cards.

To assign CoS components to interfaces:

- To assign CoS components to a single interface, associate a CoS component (for example a scheduler map named **ethernet-cos-map**) with an interface:

```
[edit class-of-service interfaces]
user@switch# set ge-0/0/20 scheduler-map ethernet-cos-map
```

- To assign a CoS component to multiple interfaces, associate a CoS component (for example, a rewrite rule named **customup-rw**) to all Gigabit Ethernet interfaces on the switch, use wild characters for the interface name and logical-interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set ge-* unit * rewrite-rules ieee-802.1 customup-rw
```

- Related Documentation**
- [Assigning CoS Components to Interfaces \(J-Web Procedure\) on page 2123](#)
  - [Example: Configuring CoS on EX Series Switches on page 2075](#)
  - [Monitoring Interfaces That Have CoS Components on page 2160](#)
  - [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)

## Assigning CoS Components to Interfaces (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

After you have defined CoS components on an EX Series switch, you must assign them to logical or physical interfaces. You can use the J-Web interface to assign scheduler

maps to physical or logical interfaces and to assign forwarding classes or classifiers to logical interfaces.

To assign CoS components to interfaces:

1. Select **Configure > Class of Service > Assign to Interface**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See *Using the Commit Options to Commit Configuration Changes* for details about all commit options.

2. To configure interface association, select an interface from the list and click **Edit**. For an EX8200 Virtual Chassis configuration, select the member, the FPC, and the interface from the list, and click **Edit**.
3. Select one of the following:
  - **Associate system default scheduler map**—Associates the interface with the default scheduler map.
  - **Select the scheduler map**—Associates the interface with a configured scheduler map. Select the scheduler map from the list.



**NOTE:** On the 40-port SFP+ line card for EX8200 switches, you cannot commit your changes using the J-Web interface unless you assign the same scheduler map or the default scheduler map to all interfaces in a port group.

4. Click **OK**.
5. To manage a CoS service assignment on a logical interface, Click one of the following options:
  - **Add**—Adds a CoS service to a logical interface on a specified physical interface. Enter information as described in [Table 211 on page 2124](#).
  - **Edit**—Modifies a CoS service assignment to a logical interface. Enter information as described in [Table 211 on page 2124](#).
  - **Delete**—Deletes the CoS service assignment to a logical interface.

**Table 211: Assigning CoS Components to Logical Interfaces**

Field	Function	Your Action
Unit	Specifies the name of a logical interface. Allows you to assign CoS components while configuring a logical interface on a physical interface at the same time.	Type the interface name.  To assign CoS services to all logical interfaces configured on this physical interface, type the wildcard character (*).

Table 211: Assigning CoS Components to Logical Interfaces (*continued*)

Field	Function	Your Action
Forwarding Class	Assigns a predefined forwarding class to incoming packets on a logical interface.	To assign a forwarding class to an interface, select the forwarding class.
Classifiers	Allows you to apply classification maps to a logical interface. Classifiers assign a forwarding class and loss priority to an incoming packet based on its CoS value.	To assign a classification map to an interface, select an appropriate classifier for each CoS value type used on the interface.
Rewrite Rules	<p>Allows you to alter the CoS values in outgoing packets to meet the requirements of the targeted peer. A rewrite rule examines the forwarding class and loss priority of a packet and sets its bits to a corresponding value specified in the rule.</p> <p><b>NOTE:</b> In EX4300 switches, this option is available only when you click <b>Edit</b> button in the Configure Interface Association table.</p>	To assign rewrite rules to the interface, select the appropriate rewrite rule for each CoS value type used on the interface.

- Related Documentation**
- [Assigning CoS Components to Interfaces \(CLI Procedure\) on page 2123](#)
  - [Example: Configuring CoS on EX Series Switches on page 2075](#)
  - [Monitoring Interfaces That Have CoS Components on page 2160](#)

## Configuring Junos OS EZQoS for CoS (CLI Procedure)

You use Junos OS EZQoS on EX Series switches to eliminate the complexities involved in configuring class of service (CoS) across the network. EZQoS offers templates for key traffic classes.

When you configure EZQoS on EX Series switches, preconfigured values are assigned to all CoS parameters based on the typical application requirements. These preconfigured values are stored in a template with a unique name.



**NOTE:** Currently, we provide an EZQoS template for configuring CoS for VoIP applications. The EZQoS VoIP template is stored in `/etc/config/ezqos-voip.conf`.

To configure EZQoS using the CLI:

1. Load the EZQoS configuration file (`/etc/config/ezqos-voip.conf`):
 

```
[edit]
user@switch# load merge /etc/config/ezqos-voip.conf
```
2. Apply the EZQoS group (`ezqos-voip`):
 

```
[edit]
user@switch# set apply-groups ezqos-voip
```
3. Apply the DSCP classifier (`ezqos-dscp-classifier`) to a Gigabit Ethernet interface (`ge-0/0/0`):
 

```
[edit class-of-service interfaces]
```

- ```
user@switch# set ge-0/0/0 unit 0 classifiers dscp ezqos-dscp-classifier
```
4. Apply the scheduler map (**ezqos-voip-sched-maps**) to a Gigabit Ethernet interface (**ge-0/0/1**):

```
[edit class-of-service interfaces]
```

```
user@switch# set ge-0/0/1 scheduler-map ezqos-voip-sched-maps
```

**Related  
Documentation**

- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Understanding Junos OS EZQoS for CoS Configurations on EX Series Switches on page 2070](#)

## Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure)

EX8200 switches provide certain line cards that include oversubscribed ports. These ports are logically grouped into a port group and each port group share a certain fixed bandwidth. Because oversubscribed ports handle traffic differently than ports that provide continuous line-rate bandwidth, configuring CoS queues is different for oversubscribed ports than for line-rate ports.

Packets arriving on an oversubscribed port in a line card are directed to a high-priority, low priority, or line-rate queue. These queues are used for scheduling traffic from the port into the Packet Forwarding Engine. The fabric priority associated with the packet's forwarding class determines which queue the packet is sent to. The forwarding class of the packet in turn is determined by the behavior aggregate (BA) classifier assigned to the port. By default, the fabric priority of all forwarding classes is low. Thus all packets, with the exception of critical network packets and line-rate packets, are sent to the low-priority ingress queue by default. The critical network packets and line-rate packets do not need a BA classifier as they are always sent on the high-priority and line-rate queues, respectively.

This procedure describes how you can direct traffic into the high-priority ingress queue and thus avoid congestion at the port group.

To direct traffic to the high-priority ingress queue for a port group:

1. Create the BA classifier for the forwarding class:

```
[edit class-of-service]
```

```
user@switch# set classifiers classifier-type classifier-name
```

```
forwarding-class class-name loss-priority level code-points code-point
```

2. Assign a queue number and fabric priority to the forwarding class:

```
[edit class-of-service]
```

```
user@switch# set forwarding-classes class class-name queue-num number  
priority level
```

3. Assign the BA classifier to the physical interface:

```
[edit class-of-service]
```

```
user@switch# set interfaces interface-name unit 0  
classifiers classifier-type classifier-name
```

For example, to direct voice traffic to the high-priority ingress queue for interface **xe-1/0/2**:

```
[edit class-of-service]
user@switch# set classifiers dscp dscp1 forwarding-class cos-voice
loss-priority low code-points ef
```

```
[edit class-of-service]
user@switch# set forwarding-classes class cos-voice queue-num 5 priority high
```

```
[edit class-of-service]
user@switch# set interfaces xe-1/0/2 unit 0 classifiers dscp dscp1
```



**NOTE:** You must use a BA classifier to classify traffic for ingress queuing. Multifield (MF) classification and port classification (that is, assigning a forwarding class to the interface) are not supported for classifying traffic for ingress queuing. The BA classifier must be assigned to a physical interface, not a Layer 3 tagged interface or a routed VLAN interface (RVI).

**Related  
Documentation**

- [Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports on page 2071](#)

## Configuration Statements

- [\[edit class-of-service\] Configuration Statement Hierarchy on EX Series Switches on page 2128](#)
- [broadcast on page 2131](#)
- [buffer-size on page 2132](#)
- [classifiers on page 2133](#)
- [code-point-aliases on page 2134](#)
- [code-points on page 2134](#)
- [drop-profile-map on page 2135](#)
- [dscp on page 2136](#)
- [dscp-ipv6 on page 2137](#)
- [ethernet \(CoS for Multidestination Traffic\) on page 2138](#)
- [excess-rate \(Schedulers\) on page 2138](#)
- [family on page 2139](#)
- [forwarding-class on page 2140](#)
- [forwarding-classes on page 2141](#)
- [ieee-802.1 on page 2142](#)
- [import on page 2143](#)
- [inet-precedence on page 2144](#)
- [interfaces on page 2145](#)
- [loss-priority \(Classifiers and Rewrite Rules\) on page 2146](#)

- [policing on page 2147](#)
- [priority \(Schedulers\) on page 2148](#)
- [protocol \(Drop Profiles\) on page 2148](#)
- [rewrite-rules on page 2149](#)
- [scheduler-map on page 2150](#)
- [scheduler-maps on page 2151](#)
- [schedulers \(CoS\) on page 2152](#)
- [shaping-rate on page 2153](#)
- [transmit-rate \(EX Series Switches\) on page 2154](#)
- [unit on page 2155](#)

## **[edit class-of-service] Configuration Statement Hierarchy on EX Series Switches**

This topic lists supported and unsupported configuration statements in the **[edit class-of-service]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit class-of-service\] Hierarchy Level on page 2128](#)
- [Unsupported Statements in the \[edit class-of-service\] Hierarchy Level on page 2130](#)

### **Supported Statements in the [edit class-of-service] Hierarchy Level**

The following hierarchy shows the **[edit class-of-service]** configuration statements supported on EX Series switches:

```
class-of-service {
  classifiers {
    (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) {
          code-points [ aliases ] [ 6 bit-patterns ];
        }
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) {
      alias-name bits;
    }
  }
}
```



```

drop-profiles {
  profile-name {
    interpolate {
      drop-probability [values];
      fill-level [values]
    }
  }
}
forwarding-classes {
  class class-name
  queue queue-number;
}
interfaces interface-name {
  scheduler-map map-name;
  shaping-rate rate;
  unit (logical-unit-number | * ) {
    classifiers {
      (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) (classifier-name | default);
    }
    forwarding-class class-name ;
  }
}
rewrite-rules {
  (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) (rewrite-rule-name | default);
}
rewrite-rules {
  (dscp | dscp-ipv6 | ieee-802.1 | inet-precedence) rewrite-name {
    import (default | rewrite-name);
    forwarding-class class-name {
      loss-priority (high | low | medium-high | medium-low) code-point (alias | bits);
    }
  }
}
scheduler-maps {
  map-name {
    forwarding-class class-name {
      scheduler scheduler-name;
    }
  }
}
schedulers {
  scheduler-name {
    buffer-size (exact | percent percentage | remainder);
    drop-profile-map {
      loss-priority (any | high | medium-high | medium-low);
      protocol any;
      {
        drop-profile profile-name
      }
    }
    excess-rate {
      percent percentage;
    }
    priority (low | strict-high);
    shaping-rate (rate | percent percentage);
  }
}

```

```
        transmit-rate (EX Series Switches) (rate | percent percentage | remainder) ;
    }
}
shared-buffer {
    percent;
}
traceoptions {
    file (file-name | files files | match match | no-world-readable | size size | world-readable);
    flag ( all | asynch | chassis-scheduler | cos-adjustment | dynamic | hardware-database
        | init | parse | performance-monitor | process | restart | route-socket | show | snmp |
        util);
    no-remote-trace;
}
tri-color;
}
```

### Unsupported Statements in the [edit class-of-service] Hierarchy Level

All statements in the [edit class-of-service] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### **Related Documentation**

- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Defining CoS Code-Point Aliases \(CLI Procedure\) on page 2101](#) or [Defining CoS Code-Point Aliases \(J-Web Procedure\) on page 2101](#)
- [Defining CoS Classifiers \(CLI Procedure\) on page 2103](#) or [Defining CoS Classifiers \(J-Web Procedure\) on page 2105](#)
- [Defining CoS Forwarding Classes \(CLI Procedure\) on page 2107](#) or [Defining CoS Forwarding Classes \(J-Web Procedure\) on page 2107](#)
- [Configuring CoS Tail Drop Profiles \(CLI Procedure\)](#)
- [Defining CoS Schedulers and Scheduler Maps \(CLI Procedure\) on page 2109](#) or [Defining CoS Schedulers \(J-Web Procedure\) on page 2111](#)
- [Defining CoS Rewrite Rules \(CLI Procedure\) on page 2120](#) or [Defining CoS Rewrite Rules \(J-Web Procedure\) on page 2121](#)

---

## broadcast

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>broadcast forwarding-class-name;</code>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit class-of-service multi-destination family <a href="#">ethernet</a> ]                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify the forwarding class for the broadcast traffic belonging to the Ethernet family.                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><i>forwarding-class-name</i> —Name of the forwarding class:</p> <ul style="list-style-type: none"><li>• <b>mcast-af</b>—Default forwarding class for assured forwarding of multicast traffic.</li><li>• <b>mcast-be</b>—Default best-effort forwarding class for multicast traffic.</li><li>• <b>mcast-ef</b>—Default forwarding class for expedited forwarding of multicast traffic.</li></ul> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding CoS Schedulers on page 2060</a></li><li>• <i>Understanding CoS Forwarding Classes</i></li><li>• <i>Understanding CoS Classifiers</i></li></ul>                                                                                                                                                                                   |

## buffer-size

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | buffer-size (exact   percent <i>percentage</i>   remainder   temporal);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit class-of-service <a href="#">schedulers</a> <i>scheduler-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Specify buffer size in a scheduler configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Default</b>                  | On EX Series switches except EX4300 switches, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent, respectively. On EX4300 switches, the default scheduler transmission rate and buffer size for queues 0 through 11 are 75, 0, 0, 5, 0, 0, 0, 0, 15, 0, 0 and 5 percent, respectively, of the total available buffer.                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>exact</b>—(Except on EX8200 standalone switches and EX8200 Virtual Chassis) Enforce the exact buffer size. When this option is configured, sharing is disabled on the queue, restricting the usage to guaranteed buffers only.</p> <p><b>percent <i>percentage</i></b>—Buffer size as a percentage of the total buffer.</p> <p><b>remainder</b>—Remaining buffer available.</p> <p><b>temporal</b>—(EX4200 standalone switches, EX4200 Virtual Chassis, EX4300 standalone switches, EX4300 Virtual Chassis, EX8200 standalone switches, and EX8200 Virtual Chassis only) Buffer size as a temporal value.</p> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Defining CoS Schedulers and Scheduler Maps (CLI Procedure) on page 2109</a> or <a href="#">Defining CoS Schedulers (J-Web Procedure) on page 2111</a></li><li>• <a href="#">Understanding CoS Schedulers on page 2060</a></li></ul>                                                                                                                                                                                                                                       |

## classifiers

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> classifiers {   (dscp   dscp-ipv6   ieee-802.1   inet-precedence) classifier-name {     import (classifier-name   default);     forwarding-class class-name {       loss-priority level {         code-points [aliases] [6-bit-patterns];       }     }   } } </pre>                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit class-of-service],<br>[edit class-of-service <a href="#">interfaces</a> interface-name unit logical-unit-number]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Expanded to include EXP classifiers in Junos OS Release 10.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Apply a CoS aggregate behavior classifier to a logical interface. You can apply a default classifier or a custom classifier.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Example: Combining CoS with MPLS on EX Series Switches</a></li> <li>• <a href="#">Defining CoS Classifiers (CLI Procedure) on page 2103</a> or <a href="#">Defining CoS Classifiers (J-Web Procedure) on page 2105</a></li> <li>• <a href="#">Assigning CoS Components to Interfaces (CLI Procedure) on page 2123</a> or <a href="#">Assigning CoS Components to Interfaces (J-Web Procedure) on page 2123</a></li> <li>• <a href="#">Understanding CoS Classifiers</a></li> </ul> |

## code-point-aliases

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>code-point-aliases {<br/>    (dscp   dscp-ipv6   ieee-802.1   inet-precedence) [{<br/>        alias-name bits;<br/>    }<br/>}</pre>                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit class-of-service]                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Define an alias for a CoS marker.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Defining CoS Code-Point Aliases (CLI Procedure) on page 2101</a> or <a href="#">Defining CoS Code-Point Aliases (J-Web Procedure) on page 2101</a></li><li>• <a href="#">Understanding CoS Code-Point Aliases on page 2046</a></li></ul> |

## code-points

---

|                                 |                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>code-points [ aliases ] [ 6 bit-patterns ];</pre>                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit class-of-service classifiers (dscp   ieee-802.1   inet-precedence) forwarding-class class-name loss-priority level]                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Specify one or more DSCP code-point aliases or bit sets for association with a forwarding class.                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><i>aliases</i> —Name of the DSCP alias.</p> <p><i>6 bit-patterns</i> —Value of the code-point bits, in decimal form.</p>                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Defining CoS Classifiers (CLI Procedure) on page 2103</a> or <a href="#">Defining CoS Classifiers (J-Web Procedure) on page 2105</a></li><li>• <a href="#">Understanding CoS Classifiers</a></li></ul> |

## drop-profile-map

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | drop-profile-map <b>loss-priority</b> <i>loss-priority</i> <b>protocol</b> <i>protocol</i> drop-profile <i>profile-name</i> ;                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit class-of-service <b>schedulers</b> <i>scheduler-name</i> ]                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Define the loss priority value for the specified drop profile.                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>drop-profile</b> <i>profile-name</i> —Name of the drop profile.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Defining CoS Schedulers and Scheduler Maps (CLI Procedure) on page 2109</a> or <a href="#">Defining CoS Schedulers (J-Web Procedure) on page 2111</a></li> <li>• <a href="#">Understanding CoS Schedulers on page 2060</a></li> </ul> |

## dscp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>dscp classifier-name {<br/>    import (classifier-name   default);<br/>    forwarding-class class-name {<br/>        loss-priority level {<br/>            code-points [ aliases ] [ 6-bit-patterns ];<br/>        }<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit class-of-service classifiers],<br>[edit class-of-service <a href="#">code-point-aliases</a> ],<br>[edit class-of-service <a href="#">interfaces</a> interface-name <a href="#">unit</a> logical-unit-number classifiers],<br>[edit class-of-service <a href="#">rewrite-rules</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Define the Differentiated Services code point (DSCP) mapping that is applied to the packets.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b>classifier-name</b> —Name of the classifier.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Defining CoS Code-Point Aliases (CLI Procedure) on page 2101</a> or <a href="#">Defining CoS Code-Point Aliases (J-Web Procedure) on page 2101</a></li><li>• <a href="#">Defining CoS Classifiers (CLI Procedure) on page 2103</a> or <a href="#">Defining CoS Classifiers (J-Web Procedure) on page 2105</a></li><li>• <a href="#">Defining CoS Rewrite Rules (CLI Procedure) on page 2120</a> or <a href="#">Defining CoS Rewrite Rules (J-Web Procedure) on page 2121</a></li><li>• <a href="#">Assigning CoS Components to Interfaces (CLI Procedure) on page 2123</a> or <a href="#">Assigning CoS Components to Interfaces (J-Web Procedure) on page 2123</a></li><li>• <a href="#">Understanding CoS Classifiers</a></li></ul> |



## dscp-ipv6

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> dscp-ipv6 classifier-name {   import (classifier-name   default);   forwarding-class class-name {     loss-priority level {       code-points [aliases] [6-bit-patterns];     }   } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <p>[edit class-of-service classifiers],<br/> [edit class-of-service <b>code-point-aliases</b>],<br/> [edit class-of-service <b>interfaces</b> interface-name <b>unit</b> logical-unit-number classifiers]<br/> [edit class-of-service <b>interfaces</b> interface-name <b>unit</b> logical-unit-number <b>rewrite-rules</b>]<br/> [edit class-of-service <b>rewrite-rules</b>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Define the Differentiated Services code point (DSCP) mapping that is applied to the IPv6 packets.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>classifier-name</b>—Name of the classifier.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Defining CoS Code-Point Aliases (CLI Procedure) on page 2101</a> or <a href="#">Defining CoS Code-Point Aliases (J-Web Procedure) on page 2101</a></li> <li>• <a href="#">Defining CoS Classifiers (CLI Procedure) on page 2103</a> or <a href="#">Defining CoS Classifiers (J-Web Procedure) on page 2105</a></li> <li>• <a href="#">Defining CoS Rewrite Rules (CLI Procedure) on page 2120</a> or <a href="#">Defining CoS Rewrite Rules (J-Web Procedure) on page 2121</a></li> <li>• <a href="#">Assigning CoS Components to Interfaces (CLI Procedure) on page 2123</a> or <a href="#">Assigning CoS Components to Interfaces (J-Web Procedure) on page 2123</a></li> <li>• <a href="#">Understanding CoS Classifiers</a></li> </ul> |

## ethernet (CoS for Multidestination Traffic)

---

|                                 |                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ethernet {<br/>    <b>broadcast</b> <i>forwarding-class-name</i>;<br/>}</pre>                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit class-of-service multi-destination <b>family</b> ]                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                               |
| <b>Description</b>              | <p>Specify the Ethernet broadcast traffic family.</p> <p>The remaining statement is explained separately.</p>                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding CoS Schedulers on page 2060</a></li><li>• <a href="#">Understanding CoS Forwarding Classes</a></li><li>• <a href="#">Understanding CoS Classifiers</a></li></ul> |

## excess-rate (Schedulers)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>excess-rate {<br/>    percent <i>percentage</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service</b> on page 328 <b>schedulers</b> <i>scheduler-name</i> ]                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | (EX4300 switches only) Specify the percentage of excess bandwidth traffic to share.                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | Excess bandwidth is shared in proportion to the configured transmit rate of each queue.                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <ul style="list-style-type: none"><li>• <b>percent</b>—Percentage of the excess bandwidth to share.</li></ul>                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding CoS Schedulers on page 2060</a></li><li>• <a href="#">Defining CoS Schedulers and Scheduler Maps (CLI Procedure) on page 2109</a> or <a href="#">Defining CoS Schedulers (J-Web Procedure) on page 2111</a></li><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li></ul> |

## family

|                                 |                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>family {   ethernet {     broadcast forwarding-class-name;   }   inet {     classifiers{       (dscp   ieee-802.1   inet-precedence) classifier-name;     }   } }</pre>                                                           |
| <b>Hierarchy Level</b>          | [edit class-of-service multi-destination]                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                   |
| <b>Description</b>              | <p>Specify the multideestination traffic family.</p> <p>The remaining statements are explained separately.</p>                                                                                                                         |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.interface-control—To add this statement to the configuration.</p>                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding CoS Schedulers on page 2060</a></li> <li>• <a href="#">Understanding CoS Forwarding Classes</a></li> <li>• <a href="#">Understanding CoS Classifiers</a></li> </ul> |

## forwarding-class

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>forwarding-class <i>class-name</i> {<br/>    loss-priority <i>level</i> {<br/>        code-points [<i>aliases</i>] [<i>6-bit-patterns</i>];<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit class-of-service classifiers ( <a href="#">dscp</a>   <a href="#">ieee-802.1</a>   <a href="#">inet-precedence</a> ) <i>classifier-name</i> ],<br>[edit class-of-service <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ],<br>[edit class-of-service <a href="#">rewrite-rules</a> ] ( <a href="#">dscp</a>   <a href="#">ieee-802.1</a>   <a href="#">inet-precedence</a> ) <i>rewrite-rule-name</i> ],<br>[edit class-of-service <a href="#">scheduler-maps</a> <i>map-name</i> ],<br>[edit class-of-service host-outbound-traffic] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Define forwarding class name and option values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><i>class-name</i> —Name of the forwarding class.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Defining CoS Forwarding Classes (CLI Procedure) on page 2107</a> or <a href="#">Defining CoS Forwarding Classes (J-Web Procedure) on page 2107</a></li><li>• <a href="#">Understanding CoS Forwarding Classes</a></li></ul>                                                                                                                                                                                                                         |

## forwarding-classes

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | forwarding-classes {<br>class <i>class-name</i> queue-num <i>queue-number</i> priority (high   medium-high   low  <br>medium-low);<br>}                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit class-of-service]                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Associate the forwarding class with a queue name and number.<br><br>The statement is explained separately.                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Defining CoS Forwarding Classes (CLI Procedure) on page 2107</a> or <a href="#">Defining CoS Forwarding Classes (J-Web Procedure) on page 2107</a></li> <li>• <a href="#">Understanding CoS Forwarding Classes</a></li> </ul> |

## ieee-802.1

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ieee-802.1 classifier-name {<br/>    import (classifier-name   default);<br/>    forwarding-class class-name {<br/>        loss-priority level {<br/>            code-points [ aliases ] [ 6 bit-patterns ];<br/>        }<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit class-of-service classifiers],<br>[edit class-of-service <a href="#">code-point-aliases</a> ],<br>[edit class-of-service <a href="#">interfaces</a> interface-name <a href="#">unit</a> logical-unit-number classifiers],<br>[edit class-of-service <a href="#">rewrite-rules</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Apply an IEEE-802.1 rewrite rule.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>classifier-name</b> —Name of the classifier.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Defining CoS Classifiers (CLI Procedure) on page 2103</a> or <a href="#">Defining CoS Classifiers (J-Web Procedure) on page 2105</a></li><li>• <a href="#">Defining CoS Code-Point Aliases (CLI Procedure) on page 2101</a> or <a href="#">Defining CoS Code-Point Aliases (J-Web Procedure) on page 2101</a></li><li>• <a href="#">Defining CoS Rewrite Rules (CLI Procedure) on page 2120</a> or <a href="#">Defining CoS Rewrite Rules (J-Web Procedure) on page 2121</a></li><li>• <a href="#">Understanding CoS Classifiers</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 2067</a></li></ul> |

## import

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import (classifier-name   default);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit class-of-service classifiers (<a href="#">dscp</a>   <a href="#">ieee-802.1</a>   <a href="#">inet-precedence</a>) <i>classifier-name</i>],</p> <p>[edit class-of-service <a href="#">rewrite-rules</a> (<a href="#">dscp</a>   <a href="#">ieee-802.1</a>   <a href="#">inet-precedence</a>) <i>rewrite-name</i>]</p>                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Specify a default or previously defined classifier.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>classifier-name</b> —Name of the classifier mapping configured at the [edit class-of-service classifiers] hierarchy level.</p> <p><b>default</b>—Default classifier mapping.</p>                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Defining CoS Classifiers (CLI Procedure) on page 2103</a> or <a href="#">Defining CoS Classifiers (J-Web Procedure) on page 2105</a></li> <li>• <a href="#">Defining CoS Rewrite Rules (CLI Procedure) on page 2120</a> or <a href="#">Defining CoS Rewrite Rules (J-Web Procedure) on page 2121</a></li> <li>• <a href="#">Understanding CoS Classifiers</a></li> <li>• <a href="#">Understanding CoS Rewrite Rules on page 2067</a></li> </ul> |

## inet-precedence

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>inet-precedence classifier-name {<br/>    import (classifier-name   default);<br/>    forwarding-class class-name {<br/>        loss-priority level {<br/>            code-points [ aliases ] [ 6-bit-patterns ];<br/>        }<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit class-of-service classifiers],<br>[edit class-of-service <a href="#">code-point-aliases</a> ],<br>[edit class-of-service <a href="#">interfaces</a> interface-name <a href="#">unit</a> logical-unit-number classifiers],<br>[edit class-of-service <a href="#">rewrite-rules</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Apply an IPv4 precedence rewrite rule.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><i>classifier-name</i>—Name of the classifier.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Defining CoS Classifiers (CLI Procedure) on page 2103</a> or <a href="#">Defining CoS Classifiers (J-Web Procedure) on page 2105</a></li><li>• <a href="#">Defining CoS Code-Point Aliases (CLI Procedure) on page 2101</a> or <a href="#">Defining CoS Code-Point Aliases (J-Web Procedure) on page 2101</a></li><li>• <a href="#">Defining CoS Rewrite Rules (CLI Procedure) on page 2120</a> or <a href="#">Defining CoS Rewrite Rules (J-Web Procedure) on page 2121</a></li><li>• <a href="#">Understanding CoS Classifiers</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 2067</a></li></ul> |



## interfaces

```
Syntax interfaces {
    interface-name {
        congestion-notification-profile profile-name {
            input {
                ieee-802.1 {
                    code-point up-bits pfc;
                }
            }
        }
        scheduler-map map-name;
        unit logical-unit-number {
            forwarding-class class-name;
            classifiers {
                (dscp | ieee-802.1 | inet-precedence) (classifier-name | default);
            }
        }
    }
}
```

**Hierarchy Level** [edit class-of-service]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure interface-specific class-of-service (CoS) properties for incoming packets.

**Options** *interface-name*—Name of the interface.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring CoS on EX Series Switches on page 2075](#)
- [Defining CoS Classifiers \(CLI Procedure\) on page 2103](#) or [Defining CoS Classifiers \(J-Web Procedure\) on page 2105](#)
- [Defining CoS Forwarding Classes \(CLI Procedure\) on page 2107](#) or [Defining CoS Forwarding Classes \(J-Web Procedure\) on page 2107](#)
- [Defining CoS Schedulers and Scheduler Maps \(CLI Procedure\) on page 2109](#) or [Defining CoS Schedulers \(J-Web Procedure\) on page 2111](#)
- [Configuring Priority-Based Flow Control for an EX Series Switch \(CLI Procedure\)](#)

## loss-priority (Classifiers and Rewrite Rules)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>loss-priority level {<br/>    code-points [aliases] [6-bit-patterns   3-bit-patterns];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit class-of-service classifiers (dscp   <a href="#">ieee-802.1</a>   <a href="#">inet-precedence</a>   exp) classifier-name forwarding-class class-name],<br>[edit class-of-service rewrite-rules (dscp   <a href="#">ieee-802.1</a>   <a href="#">inet-precedence</a>   exp) rewrite-rule-name forwarding-class class-name]                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement expanded to apply to EXP classifiers in Junos OS Release 10.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Specify packet loss priority value for a specific set of code-point aliases and bit patterns.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>level</b> —Can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>high</b>—Packet has high loss priority.</li><li>• <b>medium-high</b>— (On EX3200, EX4200, EX4300, and EX4500 switches only) Code points to classify to loss priority medium-high.</li><li>• <b>low</b>—Packet has low loss priority.</li><li>• <b>medium-low</b> — (On EX3200, EX4200, EX4300, and EX4500 switches only) Code points to classify to loss priority medium-low.</li></ul> <p>The remaining statement is explained separately.</p>                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Defining CoS Classifiers (CLI Procedure) on page 2103</a> or <a href="#">Defining CoS Classifiers (J-Web Procedure) on page 2105</a></li><li>• <a href="#">Defining CoS Rewrite Rules (CLI Procedure) on page 2120</a> or <a href="#">Defining CoS Rewrite Rules (J-Web Procedure) on page 2121</a></li><li>• <a href="#">Understanding CoS Classifiers</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 2067</a></li></ul> |

## policing

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>policing (filter <i>filter-name</i>   no-automatic-policing);</code>                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit protocols mpls label-switched-path <i>lsp-name</i>]</code><br><code>[edit interfaces <i>interface-id</i> unit <i>number-of-logical-unit</i> family inet address <i>ip-address</i>]</code>                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Apply a rate-limiting policer as the specified policing filter: <ul style="list-style-type: none"> <li>• To the LSP for MPLS over CCC.</li> <li>• To the customer-edge interface for IP over MPLS.</li> </ul>                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>filter <i>filter-name</i></b> —Specify the name of the policing filter.<br><br><b>no-automatic-policing</b> —Disable automatic policing on this LSP.                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">policer on page 4859</a></li> <li>• <a href="#">Configuring Policers to Control Traffic Rates (CLI Procedure) on page 4818</a></li> <li>• <a href="#">Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure)</a></li> <li>• <a href="#">Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure)</a></li> </ul> |

## priority (Schedulers)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority</code> <i>priority</i> ;                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit class-of-service <a href="#">schedulers</a> <i>scheduler-name</i> ]                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify packet-scheduling priority value.                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>priority</i> —It can be one of the following: <ul style="list-style-type: none"><li>• <b>low</b>—Scheduler has low priority.</li><li>• <b>strict-high</b>—Scheduler has strictly high priority.</li></ul>                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Defining CoS Schedulers and Scheduler Maps (CLI Procedure) on page 2109</a> or <a href="#">Defining CoS Schedulers (J-Web Procedure) on page 2111</a></li><li>• <a href="#">Understanding CoS Schedulers on page 2060</a></li></ul> |

## protocol (Drop Profiles)

---

|                                 |                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>protocol</code> <i>protocol</i> drop-profile <i>profile-name</i> ;                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit class-of-service <a href="#">schedulers</a> <i>scheduler-name</i> ]                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                      |
| <b>Description</b>              | Specify the protocol type for the specified drop profile.                                                                                                                                                                                                                 |
| <b>Options</b>                  | <b>drop-profile</b> <i>profile-name</i> —Name of the drop profile.<br><br><i>protocol</i> —Type of protocol. It can be: <ul style="list-style-type: none"><li>• <b>any</b>—Accept any protocol type.</li></ul>                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Configuring CoS Tail Drop Profiles (CLI Procedure)</a></li><li>• <a href="#">Understanding CoS Tail Drop Profiles</a></li></ul> |

## rewrite-rules

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rewrite-rules {   (dscp   dscp-ipv6   exp   ieee-802.1   inet-precedence ) rewrite-name {     import ( default   rewrite-name);     forwarding-class class-name {       loss-priority level code-point (alias   bits);     }   } }</pre>                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit class-of-service]                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement expanded for use with global EXP classifiers in Junos OS Release 10.1 for EX Series switches.</p>                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Defining CoS Rewrite Rules (CLI Procedure) on page 2120</a> or <a href="#">Defining CoS Rewrite Rules (J-Web Procedure) on page 2121</a></li> <li>• <a href="#">Understanding CoS Rewrite Rules on page 2067</a></li> <li>• <a href="#">Understanding Using CoS with MPLS Networks on EX Series Switches</a></li> </ul> |

## scheduler-map

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>scheduler-map <i>map-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit class-of-service <a href="#">interfaces</a> ],<br>[edit class-of-service multi-destination]                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Associate a scheduler map name with an interface or with a multidestination traffic configuration.                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <i>map-name</i> —Name of the scheduler map.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Assigning CoS Components to Interfaces (CLI Procedure) on page 2123</a> or <a href="#">Assigning CoS Components to Interfaces (J-Web Procedure) on page 2123</a></li><li>• <a href="#">Understanding CoS Schedulers on page 2060</a></li><li>• <a href="#">Understanding CoS Classifiers</a></li></ul> |

## scheduler-maps

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> scheduler-maps {   map-name {     forwarding-class class-name scheduler scheduler-name;   } } </pre>                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit class-of-service]                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify a scheduler map name and associate it with the scheduler configuration and forwarding class.                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>map-name</b> —Name of the scheduler map.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Defining CoS Forwarding Classes (CLI Procedure) on page 2107</a> or <a href="#">Defining CoS Forwarding Classes (J-Web Procedure) on page 2107</a></li> <li>• <a href="#">Understanding CoS Schedulers on page 2060</a></li> <li>• <a href="#">Understanding CoS Forwarding Classes</a></li> </ul> |

## schedulers (CoS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>schedulers {<br/>  scheduler-name {<br/>    buffer-size (percent <i>percentage</i>   remainder);<br/>    drop-profile-map loss-priority <i>loss-priority</i> protocol <i>protocol</i> drop-profile <i>profile-name</i>;<br/>    excess-rate (percent <i>percentage</i>);<br/>    priority <i>priority</i>;<br/>    shaping-rate (<i>rate</i>   percent <i>percentage</i>);<br/>    transmit-rate (EX Series Switches) (<i>rate</i>   percent <i>percentage</i>   remainder);<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit class-of-service]                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Specify scheduler name and parameter values.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><i>scheduler-name</i> —Name of the scheduler.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Defining CoS Schedulers and Scheduler Maps (CLI Procedure) on page 2109</a> or <a href="#">Defining CoS Schedulers (J-Web Procedure) on page 2111</a></li><li>• <a href="#">Understanding CoS Schedulers on page 2060</a></li></ul>                                                                                                                         |



## shaping-rate

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | shaping-rate (percent <i>percentage</i>   rate);                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit class-of-service <a href="#">schedulers (CoS)</a> <i>scheduler-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Configure shaping rate to throttle the rate at which queues transmit packets.</p> <p>We recommend that you configure the shaping rate as an absolute maximum usage and not as additional usage beyond the configured transmit rate.</p>                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | If you do not include this statement, the default shaping rate is 100 percent, which is the same as no shaping at all.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>percentpercentage</b>—Shaping rate as a percentage of the available interface bandwidth.<br/> <b>Range:</b> 0 through 100 percent</p> <p><b>rate</b>—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).<br/> <b>Range:</b> 3200 through 32,000,000,000 bps<br/>         (EX4300 switches only) 8000 through 160,000,000,000 bps</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Understanding Junos OS CoS Components for EX Series Switches on page 2043</a></li> </ul>                                                                                                                                                                                                                                                                          |

## transmit-rate (EX Series Switches)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | transmit-rate ( <i>rate</i>   percent <i>percentage</i>   remainder);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit class-of-service <a href="#">schedulers</a> <i>scheduler-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Specify the transmit rate or percentage for a scheduler.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>                  | If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent.                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>rate</b> —Transmission rate, in bps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).</p> <p><b>Range:</b> 3200 through 160,000,000,000 bps</p> <p>(EX4300 switches only) 8000 through 160,000,000,000 bps</p> <p><b>percent <i>percentage</i></b> —Percentage of transmission capacity. A percentage of zero drops all packets in the queue.</p> <p><b>Range:</b> 0 through 100 percent</p> <p><b>remainder</b>—Remaining rate available</p> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li><li>• <a href="#">Defining CoS Schedulers and Scheduler Maps (CLI Procedure) on page 2109</a> or <a href="#">Defining CoS Schedulers (J-Web Procedure) on page 2111</a></li><li>• <a href="#">Understanding CoS Schedulers on page 2060</a></li></ul>                                                                                                                                                                                                        |

## unit

|                                 |                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>unit <i>logical-unit-number</i> {     forwarding-class <i>class-name</i>;     classifiers {         (<i>dscp</i>   <i>ieee-802.1</i>   <i>inet-precedence</i>) (<i>classifier-name</i>   default);     } }</pre>                                                                                                   |
| <b>Hierarchy Level</b>          | [edit class-of-service <b>interfaces</b> <i>interface-name</i> ]                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b><i>logical-unit-number</i></b> —Number of the logical unit.</p> <p><b>Range:</b> 0 through 16,385</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Assigning CoS Components to Interfaces (CLI Procedure) on page 2123</a> or <a href="#">Assigning CoS Components to Interfaces (J-Web Procedure) on page 2123</a></li> </ul> |



# Administration

- [Routine Monitoring on page 2157](#)
- [Operational Commands on page 2165](#)

## Routine Monitoring

---

- [Monitoring CoS Classifiers on page 2157](#)
- [Monitoring CoS Forwarding Classes on page 2158](#)
- [Monitoring Interfaces That Have CoS Components on page 2160](#)
- [Monitoring CoS Rewrite Rules on page 2161](#)
- [Monitoring CoS Scheduler Maps on page 2162](#)
- [Monitoring CoS Value Aliases on page 2164](#)
- [Monitoring CoS Drop Profiles on page 2164](#)

## Monitoring CoS Classifiers

**Purpose**



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to display the mapping of incoming CoS values to forwarding class and loss priority for each classifier.

**Action**

To monitor CoS classifiers in the J-Web interface, select **Monitor > Class of Service > Classifiers**.

To monitor CoS classifiers in the CLI, enter the following CLI command:

**show class-of-service classifier**

**Meaning**

[Table 212 on page 2158](#) summarizes key output fields for CoS classifiers.

Table 212: Summary of Key CoS Classifier Output Fields

| Field                      | Values                                                                                                                                                                                                                                                                                     | Additional Information                                      |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------|
| Classifier Name            | Name of a classifier.                                                                                                                                                                                                                                                                      | To display classifier assignments, click the plus sign (+). |
| CoS Value Type             | The classifiers are displayed by type: <ul style="list-style-type: none"> <li>• <b>dscp</b>—All classifiers of the DSCP type.</li> <li>• <b>ieee-802.1</b>—All classifiers of the IEEE 802.1 type.</li> <li>• <b>inet-precedence</b>—All classifiers of the IP precedence type.</li> </ul> |                                                             |
| Index                      | Internal index of the classifier.                                                                                                                                                                                                                                                          |                                                             |
| Incoming CoS Value         | CoS value of the incoming packets, in bits. These values are used for classification.                                                                                                                                                                                                      |                                                             |
| Assign to Forwarding Class | Forwarding class that the classifier assigns to an incoming packet. This class affects the forwarding and scheduling policies that are applied to the packet as it transits the switch.                                                                                                    |                                                             |
| Assign to Loss Priority    | Loss priority value that the classifier assigns to the incoming packet based on its CoS value.                                                                                                                                                                                             |                                                             |

- Related Documentation**
- [Defining CoS Classifiers \(CLI Procedure\) on page 2103](#)
  - [Defining CoS Classifiers \(J-Web Procedure\) on page 2105](#)
  - [Example: Configuring CoS on EX Series Switches on page 2075](#)

## Monitoring CoS Forwarding Classes

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

View the current assignment of class-of-service (CoS) forwarding classes to queues on the switch.

### Action

To monitor CoS forwarding classes in the J-Web interface, select **Monitor > Class of Service > Forwarding Classes**.

To monitor CoS forwarding classes in the CLI, enter the following CLI command:

```
show class-of-service forwarding-class
```

**Meaning** Table 213 on page 2159 summarizes key output fields for CoS forwarding classes.

**Table 213: Summary of Key CoS Forwarding Class Output Fields**

| Field            | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding Class | <p>Names of forwarding classes assigned to queue numbers. The following are the default forwarding classes:</p> <ul style="list-style-type: none"> <li>• <b>best-effort</b>—Provides no special CoS handling of packets. Loss priority is typically not carried in a CoS value.</li> <li>• <b>expedited-forwarding</b>—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service.</li> <li>• <b>assured-forwarding</b>—Provides high assurance for packets within the specified service profile. Excess packets are dropped.</li> <li>• <b>network-control</b>—Packets can be delayed but not dropped.</li> </ul> <p>EX8200 switches have the following additional default forwarding classes:</p> <ul style="list-style-type: none"> <li>• <b>mcast-be</b>—Provides no special CoS handling of packets.</li> <li>• <b>mcast-ef</b>—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service.</li> <li>• <b>mcast-af</b>—Provides high assurance for packets within the specified service profile. Excess packets are dropped.</li> </ul> <p>EX4300 switches supports all the forwarding classes mentioned above and the one mentioned in this section:</p> <ul style="list-style-type: none"> <li>• <b>mcast-nc</b>—Provides multicast network-control traffic.</li> </ul> |
| Queue            | <p>Queue number corresponding to the forwarding class name. The default forwarding classes are assigned as follows:</p> <ul style="list-style-type: none"> <li>• <b>best-effort</b>—0</li> <li>• <b>expedited-forwarding</b>—5</li> <li>• <b>assured-forwarding</b>—1</li> <li>• <b>network-control</b>—7</li> <li>• <b>mcast-be</b>—2</li> <li>• <b>mcast-ef</b>—4</li> <li>• <b>mcast-af</b>—6</li> </ul> <p>EX4300 switches have the following queue numbers for the forwarding classes:</p> <ul style="list-style-type: none"> <li>• <b>best-effort</b>—0</li> <li>• <b>expedited-forwarding</b>—1</li> <li>• <b>assured-forwarding</b>—2</li> <li>• <b>network-control</b>—3</li> <li>• <b>mcast-be</b>—8</li> <li>• <b>mcast-ef</b>—9</li> <li>• <b>mcast-af</b>—10</li> <li>• <b>mcast-nc</b>—11</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Fabric Priority  | <p>(EX8200 switches only) Fabric priority for the forwarding class, either <b>high</b> or <b>low</b>. The fabric priority determines the priority of packets entering the switch fabric.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- Related Documentation**
- [Defining CoS Forwarding Classes \(CLI Procedure\) on page 2107](#)
  - [Defining CoS Forwarding Classes \(J-Web Procedure\) on page 2107](#)
  - [Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards \(CLI Procedure\) on page 2126](#)
  - [Example: Configuring CoS on EX Series Switches on page 2075](#)

## Monitoring Interfaces That Have CoS Components

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to display details about the physical and logical interfaces and the CoS components assigned to them.

### Action

To monitor interfaces that have CoS components in the J-Web interface, select **Monitor > Class of Service > Interface Association**.

To monitor interfaces that have CoS components in the CLI, enter the following command:

**show class-of-service interface *interface***

### Meaning

[Table 214 on page 2160](#) summarizes key output fields for CoS interfaces.

**Table 214: Summary of Key CoS Interfaces Output Fields**

| Field             | Values                                                                                            | Additional Information                                                                                 |
|-------------------|---------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| Interface         | Name of a physical interface to which CoS components are assigned.                                | To display names of logical interfaces configured on this physical interface, click the plus sign (+). |
| Scheduler Map     | Name of the scheduler map associated with this interface.                                         |                                                                                                        |
| Queues Supported  | Number of queues you can configure on the interface.                                              |                                                                                                        |
| Queues in Use     | Number of queues currently configured.                                                            |                                                                                                        |
| Logical Interface | Name of a logical interface on the physical interface to which CoS components are assigned.       |                                                                                                        |
| Object            | Category of an object—for example, <b>classifier</b> , <b>scheduler-map</b> , or <b>rewrite</b> . |                                                                                                        |
| Name              | Name that you have given to an object—for example, <b>ba-classifier</b> .                         |                                                                                                        |



Table 214: Summary of Key CoS Interfaces Output Fields (*continued*)

| Field | Values                                                              | Additional Information |
|-------|---------------------------------------------------------------------|------------------------|
| Type  | Type of an object—for example, <b>dscp</b> for a classifier.        |                        |
| Index | Index of this interface or the internal index of a specific object. |                        |

**Related Documentation**

- [Assigning CoS Components to Interfaces \(CLI Procedure\) on page 2123](#)
- [Assigning CoS Components to Interfaces \(J-Web Procedure\) on page 2123](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)

## Monitoring CoS Rewrite Rules

**Purpose**



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to display information about CoS value rewrite rules, which are based on the forwarding class and loss priority.

**Action**

To monitor CoS rewrite rules in the J-Web interface, select **Monitor > Class of Service > Rewrite Rules**.

To monitor CoS rewrite rules in the CLI, enter the following command:

```
show class-of-service rewrite-rules
```

**Meaning**

[Table 215 on page 2161](#) summarizes key output fields for CoS rewrite rules.

Table 215: Summary of Key CoS Rewrite Rules Output Fields

| Field             | Values                                                                                                                                                                                                                                                      | Additional Information                                                                             |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| Rewrite Rule Name | Names of rewrite rules.                                                                                                                                                                                                                                     |                                                                                                    |
| CoS Value Type    | Rewrite rule type: <ul style="list-style-type: none"> <li>• <b>dscp</b>—For IPv4 DiffServ traffic.</li> <li>• <b>exp</b>—For MPLS traffic.</li> <li>• <b>ieee-802.1</b>—For Layer 2 traffic.</li> <li>• <b>inet-precedence</b>—For IPv4 traffic.</li> </ul> | To display forwarding classes, loss priorities, and rewritten CoS values, click the plus sign (+). |
| Index             | Internal index for this particular rewrite rule.                                                                                                                                                                                                            |                                                                                                    |

Table 215: Summary of Key CoS Rewrite Rules Output Fields (*continued*)

| Field                | Values                                                                                                 | Additional Information                                                                                           |
|----------------------|--------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| Forwarding Class     | Forwarding class that is used to determine CoS values for rewriting in combination with loss priority. | Rewrite rules are applied to CoS values in outgoing packets based on forwarding class and loss priority setting. |
| Loss Priority        | Loss priority that is used to determine CoS values for rewriting in combination with forwarding class. |                                                                                                                  |
| Rewrite CoS Value To | Value that the CoS value is rewritten to.                                                              |                                                                                                                  |

- Related Documentation**
- [Defining CoS Rewrite Rules \(CLI Procedure\) on page 2120](#)
  - [Defining CoS Rewrite Rules \(J-Web Procedure\) on page 2121](#)
  - [Example: Configuring CoS on EX Series Switches on page 2075](#)

## Monitoring CoS Scheduler Maps

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to display assignments of CoS forwarding classes to schedulers.

### Action

To monitor CoS scheduler maps in the J-Web interface, select **Monitor > Class of Service > Scheduler Maps**.

To monitor CoS scheduler maps in the CLI, enter the following CLI command:

```
show class-of-service scheduler-map
```

### Meaning

[Table 216 on page 2162](#) summarizes key output fields for CoS scheduler maps.

Table 216: Summary of Key CoS Scheduler Maps Output Fields

| Field            | Values                                                                   | Additional Information                |
|------------------|--------------------------------------------------------------------------|---------------------------------------|
| Scheduler Map    | Name of a scheduler map.                                                 | For details, click the plus sign (+). |
| Index            | Index of a specific object—scheduler maps, schedulers, or drop profiles. |                                       |
| Scheduler Name   | Name of a scheduler.                                                     |                                       |
| Forwarding Class | Forwarding classes this scheduler is assigned to.                        |                                       |

Table 216: Summary of Key CoS Scheduler Maps Output Fields (*continued*)

| Field             | Values                                                                                                                                                                                                                                                                                                                                                                                                         | Additional Information |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| Transmit Rate     | Configured transmit rate of the scheduler in bits per second (bps). The rate value can be either of the following: <ul style="list-style-type: none"> <li>A percentage—The scheduler receives the specified percentage of the total interface bandwidth.</li> <li><b>remainder</b>— The scheduler receives the remaining bandwidth of the interface after bandwidth allocation to other schedulers.</li> </ul> |                        |
| Buffer Size       | Delay buffer size in the queue or the amount of transmit delay (in milliseconds). The buffer size can be either of the following: <ul style="list-style-type: none"> <li>A percentage—The buffer is a percentage of the total buffer allocation.</li> <li><b>remainder</b>—The buffer is sized according to what remains after other scheduler buffer allocations.</li> </ul>                                  |                        |
| Priority          | Scheduling priority of a queue: <ul style="list-style-type: none"> <li><b>strict-high</b>—Packets in this queue are transmitted first.</li> <li><b>low</b>—Packets in this queue are transmitted last.</li> </ul>                                                                                                                                                                                              |                        |
| Excess rate       | The percentage of excess bandwidth traffic to share.                                                                                                                                                                                                                                                                                                                                                           |                        |
| Drop Profiles     | Name and index of a drop profile that is assigned to a specific loss priority and protocol pair.                                                                                                                                                                                                                                                                                                               |                        |
| Loss Priority     | Packet loss priority corresponding to a drop profile.                                                                                                                                                                                                                                                                                                                                                          |                        |
| Protocol          | Transport protocol corresponding to a drop profile.                                                                                                                                                                                                                                                                                                                                                            |                        |
| Drop Profile Name | Name of the drop profile.                                                                                                                                                                                                                                                                                                                                                                                      |                        |
| Index             | Index of a specific object—scheduler maps, schedulers, or drop profiles.                                                                                                                                                                                                                                                                                                                                       |                        |

- Related Documentation**
- [Defining CoS Schedulers and Scheduler Maps \(CLI Procedure\) on page 2109](#)
  - [Defining CoS Schedulers \(J-Web Procedure\) on page 2111](#)
  - [Example: Configuring CoS on EX Series Switches on page 2075](#)

## Monitoring CoS Value Aliases

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to display information about the CoS value aliases that the system is currently using to represent DSCP, IEEE 802.1p, and IPv4 precedence bits.

### Action

To monitor CoS value aliases in the J-Web interface, select **Monitor > Class of Service > CoS Value Aliases**.

To monitor CoS value aliases in the CLI, enter the following command:

```
show class-of-service code-point-aliases
```

### Meaning

Table 217 on page 2164 summarizes key output fields for CoS value aliases.

**Table 217: Summary of Key CoS Value Alias Output Fields**

| Field           | Values                                                                                                                                                                                                                                                                                                                                 | Additional Information                                        |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------|
| CoS Value Type  | Type of the CoS value: <ul style="list-style-type: none"> <li><b>dscp</b>—Examines Layer 3 packet headers for IP packet classification.</li> <li><b>ieee-802.1</b>—Examines Layer 2 packet headers for packet classification.</li> <li><b>inet-precedence</b>—Examines Layer 3 packet headers for IP packet classification.</li> </ul> | To display aliases and bit patterns, click the plus sign (+). |
| CoS Value Alias | Name given to a set of bits—for example, <b>af11</b> is a name for <b>001010</b> bits.                                                                                                                                                                                                                                                 |                                                               |
| CoS Value       | Set of bits associated with an alias.                                                                                                                                                                                                                                                                                                  |                                                               |

### Related Documentation

- [Defining CoS Code-Point Aliases \(CLI Procedure\) on page 2101](#)
- [Defining CoS Code-Point Aliases \(J-Web Procedure\) on page 2101](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)

## Monitoring CoS Drop Profiles

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to view data point information for each CoS random early detection (RED) drop profile on the EX8200 switch.

**Action** To monitor CoS RED drop profiles in the J-Web interface, select **Monitor > Class of Service > RED Drop Profiles**.

To monitor CoS RED drop profiles in the CLI, enter the following CLI command:  
`show class-of-service drop-profile`

**Meaning** [Table 218 on page 2165](#) summarizes the key output fields for CoS RED drop profiles.

**Table 218: Summary of the Key Output Fields for CoS Red Drop Profiles**

| Field                 | Values                                                                                                                                                                                                                                                                                                                                 | Additional Information                                                                             |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| RED Drop Profile Name | Name of the RED drop profile.<br><br>A drop profile consists of pairs of values between 0 and 100, one for queue buffer fill level and the other for drop probability, that determine the relationship between a buffer's fullness and the likelihood it will drop packets.                                                            | To display profile values, click the plus sign (+).                                                |
| Graph RED Profile     | Links to a graph of a RED curve that the system uses to determine the drop probability based on queue buffer fullness.                                                                                                                                                                                                                 | The x axis represents the queue buffer fill level, and the y axis represents the drop probability. |
| Type                  | Type of a specific drop profile: <ul style="list-style-type: none"> <li><b>interpolated</b>—The two coordinates (x and y) of the graph are interpolated to produce a smooth profile.</li> <li><b>segmented</b>—The two coordinates (x and y) of the graph are represented by line fragments to produce a segmented profile.</li> </ul> |                                                                                                    |
| Index                 | Internal index of this drop profile.                                                                                                                                                                                                                                                                                                   |                                                                                                    |
| Fill Level            | Percentage fullness of a buffer queue. This value is the x coordinate of the RED drop profile graph.                                                                                                                                                                                                                                   |                                                                                                    |
| Drop Probability      | Drop probability of a packet corresponding to a specific queue buffer fill level. This value is the y coordinate of the RED drop profile graph.                                                                                                                                                                                        |                                                                                                    |

**Related Documentation**

- [Defining CoS Drop Profiles \(J-Web Procedure\) on page 2118](#)
- [Example: Configuring CoS on EX Series Switches on page 2075](#)

## Operational Commands

- `show class-of-service`
- `show class-of-service classifier`
- `show class-of-service code-point-aliases`
- `show class-of-service drop-profile`

- `show class-of-service forwarding-class`
- `show pfe statistics traffic cpu`
- `show pfe statistics traffic egress-queues`
- `show pfe statistics traffic multicast`

## show class-of-service

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show class-of-service</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>EXP classifiers added in Junos OS Release 10.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Display the class-of-service (CoS) information.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Monitoring CoS Value Aliases on page 2164</a></li> <li>• <a href="#">Monitoring CoS Classifiers on page 2157</a></li> <li>• <a href="#">Monitoring CoS Forwarding Classes on page 2158</a></li> <li>• <a href="#">Monitoring CoS Scheduler Maps on page 2162</a></li> <li>• <a href="#">Monitoring CoS Rewrite Rules on page 2161</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show class-of- service on page 2168</a><br><a href="#">show class-of-service rewrite-rule on page 2171</a>                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | Table 219 on page 2167 lists the output fields for the <b>show class-of-service</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                         |

**Table 219: show class-of-service Output Fields**

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Forwarding class</b> | The forwarding class configuration: <ul style="list-style-type: none"> <li>• <b>Forwarding class</b>—Name of the forwarding class.</li> <li>• <b>ID</b>—Forwarding class ID.</li> <li>• <b>Queue</b>—Queue number.</li> <li>• <b>Fabric Priority</b>—(EX8200 switches only) Fabric priority: either <b>high</b> or <b>low</b>. The fabric priority determines which CoS ingress queues packets are sent to.</li> </ul> | All levels      |
| <b>Code point type</b>  | The type of code-point alias: <ul style="list-style-type: none"> <li>• <b>dscp</b>—Aliases for DiffServ code point (DSCP) values.</li> <li>• <b>ieee-802.1</b>—Aliases for IEEE 802.1p values.</li> <li>• <b>inet-precedence</b>—Aliases for IP precedence values.</li> <li>• <b>exp</b>—Aliases for experimental (EXP) values.</li> </ul>                                                                             | All levels      |
| <b>Alias</b>            | Names given to CoS values.                                                                                                                                                                                                                                                                                                                                                                                             | All levels      |
| <b>Bit pattern</b>      | Set of bits associated with an alias.                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |
| <b>Classifier</b>       | Name of the classifier.                                                                                                                                                                                                                                                                                                                                                                                                | All levels      |

Table 219: show class-of-service Output Fields (*continued*)

| Field Name                | Field Description                                                                                        | Level of Output |
|---------------------------|----------------------------------------------------------------------------------------------------------|-----------------|
| <b>Code point</b>         | Code-point values.                                                                                       | All levels      |
| <b>Loss priority</b>      | Loss priority assigned to specific CoS values and aliases of the classifier.                             | All levels      |
| <b>Rewrite rule</b>       | Name of the rewrite-rule.                                                                                | All levels      |
| <b>Drop profile</b>       | Name of the drop profile.                                                                                | All levels      |
| <b>Type</b>               | Type of drop profile. EX Series switches support only the <b>discrete</b> type of drop profile.          | All levels      |
| <b>Fill level</b>         | Percentage of queue buffer fullness of <i>high</i> packets beyond which <i>high</i> packets are dropped. | All levels      |
| <b>Scheduler</b>          | Name of the scheduler.                                                                                   | All levels      |
| <b>Transmit rate</b>      | Transmission rate of the scheduler.                                                                      | All levels      |
| <b>Excess rate</b>        | Percentage of excess bandwidth traffic to share.                                                         | All levels      |
| <b>Buffer size</b>        | Delay buffer size in the queue.                                                                          | All levels      |
| <b>Drop profiles</b>      | Drop profiles configured for the specified scheduler.                                                    | All levels      |
| <b>Protocol</b>           | Transport protocol corresponding to the drop profile.                                                    | All levels      |
| <b>Name</b>               | Name of the drop profile.                                                                                | All levels      |
| <b>Queues supported</b>   | Number of queues that can be configured on the interface.                                                | All levels      |
| <b>Queues in use</b>      | Number of queues currently configured.                                                                   | All levels      |
| <b>Physical interface</b> | Name of the physical interface.                                                                          | All levels      |
| <b>Scheduler map</b>      | Name of the scheduler map.                                                                               | All levels      |
| <b>Index</b>              | Internal index of a specific object.                                                                     | All levels      |

## Sample Output

### show class-of- service

```

user@switch> show class-of-service
Forwarding class      ID      Queue
best-effort           0        0
expedited-forwarding  1        5
assured-forwarding    2        1
network-control       3        7

```



Code point type: dscp

| Alias | Bit pattern |
|-------|-------------|
| af11  | 001010      |
| af12  | 001100      |
| ...   | ...         |

Code point type: ieee-802.1

| Alias | Bit pattern |
|-------|-------------|
| af11  | 010         |
| ...   | ...         |

Code point type: inet-precedence

| Alias | Bit pattern |
|-------|-------------|
| af11  | 001         |
| ...   | ...         |

Classifier: dscp-default, Code point type: dscp, Index: 7

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 000000     | best-effort      | low           |
| 000001     | best-effort      | low           |
| ...        | ...              | ...           |

Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |
| 101        | best-effort      | low           |
| 110        | network-control  | low           |
| 111        | network-control  | low           |

Classifier: ipprec-default, Code point type: inet-precedence, Index: 12

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |
| 101        | best-effort      | low           |
| 110        | network-control  | low           |
| 111        | network-control  | low           |

Classifier: ieee8021p-untrust, Code point type: ieee-802.1, Index: 16

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |
| 101        | best-effort      | low           |
| 110        | best-effort      | low           |
| 111        | best-effort      | low           |

Rewrite rule: dscp-default, Code point type: dscp, Index: 27

| Forwarding class     | Loss priority | Code point |
|----------------------|---------------|------------|
| best-effort          | low           | 000000     |
| best-effort          | high          | 000000     |
| expedited-forwarding | low           | 101110     |

|                      |      |        |
|----------------------|------|--------|
| expedited-forwarding | high | 101110 |
| assured-forwarding   | low  | 001010 |
| assured-forwarding   | high | 001100 |
| network-control      | low  | 110000 |
| network-control      | high | 111000 |

Rewrite rule: ieee8021p-default, Code point type: ieee-802.1, Index: 30

| Forwarding class     | Loss priority | Code point |
|----------------------|---------------|------------|
| best-effort          | low           | 000        |
| best-effort          | high          | 001        |
| expedited-forwarding | low           | 100        |
| expedited-forwarding | high          | 101        |
| assured-forwarding   | low           | 010        |
| assured-forwarding   | high          | 011        |
| network-control      | low           | 110        |
| network-control      | high          | 111        |

Rewrite rule: ipprec-default, Code point type: inet-precedence, Index: 31

| Forwarding class     | Loss priority | Code point |
|----------------------|---------------|------------|
| best-effort          | low           | 000        |
| best-effort          | high          | 000        |
| expedited-forwarding | low           | 101        |
| expedited-forwarding | high          | 101        |
| assured-forwarding   | low           | 001        |
| assured-forwarding   | high          | 001        |
| network-control      | low           | 110        |
| network-control      | high          | 111        |

Drop profile: <default-drop-profile>, Type: discrete, Index: 1

Fill level  
100

Scheduler map: <default>, Index: 2

Scheduler: <default-be>, Forwarding class: best-effort, Index: 20  
Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent,  
Priority: low

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| High          | non-TCP  | 1     | <default-drop-profile> |
| High          | TCP      | 1     | <default-drop-profile> |

Scheduler: <default-nc>, Forwarding class: network-control, Index: 22  
Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,  
Priority: low

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| High          | non-TCP  | 1     | <default-drop-profile> |
| High          | TCP      | 1     | <default-drop-profile> |

Physical interface: ge-0/0/0, Index: 129

Queues supported: 8, Queues in use: 4

Scheduler map: <default>, Index: 2

Physical interface: ge-0/0/1, Index: 130

Queues supported: 8, Queues in use: 4

Scheduler map: <default>, Index: 2

... ..

Fabric priority: low

Scheduler: <default-fabric>, Index: 23

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| High          | non-TCP  | 1     | <default-drop-profile> |
| High          | TCP      | 1     | <default-drop-profile> |

Fabric priority: high

Scheduler: <default-fabric>, Index: 23

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| High          | non-TCP  | 1     | <default-drop-profile> |
| High          | TCP      | 1     | <default-drop-profile> |

### show class-of-service rewrite-rule

user@switch> show class-of-service rewrite-rule

Rewrite rule: dscp-default, Code point type: dscp, Index: 31

| Forwarding class     | Loss priority | Code point |
|----------------------|---------------|------------|
| best-effort          | low           | 000000     |
| best-effort          | high          | 000000     |
| expedited-forwarding | low           | 101110     |
| expedited-forwarding | high          | 101110     |
| fw-class             | low           | 001010     |
| fw-class             | high          | 001100     |
| network-control      | low           | 110000     |
| network-control      | high          | 111000     |

Rewrite rule: exp-default, Code point type: exp, Index: 33

| Forwarding class     | Loss priority | Code point |
|----------------------|---------------|------------|
| best-effort          | low           | 000        |
| best-effort          | high          | 001        |
| expedited-forwarding | low           | 010        |
| expedited-forwarding | high          | 011        |
| fw-class             | low           | 100        |
| fw-class             | high          | 101        |
| network-control      | low           | 110        |
| network-control      | high          | 111        |

Rewrite rule: ieee8021p-default, Code point type: ieee-802.1, Index: 34

| Forwarding class     | Loss priority | Code point |
|----------------------|---------------|------------|
| best-effort          | low           | 000        |
| best-effort          | high          | 001        |
| expedited-forwarding | low           | 010        |
| expedited-forwarding | high          | 011        |
| fw-class             | low           | 100        |
| fw-class             | high          | 101        |
| network-control      | low           | 110        |
| network-control      | high          | 111        |

Rewrite rule: ipprec-default, Code point type: inet-precedence, Index: 35

| Forwarding class     | Loss priority | Code point |
|----------------------|---------------|------------|
| best-effort          | low           | 000        |
| best-effort          | high          | 000        |
| expedited-forwarding | low           | 101        |
| expedited-forwarding | high          | 101        |
| fw-class             | low           | 001        |
| fw-class             | high          | 001        |
| network-control      | low           | 110        |
| network-control      | high          | 111        |

## show class-of-service classifier

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show class-of-service classifier &lt;name <i>name</i>&gt; &lt;type dscp   type dscp-ipv6   type exp   type ieee-802.1   type inet-precedence&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | For each class-of-service (CoS) classifier, display the mapping of code point value to forwarding class and loss priority.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>none</b>—Display all classifiers.</p> <p><b>name <i>name</i></b>—(Optional) Display named classifier.</p> <p><b>type dscp</b>—(Optional) Display all classifiers of the Differentiated Services code point (DSCP) type.</p> <p><b>type dscp-ipv6</b>—(Optional) Display all classifiers of the DSCP for IPv6 type.</p> <p><b>type exp</b>—(Optional) Display all classifiers of the MPLS experimental (EXP) type.</p> <p><b>type ieee-802.1</b>—(Optional) Display all classifiers of the ieee-802.1 type.</p> <p><b>type inet-precedence</b>—(Optional) Display all classifiers of the inet-precedence type.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>    | <p><a href="#">show class-of-service classifier type ieee-802.1 on page 2173</a></p> <p><a href="#">show class-of-service classifier type ieee-802.1 (QFX Series) on page 2173</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>            | Table 220 on page 2172 describes the output fields for the <b>show class-of-service classifier</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Table 220: show class-of-service classifier Output Fields**

| Field Name              | Field Description                                                                                                                                                       |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Classifier</b>       | Name of the classifier.                                                                                                                                                 |
| <b>Code point type</b>  | Type of the classifier: <b>exp</b> (not on EX Series switch), <b>dscp</b> , <b>dscp-ipv6</b> (not on EX Series switch), <b>ieee-802.1</b> , or <b>inet-precedence</b> . |
| <b>Index</b>            | Internal index of the classifier.                                                                                                                                       |
| <b>Code point</b>       | Code point value used for classification                                                                                                                                |
| <b>Forwarding class</b> | Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.                                        |

Table 220: show class-of-service classifier Output Fields (*continued*)

| Field Name    | Field Description                                                                                                                                                                                                 |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Loss priority | Loss priority value used for classification. For most platforms, the value is <b>high</b> or <b>low</b> . For some platforms, the value is <b>high</b> , <b>medium-high</b> , <b>medium-low</b> , or <b>low</b> . |

## Sample Output

### show class-of-service classifier type ieee-802.1

```

user@host> show class-of-service classifier type ieee-802.1
Classifier: ieee802.1-default, Code point type: ieee-802.1, Index: 3
Code Point      Forwarding Class      Loss priority
000             best-effort           low
001             best-effort           high
010             expedited-forwarding  low
011             expedited-forwarding  high
100             assured-forwarding    low
101             assured-forwarding    medium-high
110             network-control       low
111             network-control       high

Classifier: users-ieee802.1, Code point type: ieee-802.1
Code point      Forwarding class      Loss priority
100             expedited-forwarding  low

```

### show class-of-service classifier type ieee-802.1 (QFX Series)

```

user@switch> show class-of-service classifier type ieee-802.1
Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
Code point      Forwarding class      Loss priority
000             best-effort           low
001             best-effort           low
010             best-effort           low
011             fcoe                  low
100             no-loss               low
101             best-effort           low
110             network-control       low
111             network-control       low

Classifier: ieee-mcast, Code point type: ieee-802.1, Index: 46
Code point      Forwarding class      Loss priority
000             mcast                 low
001             mcast                 low
010             mcast                 low
011             mcast                 low
100             mcast                 low
101             mcast                 low
110             mcast                 low
111             mcast                 low

```

## show class-of-service code-point-aliases

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show class-of-service code-point-aliases</code><br><code>&lt;dscp   dscp-ipv6   exp   ieee-802.1   inet-precedence&gt;</code>                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>none</b>—Display code point aliases of all code point types.</p> <p><b>dscp</b>—(Optional) Display Differentiated Services code point (DSCP) aliases.</p> <p><b>dscp-ipv6</b>—(Optional) Display IPv6 DSCP aliases.</p> <p><b>exp</b>—(Optional) Display MPLS EXP code point aliases.</p> <p><b>ieee-802.1</b>—(Optional) Display IEEE-802.1 code point aliases.</p> <p><b>inet-precedence</b>—(Optional) Display IPv4 precedence code point aliases.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service code-point-aliases exp on page 2175</a>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 221 on page 2174</a> describes the output fields for the <b>show class-of-service code-point-aliases</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                     |

**Table 221: show class-of-service code-point-aliases Output Fields**

| Field Name             | Field Description                                                                                                                                                                                                           |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Code point type</b> | Type of the code points displayed: <b>dscp</b> , <b>dscp-ipv6</b> (not on EX Series switch), <b>exp</b> (not on EX Series switch or the QFX Series), <b>ieee-802.1</b> , or <b>inet-precedence</b> (not on the QFX Series). |
| <b>Alias</b>           | Alias for a bit pattern.                                                                                                                                                                                                    |
| <b>Bit pattern</b>     | Bit pattern for which the alias is displayed.                                                                                                                                                                               |

## Sample Output

`show class-of-service code-point-aliases exp`

```
user@host> show class-of-service code-point-aliases exp
Code point type: exp
  Alias      Bit pattern
  af11       100
  af12       101
  be         000
  be1        001
  cs6        110
  cs7        111
  ef         010
  ef1        011
  nc1        110
  nc2        111
```

## show class-of-service drop-profile

|                                 |                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show class-of-service drop-profile</code><br><code>&lt;profile-name <i>profile-name</i>&gt;</code>                                                                                                                              |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                              |
| <b>Description</b>              | Display data points for each class-of-service (CoS) random early detection (RED) drop profile.                                                                                                                                        |
| <b>Options</b>                  | <b>none</b> —Display all drop profiles.<br><br><b>profile-name <i>profile-name</i></b> —(Optional) Display the specified profile only.                                                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service drop-profile on page 2177</a><br><a href="#">show class-of-service drop-profile (EX4200 Switch) on page 2177</a><br><a href="#">show class-of-service drop-profile (EX8200 Switch) on page 2177</a> |
| <b>Output Fields</b>            | <a href="#">Table 222 on page 2176</a> describes the output fields for the <b>show class-of-service drop-profile</b> command. Output fields are listed in the approximate order in which they appear.                                 |

**Table 222: show class-of-service drop-profile Output Fields**

| Field Name              | Field Description                                                                                                                                     |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Drop profile</b>     | Name of a drop profile.                                                                                                                               |
| <b>Type</b>             | Type of drop profile: <ul style="list-style-type: none"> <li><b>discrete</b> (default)</li> <li><b>interpolated</b> (EX8200 switches only)</li> </ul> |
| <b>Index</b>            | Internal index of this drop profile.                                                                                                                  |
| <b>Fill Level</b>       | Percentage fullness of a queue.                                                                                                                       |
| <b>Drop probability</b> | Drop probability at this fill level.                                                                                                                  |



## Sample Output

### show class-of-service drop-profile

```

user@host> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
    100         100
Drop profile: user-drop-profile, Type: interpolated, Index: 2989
  Fill level    Drop probability
    0           0
    1           1
    2           2
    4           4
    5           5
    6           6
    8           8
   10          10
   12          15
   14          20
   15          23
... 64 entries total
   90          96
   92          96
   94          97
   95          98
   96          98
   98          99
   99          99
  100          100

```

### show class-of-service drop-profile (EX4200 Switch)

```

user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level
    100
Drop profile: dp1, Type: discrete, Index: 40496
  Fill level
    10

```

### show class-of-service drop-profile (EX8200 Switch)

```

user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
    100         100
Drop profile: dp1, Type: interpolated, Index: 40496
  Fill level    Drop probability
    0           0
    1           80
    2           90
    4           90
    5           90
    6           90
    8           90
   10           90
   12           91
   14           91
   15           91
   16           91

```

|                                                 |                  |
|-------------------------------------------------|------------------|
| 18                                              | 91               |
| 20                                              | 91               |
| 22                                              | 92               |
| 24                                              | 92               |
| 25                                              | 92               |
| 26                                              | 92               |
| 28                                              | 92               |
| 30                                              | 92               |
| 32                                              | 93               |
| 34                                              | 93               |
| 35                                              | 93               |
| 36                                              | 93               |
| 38                                              | 93               |
| 40                                              | 93               |
| 42                                              | 94               |
| 44                                              | 94               |
| 45                                              | 94               |
| 46                                              | 94               |
| 48                                              | 94               |
| 49                                              | 94               |
| 51                                              | 95               |
| 52                                              | 95               |
| 54                                              | 95               |
| 55                                              | 95               |
| 56                                              | 95               |
| 58                                              | 95               |
| 60                                              | 95               |
| 62                                              | 96               |
| 64                                              | 96               |
| 65                                              | 96               |
| 66                                              | 96               |
| 68                                              | 96               |
| 70                                              | 96               |
| 72                                              | 97               |
| 74                                              | 97               |
| 75                                              | 97               |
| 76                                              | 97               |
| 78                                              | 97               |
| 80                                              | 97               |
| 82                                              | 98               |
| 84                                              | 98               |
| 85                                              | 98               |
| 86                                              | 98               |
| 88                                              | 98               |
| 90                                              | 98               |
| 92                                              | 99               |
| 94                                              | 99               |
| 95                                              | 99               |
| 96                                              | 99               |
| 98                                              | 99               |
| 99                                              | 99               |
| 100                                             | 100              |
| Drop profile: dp2, Type: discrete, Index: 40499 |                  |
| Fill level                                      | Drop probability |
| 10                                              | 5                |
| 50                                              | 50               |

## show class-of-service forwarding-class

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show class-of-service forwarding-class                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Display information about forwarding classes, including the mapping of forwarding classes to queue numbers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Example: Configuring Forwarding Classes</a></li> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS)</a></li> <li>• <a href="#">Monitoring CoS Forwarding Classes on page 2158</a></li> <li>• <a href="#">Defining CoS Forwarding Classes (CLI Procedure) on page 2107</a></li> <li>• <a href="#">Defining CoS Forwarding Class Sets</a></li> <li>• <a href="#">Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) on page 2126</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show class-of-service forwarding-class on page 2180</a><br><a href="#">show class-of-service forwarding-class (EX8200 Switch) on page 2180</a><br><a href="#">show class-of-service forwarding-class (QFX Series) on page 2180</a>                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | Table 223 on page 2179 describes the output fields for the <b>show class-of-service forwarding-class</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

**Table 223: show class-of-service forwarding-class Output Fields**

| Field Name               | Field Description                                                                                                                                                  |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Forwarding class</b>  | Name of the forwarding class.                                                                                                                                      |
| <b>ID</b>                | Forwarding class identifier.                                                                                                                                       |
| <b>Queue</b>             | CoS queue mapped to the forwarding class.                                                                                                                          |
| <b>Policing priority</b> | Not supported on EX Series switches or the QFX Series and can be ignored.                                                                                          |
| <b>Fabric priority</b>   | (EX8200 switches only) Fabric priority for the forwarding class, either <b>high</b> or <b>low</b> . Determines the priority of packets entering the switch fabric. |

Table 223: show class-of-service forwarding-class Output Fields (*continued*)

| Field Name     | Field Description                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>No-Loss</b> | <p>(QFX Series only) Packet loss attribute to differentiate lossless forwarding classes from lossy forwarding classes:</p> <ul style="list-style-type: none"> <li>Disabled—Lossless transport is not configured on the forwarding class (packet drop attribute is <b>drop</b>).</li> <li>Enabled—Lossless transport is configured on the forwarding class (packet drop attribute is <b>no-loss</b>).</li> </ul> |

## Sample Output

### show class-of-service forwarding-class

```

user@switch> show class-of-service forwarding-class
Forwarding class      ID      Queue Policing priority
best-effort           0        0      normal
expedited-forwarding  1        5      normal
assured-forwarding    2        1      normal
network-control       3        7      normal

```

## Sample Output

### show class-of-service forwarding-class (EX8200 Switch)

```

user@switch> show class-of-service forwarding-class
Forwarding class      ID      Queue Fabric priority
best-effort           0        0      low
expedited-forwarding  1        5      low
assured-forwarding    2        1      low
network-control       3        7      low
mcast-be              4        2      low
mcast-ef              5        4      low
mcast-af              6        6      low

```

## Sample Output

### show class-of-service forwarding-class (QFX Series)

```

user@switch> show class-of-service forwarding-class
Forwarding class      ID      Queue Policing priority No-Loss
best-effort           0        0      normal      Disabled
fcoe                  1        3      normal      Enabled
no-loss               2        4      normal      Enabled
network-control       3        7      normal      Disabled
mcast                 8        8      normal      Disabled

```

## show pfe statistics traffic cpu

**Syntax** `show pfe statistics traffic cpu <fpc fpc-slot>`

**Release Information** Command introduced in Junos OS Release 9.5 for EX Series switches.

**Description** (On EX8200 switches only) Display count of multidestination packets ingressing from the physical interface to the CPU.



**NOTE:** Multidestination packets include unknown unicast, broadcast, and multicast packets.

**Options** **none**—Displays the count of packets ingressing from all the physical interfaces (line cards) to the CPU.

**fpc fpc-slot**—(Optional) Displays the count of packets ingressing from the physical interface, referred to by the slot number, to the CPU.

On an EX8200 switch, the FPC slot number is the slot number for the line card. Possible values are **0** through **7** on the EX8208 switch and **0** through **15** on the EX8216 switch.

**Required Privilege Level** view

- Related Documentation**
- [show pfe statistics traffic multicast on page 2187](#)
  - [show pfe statistics traffic egress-queues on page 2185](#)
  - [show interfaces queue on page 2905](#)
  - [Monitoring Interface Status and Traffic on page 2835](#)
  - [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)

**List of Sample Output** [show pfe statistics traffic cpu \(EX8208 Switch\) on page 2182](#)

**Output Fields** [Table 224 on page 2181](#) lists the output fields for the `show pfe statistics traffic cpu` command. Output fields are listed in the approximate order in which they appear.

**Table 224: show pfe statistics traffic cpu Output Fields**

| Field Name         | Field Description                       |
|--------------------|-----------------------------------------|
| Queue              | CoS queue number.                       |
| Forwarding classes | Forwarding class name.                  |
| Queued Packets     | Number of packets queued to this queue. |

Table 224: show pfe statistics traffic cpu Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                               |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queued Bytes         | Number of bytes queued to this queue.                                                                                                                                                                                                                                           |
| Packets              | Number of packets transmitted by this queue.                                                                                                                                                                                                                                    |
| Bytes                | Number of bytes transmitted by this queue.                                                                                                                                                                                                                                      |
| Tail-dropped packets | Count of packets dropped at the tail end of the queue because of lack of buffer space.                                                                                                                                                                                          |
| RED-dropped packets  | Number of packets dropped because of Random Early Discard (RED): <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> |
| RED-dropped bytes    | Number of bytes dropped because of Random Early Discard (RED): <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority bytes dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority bytes dropped because of RED.</li> </ul>       |

## Sample Output

### show pfe statistics traffic cpu (EX8208 Switch)

```

user@switch> show pfe statistics traffic cpu

Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : Not Available
    Bytes        : Not Available
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets :                0
    RED-dropped bytes  :                0                0 bps
    Low           :                0                0 bps
    High          :                0                0 bps
    RED-dropped packets :                0                0 pps
    Low           :                0                0 pps
    High          :                0                0 pps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      : Not Available
    Bytes        : Not Available
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets :                0
    RED-dropped bytes  :                0                0 bps
    Low           :                0                0 bps
    High          :                0                0 bps
    RED-dropped packets :                0                0 pps
    Low           :                0                0 pps
    High          :                0                0 pps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      : Not Available

```

```

Bytes          : Not Available
Packets        :                0                0 pps
Bytes          :                0                0 bps
Tail-dropped packets :                0
RED-dropped bytes :                0                0 bps
  Low          :                0                0 bps
  High         :                0                0 bps
RED-dropped packets :                0                0 pps
  Low          :                0                0 pps
  High         :                0                0 pps
Queue: 3, Forwarding classes: network-control
Queued:
Packets        : Not Available
Bytes          : Not Available
Packets        :                0                0 pps
Bytes          :                0                0 bps
Tail-dropped packets :                0
RED-dropped bytes :                0                0 bps
  Low          :                0                0 bps
  High         :                0                0 bps
RED-dropped packets :                0                0 pps
  Low          :                0                0 pps
  High         :                0                0 pps
Queue: 4
Packets        : Not Available
Bytes          : Not Available
Packets        :                0                0 pps
Bytes          :                0                0 bps
Tail-dropped packets :                0
RED-dropped bytes :                0                0 bps
  Low          :                0                0 bps
  High         :                0                0 bps
RED-dropped packets :                0                0 pps
  Low          :                0                0 pps
  High         :                0                0 pps
Queue: 5
Packets        : Not Available
Bytes          : Not Available
Packets        :                0                0 pps
Bytes          :                0                0 bps
Tail-dropped packets :                0
RED-dropped bytes :                0                0 bps
  Low          :                0                0 bps
  High         :                0                0 bps
RED-dropped packets :                0                0 pps
  Low          :                0                0 pps
  High         :                0                0 pps
Queue: 6
Packets        : Not Available
Bytes          : Not Available
Packets        :                0                0 pps
Bytes          :                0                0 bps
Tail-dropped packets :                0
RED-dropped bytes :                0                0 bps
  Low          :                0                0 bps
  High         :                0                0 bps
RED-dropped packets :                0                0 pps
  Low          :                0                0 pps
  High         :                0                0 pps
Queue: 7
Packets        : Not Available

```

|                      |   |               |       |
|----------------------|---|---------------|-------|
| Bytes                | : | Not Available |       |
| Packets              | : | 0             | 0 pps |
| Bytes                | : | 0             | 0 bps |
| Tail-dropped packets | : | 0             |       |
| RED-dropped bytes    | : | 0             | 0 bps |
| Low                  | : | 0             | 0 bps |
| High                 | : | 0             | 0 bps |
| RED-dropped packets  | : | 0             | 0 pps |
| Low                  | : | 0             | 0 pps |
| High                 | : | 0             | 0 pps |



## show pfe statistics traffic egress-queues

**Syntax** `show pfe statistics traffic egress-queues <fpc fpc-slot>`

**Release Information** Command introduced in Junos OS Release 9.5 for EX Series switches.

**Description** (On EX8200 switches only) Display count of multidestination packets dropped on egress ports when the egress queues are oversubscribed due to multidestination traffic.



**NOTE:** Multidestination packets include unknown unicast, broadcast, and multicast packets.

**Options** **none**—Displays count of packets dropped on egress ports of all physical interfaces (line cards) when egress queues are oversubscribed due to multidestination traffic.

**fpc fpc-slot**—(Optional) Displays count of packets dropped on egress ports of the physical interface (line card) referred to by the slot number.



**NOTE:** On an EX8200 switch, the FPC slot number is the slot number for the line card. Possible values are 0 through 7 on the EX8208 switch and 0 through 15 on the EX8216 switch.

**Required Privilege Level** view

**Related Documentation**

- [show pfe statistics traffic cpu on page 2181](#)
- [show pfe statistics traffic multicast on page 2187](#)
- [show interfaces queue on page 2905](#)
- [Monitoring Interface Status and Traffic on page 2835](#)
- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)

**List of Sample Output** [show pfe statistics traffic egress-queues fpc 4 \(EX8208 Switch\) on page 2186](#)

**Output Fields** [Table 225 on page 2185](#) lists the output fields for the **show pfe statistics traffic egress-queues** command. Output fields are listed in the approximate order in which they appear.

**Table 225: show pfe statistics traffic egress-queues Output Fields**

| Field Name           | Field Description                                                             |
|----------------------|-------------------------------------------------------------------------------|
| Tail-dropped packets | Number of arriving packets dropped because the output queue buffers are full. |

## Sample Output

### show pfe statistics traffic egress-queues fpc 4 (EX8208 Switch)

```
user@switch> show pfe statistics traffic egress-queues fpc 4
Tail-dropped packets :                0
```

## show pfe statistics traffic multicast

**Syntax** `show pfe statistics traffic multicast <fpc fpc-slot dev-number>`

**Release Information** Command introduced in Junos OS Release 9.5 for EX Series switches.

**Description** (On EX8200 switches only) Display class-of-service (CoS) queue information for multdestination traffic on a physical interface (line card).



**NOTE:** Multidestination packets include unknown unicast, broadcast, and multicast packets.



**NOTE:** To view statistical information for unicast traffic, use the `show interfaces queue` command.

**Options** `fpc fpc-slot dev-number`—(Optional) Displays class-of-service (CoS) queue information for multdestination traffic on the physical interface (line card) referred to by the slot number and device number.



**NOTE:** On an EX8200 switch, the FPC slot number is the slot number for the line card. Possible values for the FPC slot number are 0 through 7 on the EX8208 switch and 0 through 15 on the EX8216 switch. The value for the device number ranges from 0–5, where 0–4 values correspond to the statistics only from that specific device and the value 5 corresponds to the combined statistics from all the devices in the FPC.

**Required Privilege Level** view

**Related Documentation**

- [show pfe statistics traffic cpu on page 2181](#)
- [show pfe statistics traffic egress-queues on page 2185](#)
- [show interfaces queue on page 2905](#)
- [Monitoring Interface Status and Traffic on page 2835](#)
- [Understanding Junos OS CoS Components for EX Series Switches on page 2043](#)

**List of Sample Output** [show pfe statistics traffic multicast fpc 0 2 \(EX8208 Switch\) on page 2188](#)

**Output Fields** [Table 226 on page 2188](#) lists the output fields for the `show pfe statistics traffic multicast` command. Output fields are listed in the approximate order in which they appear.

Table 226: show pfe statistics traffic multicast Output Fields

| Field Name           | Field Description                                                                                                                                                                                                                                                               |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queue                | CoS queue number.                                                                                                                                                                                                                                                               |
| Forwarding classes   | Forwarding class name.                                                                                                                                                                                                                                                          |
| Queued Packets       | Number of packets queued to this queue.                                                                                                                                                                                                                                         |
| Queued Bytes         | Number of bytes queued to this queue.                                                                                                                                                                                                                                           |
| Packets              | Number of packets transmitted by this queue.                                                                                                                                                                                                                                    |
| Bytes                | Number of bytes transmitted by this queue.                                                                                                                                                                                                                                      |
| Tail-dropped packets | Count of packets dropped at the tail end of the queue because of lack of buffer space.                                                                                                                                                                                          |
| RED-dropped packets  | Number of packets dropped because of Random Early Discard (RED): <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> |
| RED-dropped bytes    | Number of bytes dropped because of Random Early Discard (RED): <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low-loss priority bytes dropped because of RED.</li> <li>• <b>High</b>—Number of high-loss priority bytes dropped because of RED.</li> </ul>       |

## Sample Output

### show pfe statistics traffic multicast fpc 0 2(EX8208 Switch)

```
user@switch> show pfe statistics traffic multicast fpc 0 2
```

```
Queue: 0, Forwarding classes: best-effort
```

```
Queued:
```

```

Packets          : Not Available
Bytes            : Not Available
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets :                0
RED-dropped bytes  :                0                0 bps
  Low             :                0                0 bps
  High            :                0                0 bps
RED-dropped packets :                0                0 pps
  Low             :                0                0 pps
  High            :                0                0 pps
```

```
Queue: 1, Forwarding classes: expedited-forwarding
```

```
Queued:
```

```

Packets          : Not Available
Bytes            : Not Available
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets :                0
RED-dropped bytes  :                0                0 bps
  Low             :                0                0 bps
```

```

      High                :                0                0 bps
    RED-dropped packets :                0                0 pps
      Low                :                0                0 pps
      High                :                0                0 pps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets :                0
    RED-dropped bytes    :                0                0 bps
      Low                :                0                0 bps
      High               :                0                0 bps
    RED-dropped packets :                0                0 pps
      Low                :                0                0 pps
      High               :                0                0 pps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets :                0
    RED-dropped bytes    :                0                0 bps
      Low                :                0                0 bps
      High               :                0                0 bps
    RED-dropped packets :                0                0 pps
      Low                :                0                0 pps
      High               :                0                0 pps
Queue: 4
  Packets              : Not Available
  Bytes                : Not Available
  Packets              :                0                0 pps
  Bytes                :                0                0 bps
  Tail-dropped packets :                0
  RED-dropped bytes    :                0                0 bps
    Low                :                0                0 bps
    High               :                0                0 bps
  RED-dropped packets :                0                0 pps
    Low                :                0                0 pps
    High               :                0                0 pps
Queue: 5
  Packets              : Not Available
  Bytes                : Not Available
  Packets              :                0                0 pps
  Bytes                :                0                0 bps
  Tail-dropped packets :                0
  RED-dropped bytes    :                0                0 bps
    Low                :                0                0 bps
    High               :                0                0 bps
  RED-dropped packets :                0                0 pps
    Low                :                0                0 pps
    High               :                0                0 pps
Queue: 6
  Packets              : Not Available
  Bytes                : Not Available
  Packets              :                0                0 pps
  Bytes                :                0                0 bps
  Tail-dropped packets :                0
  RED-dropped bytes    :                0                0 bps

```

|                      |   |               |       |
|----------------------|---|---------------|-------|
| Low                  | : | 0             | 0 bps |
| High                 | : | 0             | 0 bps |
| RED-dropped packets  | : | 0             | 0 pps |
| Low                  | : | 0             | 0 pps |
| High                 | : | 0             | 0 pps |
| Queue: 7             |   |               |       |
| Packets              | : | Not Available |       |
| Bytes                | : | Not Available |       |
| Packets              | : | 0             | 0 pps |
| Bytes                | : | 0             | 0 bps |
| Tail-dropped packets | : | 0             |       |
| RED-dropped bytes    | : | 0             | 0 bps |
| Low                  | : | 0             | 0 bps |
| High                 | : | 0             | 0 bps |
| RED-dropped packets  | : | 0             | 0 pps |
| Low                  | : | 0             | 0 pps |
| High                 | : | 0             | 0 pps |

# Troubleshooting Procedures

- [Troubleshooting CoS Schedulers on a 40-port SFP+ Line Card in an EX8200 Switch on page 2191](#)
- [Troubleshooting a CoS Classifier Configuration for a TCAM Space Error on page 2192](#)

## Troubleshooting CoS Schedulers on a 40-port SFP+ Line Card in an EX8200 Switch

- Problem** **Description:** After you configure a scheduler map on an interface on the 40-port SFP+ line card, you notice one or both of the following:
- All packets are being dropped on a class-of-service queue configured on the interface.
  - A message in the system log states that the interface is using the default scheduler map, not the scheduler map you configured. For example:
- ```
Sep 19 21:26:50 hostname cosd[907]: COSD_SCHED_MAP_GROUP_CONFLICT:  
Interface xe-5/0/15 cannot be bound to scheduler-map m1. It will be bound to  
default scheduler-map
```
- Cause** The ports in a 40-port SFP+ line card are divided into eight groups, each group comprising five ports. The ports in a port group share 10 gigabits of bandwidth. Because the port groups share bandwidth, only one scheduler map can be active at a time in a port group. If you configure different scheduler maps for different interfaces in a port group, you do not receive an error when you commit the configuration. Instead, default scheduler map becomes the active scheduler map for all interfaces in the port group, and messages in the system log report that the default scheduler map is in use for the affected interfaces. If the default scheduler map does not define a queue, all traffic is dropped on that queue.
- Solution** Check your CoS configuration for the interfaces in the port group. If you have different scheduler maps assigned to different interfaces in the port group:
1. Delete the scheduler map configuration for all interfaces in the port group.
  2. Determine the scheduler map that you want all interfaces in the port group to use.
  3. Assign that scheduler map to at least one interface in the port group. The remaining interfaces in the port group will adopt this scheduler map.



**BEST PRACTICE:** To prevent confusion and future configuration conflicts, explicitly assign the scheduler map to each interface in the port group.

4. After you commit the configuration, verify that the scheduler map is the active scheduler map for the interfaces in the port group by using the **show class-of-service forwarding-table scheduler-map** command.

**Related Documentation**

- *40-port SFP+ Line Card in an EX8200 Switch*
- [Defining CoS Schedulers and Scheduler Maps \(CLI Procedure\) on page 2109](#)
- [Understanding CoS Queues on EX8200 Line Cards That Include Oversubscribed Ports on page 2071](#)

---

## Troubleshooting a CoS Classifier Configuration for a TCAM Space Error

---

**Problem**    **Description:** When a CoS classifier configuration exceeds the amount of available ternary content addressable memory (TCAM) space, the switch returns the following system log message:

```
<number_of_rules_being_added> rules for <filter_name> class <filter_class> will  
not be installed, key: <bind_point>. no space in tcam db(<shared_pool_information>)
```

The switch returns this message during the commit operation if the number of classifiers defined in the CoS configuration or the number of bind points (interfaces) to which classifiers are bound causes the CoS configuration to exceed the amount of available TCAM space. However, the commit operation for the CoS configuration is completed in the CLI module.

**Solution**    When a CoS configuration exceeds the amount of available TCAM table space, you must either define fewer classifiers or bind them to fewer interfaces, or both, so that the space requirements for the CoS configuration do not exceed the available space in TCAM.



To delete classifier definitions and bind points in a CoS configuration, and to apply a new CoS classifier definition to fewer bind points:

1. Delete either the CoS classifier definition or the bind points:

- To delete the CoS classifier definition:

- For behavioral classifiers:

```
[edit class-of-service]
user@switch# delete classifier dscp d1
```

- For multifield classifiers:

```
[edit]
user@switch# delete interfaces ge-3/0/2 unit 0 family ethernet-switching filter input
ipacl
```

This command deletes a multifield classifier defined for a port. Similarly, you can delete a multifield classifier defined for a VLAN or router.

You can also delete terms defined in a single multifield classifier:

```
[edit]
user@switch# delete firewall family inet filter f1 term t1
```

In both these examples (for behavioral and multifield classifiers), the assumption is that too many classifier definitions resulted in the error message.

- To delete the bind points:

```
[edit class-of-service]
user@switch# delete class-of-service interfaces ge-0/0/0
user@switch# delete class-of-service interfaces ge-0/0/1
user@switch# delete class-of-service interfaces ge-0/0/2
user@switch# delete class-of-service interfaces ge-0/0/3
user@switch# delete class-of-service interfaces ge-0/0/4
user@switch# delete class-of-service interfaces ge-0/0/5
user@switch# delete class-of-service interfaces ge-0/0/6
user@switch# delete class-of-service interfaces ge-0/0/7
user@switch# delete class-of-service interfaces ge-0/0/8
```

Here the assumption is that too many bind points (nine) in the configuration resulted in the error message.

2. Commit the operation:

```
[edit]
user@switch# commit
```

3. Define fewer classifiers in the CoS configuration or bind classifiers to fewer interfaces, or both, so that the CoS classifier configuration does not exceed the amount of available TCAM space on the switch:

- To define CoS classifiers:

- For behavioral classifiers:

```
[edit]
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc1 loss-priority
low code-points 000001
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc2 loss-priority
low code-points 000010
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc3 loss-priority
low code-points 000011
```

```
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc4 loss-priority
low code-points 000100
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc5 loss-priority
low code-points 000101
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc6 loss-priority
low code-points 000110
user@switch# set class-of-service classifiers dscp d2 forwarding-class fc7 loss-priority
low code-points 000111
```

- For multifield Classifiers:

```
[edit]
user@switch# set firewall family inet filter f1 term t1 from protocol tcp
user@switch# set firewall family inet filter f1 term t1 then loss-priority high
user@switch# set firewall family inet filter f1 term t1 then forwarding-class best-effort
user@switch# set firewall family inet filter f1 term t2 from protocol udp
user@switch# set firewall family inet filter f1 term t2 then loss-priority high
user@switch# set firewall family inet filter f1 term t2 then forwarding-class
assured-forwarding
user@switch# set firewall family inet filter f1 term t3 from source-port ssh
user@switch# set firewall family inet filter f1 term t3 then loss-priority low
user@switch# set firewall family inet filter f1 term t3 then forwarding-class fc8
user@switch# set class-of-service forwarding-classes best-effort, assured-forwarding,
fc8
```

- To bind classifiers to fewer interfaces:

```
[edit]
user@switch# set class-of-service interfaces ge-0/0/0 unit 0 classifiers dscp d2
user@switch# set class-of-service interfaces ge-0/0/1 unit 0 classifiers dscp d2
user@switch# set class-of-service interfaces ge-0/0/2 unit 0 forwarding-class best-effort
user@switch# set class-of-service interfaces ge-0/0/3 unit 0 forwarding-class
assured-forwarding
user@switch# set class-of-service interfaces ge-0/0/4 unit 0 forwarding-class fc8
```

4. Commit the operation:

```
[edit]
user@switch# commit
```

5. Check system log for an error message. If an error message is not logged, then your classifier configuration has not exceeded the TCAM space limit.

If an error message is logged, then repeat this procedure by defining fewer classifiers or binding classifiers to fewer bind points.

**Related  
Documentation**

- *Understanding CoS Classifiers*
- [Defining CoS Classifiers \(CLI Procedure\) on page 2103](#)

## PART 11

# Device Security

- [Overview on page 2197](#)
- [Configuration on page 2201](#)
- [Administration on page 2233](#)



## CHAPTER 34

# Overview

- [Storm Control Overview on page 2197](#)
- [Unknown Unicast Forwarding Overview on page 2199](#)

## Storm Control Overview

---

- [Understanding Storm Control on Switching Devices on page 2197](#)

## Understanding Storm Control on Switching Devices



**NOTE:** This topic uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switching device is an EX Series switch and runs software that does not support ELS, see *Understanding Storm Control on EX Series Switches*. If your switching device is an EX Series switch and runs software that does support ELS, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

A traffic storm is generated when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own messages on the network. This, in turn, prompts further responses, creating a snowball effect. The LAN is suddenly flooded with packets, creating unnecessary traffic that leads to poor network performance or even a complete loss of network service.

Storm control enables the switching device to monitor traffic levels and to drop broadcast, multicast, and unknown unicast packets when a specified traffic level—called the *storm control level* or *storm control bandwidth*—is exceeded, thus preventing packets from proliferating and degrading the LAN. As an alternative to having the switching device drop packets, you can configure storm control to shut down interfaces or temporarily disable interfaces (see the [action-shutdown](#) statement and the [recovery-timeout](#) statement) when the storm control level is exceeded.



**NOTE:** On Juniper Networks EX4300 Ethernet Switches, the factory default configuration enables storm control on all Layer 2 interfaces, with the storm control level set to 80 percent of the combined broadcast, multicast, and unknown unicast traffic streams.

Storm control is not enabled by default on Juniper Networks EX9200 Ethernet Switches.

Storm control is not enabled by default on Juniper Networks MX Series routers.

You can customize the storm control level for a specific interface by explicitly configuring either bandwidth level or bandwidth percentage.

- **Bandwidth level**—Configures the storm control level as the bandwidth in kilobits per second of the applicable traffic streams on that interface.
- **Bandwidth percentage**—Configures the storm control level as a percentage of the available bandwidth used by the combined applicable traffic streams that are subject to storm control on that interface.



**NOTE:** You cannot configure both bandwidth level and bandwidth percentage for the same interface.

You can disable storm control selectively for broadcast, multicast, or unknown unicast traffic, or any combination of traffic types. When disabling storm control for multicast traffic, you can specify the traffic to be either registered multicast or unregistered multicast. Registered multicast MAC addresses are multicast MAC addresses that are within the range 01-00-5E-00-00-00 through 01-00-5E-7F-FF-FF. This range has been reserved by the Internet Assigned Numbers Association (IANA) for multicast Ethernet addresses. Multicast MAC addresses that are outside this range are called unregistered multicast addresses.

The sending and receiving of broadcast, multicast, and unicast packets are part of normal LAN operation. Therefore, to recognize a storm, you must be able to identify when traffic has reached a level that is abnormal for your LAN. Suspect a storm when operations begin timing out and network response times slow down. As more packets flood the LAN, network users might be unable to access servers or e-mail.

Monitor the level of broadcast, multicast, and unknown unicast traffic in the LAN when it is operating normally. Use this data as a benchmark to determine when traffic levels are too high. Then configure storm control to set the level at which you want the switching device to drop broadcast traffic, multicast traffic, unknown unicast traffic, or two or all three of those traffic types.



**NOTE:** When you configure storm control level on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth or level. For example, if you configure a storm control level of 15,000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30,000 Kbps of combined traffic streams. Traffic might include broadcast, multicast, and unknown unicast traffic, depending upon the configuration.

**Related Documentation**

- [Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches on page 2201](#)
- [Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers](#)
- [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\) on page 2206](#)
- [Configuring or Disabling Storm Control \(CLI Procedure\) on page 2207](#)

## Unknown Unicast Forwarding Overview

- [Understanding Unknown Unicast Forwarding on page 2199](#)

### Understanding Unknown Unicast Forwarding

Unknown unicast traffic consists of unicast packets with unknown destination MAC addresses. By default, the switch floods these unicast packets that are traveling in a VLAN to all interfaces that are members of the VLAN. Forwarding this type of traffic can create unnecessary traffic that leads to poor network performance or even a complete loss of network service. This is known as a traffic storm.

To prevent a storm, you can disable the flooding of unknown unicast packets to all VLAN interfaces by configuring one VLAN or all VLANs to forward all unknown unicast traffic to a specific interface. This channels the unknown unicast traffic to a single interface.

**Related Documentation**

- [Configuring Unknown Unicast Forwarding \(CLI Procedure\)](#)
- [Configuring Unknown Unicast Forwarding \(CLI Procedure\) on page 2205](#)
- [Understanding Storm Control on EX Series Switches](#)
- [Understanding Storm Control on Switching Devices on page 2197](#)
- [Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches](#)
- [Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches on page 2201](#)





# Configuration

- [Configuration Examples on page 2201](#)
- [Configuration Tasks on page 2204](#)
- [Configuration Statements on page 2211](#)

## Configuration Examples

---

- [Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches on page 2201](#)

### Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

Storm control enables you to prevent network outages caused by broadcast storms on the LAN. You can configure storm control on an EX Series switch to rate-limit broadcast traffic, multicast traffic, and unknown unicast traffic at a specified level and to have packets dropped when the specified traffic level is exceeded, thereby preventing packets from proliferating and degrading the LAN.



**NOTE:** On EX4300 switches, the factory default configuration enables storm control on all Layer 2 interfaces, with the storm control level set to 80 percent of the available bandwidth used by the applicable traffic streams on that interface.

This example shows how to configure storm control on an EX Series switch running Junos OS with ELS.

- [Requirements on page 2202](#)
- [Overview and Topology on page 2202](#)

- [Configuration on page 2202](#)
- [Verification on page 2203](#)

## Requirements

---

This example uses the following hardware and software components:

- One EX Series switch running Junos OS with ELS
- Junos OS Release 13.2 or later for EX Series switches

## Overview and Topology

---

A storm is generated when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own messages on the network. This, in turn, prompts further responses, creating a snowball effect and resulting in a broadcast storm that can cause network outages.

You can use storm control to prevent broadcast storms by specifying the amount, also known as the *storm control level*, of broadcast traffic, multicast traffic, and unknown unicast traffic to be allowed on an interface. You specify the storm control level as the traffic rate in kilobits per second (Kbps) of the combined applicable traffic streams or as the percentage of available bandwidth used by the combined applicable traffic streams.

Storm control monitors the level of applicable incoming traffic and compares it with the level that you specify. If the combined level of the applicable traffic exceeds the specified level, the switch drops packets for the controlled traffic types. As an alternative to having the switch drop packets, you can configure storm control to shut down interfaces or temporarily disable interfaces (see the [action-shutdown](#) statement or the [recovery-timeout](#) statement) when the storm control level is exceeded.

The topology used in this example consists of one switch connected to various network devices. This example shows how to configure the storm control level on interface ge-0/0/0 by setting the level to a traffic rate of 15,000 Kbps, based on the traffic rate of the combined applicable traffic streams. If the combined traffic exceeds this level, the switch drops packets for the controlled traffic types to prevent a network outage.

## Configuration

---

### CLI Quick Configuration

To quickly configure storm control based on the traffic rate in Kbps of the combined traffic streams, copy the following command and paste it into the switch terminal window:

```
[edit]
set forwarding-options storm-control-profiles sc all bandwidth-level 15000
set interfaces ge-0/0/0 unit 0 family ethernet-switching storm-control sc
```

### Step-by-Step Procedure

To configure storm control:

1. Configure a storm control profile, **sc**, and specify the traffic rate in Kbps of the combined traffic streams:  
  
[edit]  
user@switch> set forwarding-options storm-control-profiles sc all bandwidth-level 15000
2. Bind the storm control profile, **sc**, to a logical interface:

```
[edit]
user@switch> set interfaces ge-0/0/0 unit 0 family ethernet-switching storm-control sc
```

**Results** Display the results of the configuration:

```
[edit forwarding-options]
user@switch> show storm-control-profiles sc
all {
    bandwidth 15000;
}

[edit]
user@switch> show interfaces ge-0/0/0
unit 0 {
    family ethernet-switching {
        vlan {
            members default;
        }
    }
    storm-control sc;
}
}
```

### Verification

#### *Verifying That the Storm Control Configuration Is in Effect*

**Purpose** Confirm that storm control is limiting the rate of traffic on the interface.

**Action** Use the **show interfaces ge-0/0/0 detail** operational mode command to view traffic statistics on the storm controlled interface. The input rate (bps) must not exceed the storm control limit.

```
user@switch> show interfaces ge-0/0/0 detail
Physical interface: ge-0/0/0, Enabled, Physical link is Up
Interface index: 160, SNMP ifIndex: 503, Generation: 163
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: b0:c6:9a:67:90:84, Hardware address: b0:c6:9a:67:90:84
Last flapped   : 2013-05-16 22:46:42 UTC (14w3d 03:13 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   :          312742788          512 bps
Output bytes  :          245552919           0 bps
Input packets :          3550009          1 pps
Output packets:          2622101           0 pps
IPv6 transit statistics:
Input bytes   :                   0
Output bytes  :                   0
Input packets :                   0
Output packets:                   0
Egress queues: 8 supported, 4 in use
```

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	1	0
1 assured-fow	0	0	0
5 expedited-fo	0	0	0
7 network-cont	0	2622100	0

Queue number:	Mapped forwarding classes
0	best-effort
1	assured-forwarding
5	expedited-forwarding
7	network-control

Active alarms : None  
Active defects : None  
Interface transmit statistics: Disabled

**Meaning** The **Input bytes** field shows the ingress traffic rate in bytes per second (bps). The input rate is within the storm control limit of 15,000 Kbps.

- Related Documentation**
- [Configuring or Disabling Storm Control \(CLI Procedure\) on page 2207](#)
  - [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\) on page 2206](#)
  - [Understanding Storm Control on Switching Devices on page 2197](#)

---

## Configuration Tasks

- [Configuring Unknown Unicast Forwarding \(CLI Procedure\) on page 2205](#)
- [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\) on page 2206](#)
- [Configuring or Disabling Storm Control \(CLI Procedure\) on page 2207](#)

## Configuring Unknown Unicast Forwarding (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches or QFX Series with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see *Configuring Unknown Unicast Forwarding (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3

Unknown unicast traffic consists of packets with unknown destination MAC addresses. By default, the switch floods these packets to all interfaces associated with a VLAN. Forwarding such traffic to interfaces on the switch can create a security issue.

To prevent flooding unknown unicast traffic across the switch, configure unknown unicast forwarding to direct all unknown unicast packets within a VLAN out to a specific interface. You can configure each VLAN to divert unknown unicast traffic to different interfaces or use one interface for multiple VLANs.

To configure unknown unicast forwarding options:

- Configure unknown unicast forwarding for a specific VLAN (here, the VLAN name is employee), and specify the interface to which all unknown unicast traffic will be forwarded:

```
[edit switch-options]
user@switch# set unknown-unicast-forwarding vlan vlan-name interface ge-x/y/z.0
```

### Related Documentation

- [Verifying That Unknown Unicast Packets Are Forwarded to a Single Interface on page 2233](#)
- [Understanding Unknown Unicast Forwarding on page 2199](#)

## Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)



**NOTE:** This example uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switching device is an EX Series switch and runs software that does not support ELS, see *Understanding Storm Control on EX Series Switches*. If your switching device is an EX Series switch and runs software that does support ELS, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

An Ethernet switching access interface on a switching device might shut down or be disabled as a result of one of the following port-security or storm-control configurations:

- MAC limiting—(Not supported for MX Series routers) The **mac-limit** statement is configured with the **action-shutdown** statement.
- MAC move limiting—(Not supported on EX9200 switches or MX Series routers) The **mac-move-limit** statement is configured with the **action-shutdown** statement.
- Storm control—The **storm-control** statement is configured with the **action-shutdown** statement.

You can configure the switching device to automatically restore the disabled interfaces to service after a specified period of time. The specified time configured in the **recovery-timeout** statement applies to all the interfaces that have been disabled due to MAC limiting, MAC move limiting, or storm control errors.



**NOTE:** To enable autorecovery, specify the recovery timeout value for the interfaces to recover automatically. There is no default recovery timeout. If you do not specify a timeout value, you need to use the [clear ethernet-switching recovery-timeout](#) command for EX Series switches and the [clear bridge recovery-timeout](#) command for MX Series routers to clear the errors and restore the interfaces to service.

To specify the recovery timeout period for the interface:

- Set the **recovery-timeout** statement.

For EX Series switches:

```
[edit interfaces interface-name family unit 0 ethernet-switching]  
user@switch# set recovery-timeout seconds
```

For MX Series routers:

```
[edit interfaces interface-name family unit 0 bridge]  
user@switch# set recovery-timeout seconds
```

**Related  
Documentation**

- [Configuring MAC Limiting \(CLI Procedure\) on page 2361](#)
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 4568](#)
- [Configuring or Disabling Storm Control \(CLI Procedure\) on page 2207](#)

## Configuring or Disabling Storm Control (CLI Procedure)



**NOTE:** This task uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switching device is an EX Series switch and runs software that does not support ELS, see *Understanding Storm Control on EX Series Switches*. If your switching device is an EX Series switch and runs software that does support ELS, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

On EX4300 switches, the factory default configuration enables storm control on all Layer 2 switch interfaces. The default storm control level is set to 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams.

Storm control is not enabled by default on EX9200 switches or MX Series routers.

You can customize the storm control level for a specific interface. You specify the storm control level as the traffic rate in kilobits per second (Kbps) of the combined traffic streams or as the percentage of available bandwidth used by the combined traffic streams.

You can selectively disable storm control for broadcast, multicast, or unknown unicast traffic on all interfaces or on a specified interface. You can additionally disable storm control on registered or unregistered multicast traffic.

This topic describes the following tasks. In these tasks, you use the **[edit interfaces *interface-name* unit 0 family ethernet-switching]** hierarchy level to bind the storm control profile for EX Series switches and the **[edit interfaces *interface-name* unit 0 family bridge]** hierarchy level to bind the storm control profile for MX Series routers.

- [Configuring Storm Control on page 2208](#)
- [Disabling Storm Control on Broadcast Traffic on page 2208](#)
- [Disabling Storm Control on All Multicast Traffic on page 2209](#)
- [Disabling Storm Control on Registered Multicast Traffic on page 2209](#)
- [Disabling Storm Control on Unregistered Multicast Traffic on page 2209](#)
- [Disabling Storm Control on Unknown Unicast Traffic on page 2210](#)
- [Disabling Storm Control on Multiple Types of Traffic on page 2210](#)

## Configuring Storm Control

---

You can configure storm control for a specific interface. The storm control level can be customized by explicitly configuring either the bandwidth level or the bandwidth percentage.

- **bandwidth-level**—Configures the storm control level as the bandwidth in kilobits per second of the combined traffic streams.
- **bandwidth-percentage**—Configures the storm control level as a percentage of the available bandwidth used by the combined traffic streams.

To configure storm control:

1. Create a storm control profile and set the storm control level as the traffic rate in kilobits per second of the combined traffic streams:

```
[edit forwarding-options]
user@device# set storm-control-profiles profile-name all bandwidth-level kbps
```

2. Bind the storm control profile to a logical interface:

For EX Series switches:

```
[edit]
user@device# set interfaces interface-name unit 0 family ethernet-switching storm-control
profile-name
```

For MX Series routers:

```
[edit]
user@device# set interfaces interface-name unit 0 family bridge storm-control profile-name
```

## Disabling Storm Control on Broadcast Traffic

---

To disable storm control on broadcast traffic:

1. Create a storm control profile with the storm control level set as the traffic rate in kilobits per second of the combined traffic streams and exclude broadcast traffic:

```
[edit forwarding-options]
user@device# set storm-control-profiles profile-name all bandwidth-level kbps no-broadcast
```

2. Bind the storm control profile to a logical interface:

For EX Series switches:

```
[edit]
user@device# set interfaces interface-name unit 0 family ethernet-switching storm-control
profile-name
```

For MX Series routers:

```
[edit]
user@device# set interfaces interface-name unit 0 family bridge storm-control profile-name
```



### Disabling Storm Control on All Multicast Traffic

To disable storm control on all multicast traffic:

1. Create a storm control profile with the storm control level set as the traffic rate in kilobits per second of the combined traffic streams but exclude multicast traffic:  

```
[edit forwarding-options]
user@device# set storm-control-profiles profile-name all bandwidth-level kbps no-multicast
```
2. Bind the storm control profile to a logical interface:

For EX Series switches:

```
[edit]
user@device# set interfaces interface-name unit 0 family ethernet-switching storm-control
profile-name
```

For MX Series routers:

```
[edit]
user@device# set interfaces interface-name unit 0 family bridge storm-control profile-name
```

### Disabling Storm Control on Registered Multicast Traffic

To disable storm control on only registered multicast traffic:

1. Create a storm control profile with the storm control level set as the traffic rate in kilobits per second of the combined traffic streams but exclude registered multicast traffic:

```
[edit forwarding-options]
user@device# set storm-control-profiles profile-name all bandwidth-level kbps
no-registered-multicast
```

2. Bind the storm control profile to a logical interface:

For EX Series switches:

```
[edit]
user@device# set interfaces interface-name unit 0 family ethernet-switching storm-control
profile-name
```

For MX Series routers:

```
[edit]
user@device# set interfaces interface-name unit 0 family bridge storm-control profile-name
```

### Disabling Storm Control on Unregistered Multicast Traffic

To disable storm control on only unregistered multicast traffic:

1. Create a storm control profile with the storm control level set as the traffic rate in kilobits per second of the combined traffic streams but exclude unregistered multicast traffic:

```
[edit forwarding-options]
user@device# set storm-control-profiles profile-name all bandwidth-level kbps
no-unregistered-multicast
```

2. Bind the storm control profile to a logical interface:

For EX Series switches:

```
[edit]
user@device# set interfaces interface-name unit 0 family ethernet-switching storm-control
profile-name
```

For MX Series routers:

```
[edit]
user@device# set interfaces interface-name unit 0 family bridge storm-control profile-name
```

---

### Disabling Storm Control on Unknown Unicast Traffic

To disable storm control on only unknown unicast traffic:

1. Create a storm control profile with the storm control level set as the traffic rate in kilobits per second of the combined traffic streams but exclude unregistered multicast traffic:

```
[edit forwarding-options]
user@device# set storm-control-profiles profile-name all bandwidth-level kbps
no-unknown-unicast
```

2. Bind the storm control profile to a logical interface:

For EX Series switches:

```
[edit]
user@device# set interfaces interface-name unit 0 family ethernet-switching storm-control
profile-name
```

For MX Series routers:

```
[edit]
user@device# set interfaces interface-name unit 0 family bridge storm-control profile-name
```

---

### Disabling Storm Control on Multiple Types of Traffic

To disable storm control on broadcast and multicast traffic:

1. Create a storm control profile with the storm control level set as the traffic rate in kilobits per second of the combined traffic streams but exclude broadcast and multicast traffic:

```
[edit forwarding-options]
user@device# set storm-control-profiles profile-name all bandwidth-level kbps no-broadcast
no-multicast
```

2. Bind the storm control profile to a logical interface:

For EX Series switches:

```
[edit]
user@device# set interfaces interface-name unit 0 family ethernet-switching storm-control
profile-name
```

For MX Series routers:

```
[edit]
user@device# set interfaces interface-name unit 0 family bridge storm-control profile-name
```

#### Related Documentation

- [Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches on page 2201](#)
- [Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers](#)

- [Understanding Storm Control on Switching Devices on page 2197](#)

## Configuration Statements

---

- [\[edit forwarding-options storm-control-profiles\] Configuration Statement Hierarchy for EX Series Switches on page 2211](#)
- [\[edit switch-options\] Configuration Statement Hierarchy on EX Series Switches on page 2212](#)
- [action-shutdown on page 2215](#)
- [bandwidth-level on page 2217](#)
- [bandwidth-percentage on page 2218](#)
- [interface \(Unknown Unicast Forwarding\) on page 2219](#)
- [no-broadcast on page 2220](#)
- [no-multicast on page 2222](#)
- [no-registered-multicast on page 2224](#)
- [no-unknown-unicast on page 2225](#)
- [no-unregistered-multicast on page 2227](#)
- [recovery-timeout on page 2228](#)
- [storm-control on page 2229](#)
- [storm-control-profiles on page 2230](#)
- [unknown-unicast-forwarding on page 2231](#)

### [edit forwarding-options storm-control-profiles] Configuration Statement Hierarchy for EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit forwarding-options storm-control-profiles]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit forwarding-options storm-control-profiles\] Hierarchy Level on page 2212](#)
- [Unsupported Statements in the \[edit forwarding-options storm-control-profiles\] Hierarchy Level on page 2212](#)

### Supported Statements in the [edit forwarding-options storm-control-profiles] Hierarchy Level

---

The following hierarchy shows the [edit forwarding-options storm-control-profiles] configuration statements supported on EX Series switches:

```
forwarding-options {
  storm-control-profiles profile-name{
    action-shutdown;
    all {
      bandwidth-level bandwidth-level;
      bandwidth-percentage bandwidth-percentage;
      no-broadcast;
      no-multicast;
      no-registered-multicast;
      no-unknown-unicast;
      no-unregistered-multicast;
    }
  }
}
```

### Unsupported Statements in the [edit forwarding-options storm-control-profiles] Hierarchy Level

---

All statements in the [edit forwarding-options storm-control-profiles] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit forwarding-options\] Configuration Statement Hierarchy on EX Series Switches on page 2698](#)

## [edit switch-options] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit switch-options] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit switch-options\] Hierarchy Level on page 2213](#)
- [Unsupported Statements in the \[edit switch-options\] Hierarchy Level on page 2213](#)

### Supported Statements in the [edit switch-options] Hierarchy Level

The following hierarchy shows the [edit switch-options] configuration statements supported on EX Series switches:

```
switch-options {
  authentication-whitelist mac-address {
    interface interface-name;
    vlan-assignment (vlan-id | vlan-name);
  }
  interface interface-name {
    interface-mac-limit number {
      packet-action action;
    }
    no-mac-learning;
    persistent-learning
  }
  no-mac-learning;
  redundant-trunk-group {
    group name {
      description text;
      interface interface-name {
        primary;
      }
      preempt-cutover-timer seconds
    }
  }
  unknown-unicast-forwarding {
    vlan (all | vlan-name | vlan-tag) {
      interface interface-name;
    }
  }
  voip {
    interface (all | [interface-name | access-ports]) {
      forwarding-class (assured-forwarding | best-effort | expedited-forwarding | mcast-af
        | mcast-be | mcast-ef | mcast-nc | network-control);
      vlan vlan-name;
    }
  }
}
```

### Unsupported Statements in the [edit switch-options] Hierarchy Level

All statements in the [edit switch-options] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 227: Unsupported [edit switch-options] Configuration Statements on EX Series Switches**

Statement	Hierarchy Level
NOTE: Variables, such as <i>filename</i> , are not shown in the statements or hierarchies.	
port-error-disable	[edit switch-options]

Table 227: Unsupported [edit switch-options] Configuration Statements on EX Series Switches *(continued)*

Statement	Hierarchy Level
disable-timeout	[edit switch-options port-error-disable]

---

## action-shutdown


<b>Syntax</b>	action-shutdown;
<b>Hierarchy Level</b>	<ul style="list-style-type: none"> <li>For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers): [edit forwarding-options <b>storm-control-profiles</b> <i>profile-name</i>]</li> <li>For platforms without ELS: [edit ethernet-switching-options storm-control]</li> </ul>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Hierarchy level [edit forwarding-options] introduced in Junos OS Release 13.2X50-D10. (See “Getting Started with Enhanced Layer 2 Software” on page 3 for information about ELS.)</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p>
<b>Description</b>	<p>Shut down or temporarily disable interfaces when the storm control level is exceeded, as follows:</p> <ul style="list-style-type: none"> <li>If you set both the <b>action-shutdown</b> and the <b>port-error-disable</b> statements, the interfaces are disabled temporarily and recover automatically when the disable timeout expires. (The <b>port-error-disable</b> statement is not available for MX Series routers.)</li> <li>If you set both the <b>action-shutdown</b> and the <b>recovery-timeout</b> statements, the interfaces are disabled temporarily and recover automatically when the recovery timeout expires.</li> <li>If you set the <b>action-shutdown</b> statement and do not specify the <b>port-error-disable</b> statement (the <b>port-error-disable</b> statement is not available for MX Series routers), the interfaces that are enabled for storm control are shut down when the storm control level is exceeded and they do not recover automatically from that port-error condition. You must issue the <b>clear ethernet-switching port-error</b> command to clear the port error and restore the interfaces to service. (The <b>clear ethernet-switching port-error</b> command is not available for MX Series routers.)</li> <li>If you set the <b>action-shutdown</b> statement and do not specify the <b>recovery-timeout</b> statement, the interfaces that are enabled for storm control are shut down when the storm control level is exceeded and they do not recover automatically from that port-error condition. For EX Series switches you must issue the <b>clear ethernet-switching recovery-timeout</b> command and for MX Series routers you must issue the <b>clear bridge recovery-timeout</b> command to clear the port error and restore the interfaces to service.</li> </ul>
<b>Default</b>	The <b>action-shutdown</b> option is not enabled by default. The switching device drops packets for the controlled traffic types if the ingress rate of the combined traffic streams exceeds the specified storm control level. Depending upon the configuration, applicable traffic could include broadcast, unknown unicast, and multicast traffic.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

**Related  
Documentation**

- *port-error-disable*
- *disable-timeout*
- [recovery-timeout on page 2228](#)
- *clear ethernet-switching port-error*
- *clear bridge recovery-timeout*
- [clear ethernet-switching recovery-timeout on page 2235](#)
- *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*
- *Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers*
- [Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches on page 2201](#)
- *Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*
- [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\) on page 2206](#)




## bandwidth-level

<b>Syntax</b>	<code>bandwidth-level <i>kbps</i>;</code>
<b>Hierarchy Level</b>	[edit forwarding-options <a href="#">storm-control-profiles</a> <i>profile-name</i> all]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1 for MX Series routers.
<b>Description</b>	Configure the storm control level as the bandwidth in kilobits per second of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams.
<div>  <p><b>NOTE:</b> When you configure storm control level on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control level of 15,000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30,000 Kbps of combined broadcast, multicast, and unknown unicast traffic.</p> </div>	
<b>Default</b>	<p>On EX4300 switches—If you do not specify the storm control level using either the <b>bandwidth-level</b> or the <b>bandwidth-percentage</b> statements, the storm control level defaults to 80 percent of the available bandwidth used by the combined broadcast, unknown unicast, and multicast traffic streams.</p> <p>On EX9200 switches—Storm control is not enabled by default.</p> <p>On MX Series routers—Storm control is not enabled by default.</p>
<b>Options</b>	<p><b>bandwidth-level <i>kbps</i></b>—Traffic rate in kilobits per second of the combined broadcast, multicast, and unknown unicast traffic streams.</p> <p><b>Range:</b> 100 through 10,000,000</p> <p><b>Default:</b> None</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">bandwidth-percentage on page 2218</a></li> <li>• <a href="#">Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches on page 2201</a></li> <li>• <a href="#">Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers</a></li> <li>• <a href="#">Configuring or Disabling Storm Control (CLI Procedure) on page 2207</a></li> </ul>

## bandwidth-percentage

---

<b>Syntax</b>	<code>bandwidth-percentage <i>percentage</i>;</code>
<b>Hierarchy Level</b>	[edit forwarding-options <a href="#">storm-control-profiles</a> <i>profile-name</i> all]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. Statement introduced in Junos OS Release 14.1 for MX Series routers.
<b>Description</b>	Configure the storm control level as the percentage of available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams on an interface. The storm control level is configured as part of the storm control profile.
<hr/>	
<div> <b>NOTE:</b> When you configure storm control level on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control level of 15,000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30,000 Kbps of combined broadcast, multicast, and unknown unicast traffic.</div> <hr/>	
<b>Default</b>	On EX4300 switches—The storm control level is 80 percent of the available bandwidth used by the combined broadcast, unknown unicast, and multicast traffic streams.  On EX9200 switches—Storm control is not enabled by default.  On MX Series routers—Storm control is not enabled by default.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">bandwidth-level on page 2217</a></li><li>• <a href="#">Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches on page 2201</a></li><li>• <a href="#">Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers</a></li><li>• <a href="#">Configuring or Disabling Storm Control (CLI Procedure) on page 2207</a></li></ul>

## interface (Unknown Unicast Forwarding)

<b>Syntax</b>	<code>interface <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	<ul style="list-style-type: none"> <li>For platforms with ELS: [edit switch-options <b>unknown-unicast-forwarding</b> vlan <i>vlan-name</i>]</li> <li>For platforms without ELS: [edit ethernet-switching-options <b>unknown-unicast-forwarding</b> vlan <i>vlan-name</i>]</li> </ul>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Hierarchy level <b>[edit switch-options]</b> introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)</p>
<b>Description</b>	Specify the interface to which unknown unicast packets will be forwarded.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><code>show vlans</code></li> <li><a href="#">show ethernet-switching table on page 2238</a></li> <li><a href="#">Configuring Unknown Unicast Forwarding (CLI Procedure)</a></li> <li><a href="#">Understanding Unknown Unicast Forwarding on page 2199</a></li> </ul>

## no-broadcast

---

<b>Syntax</b>	no-broadcast;
<b>Hierarchy Level</b>	<ul style="list-style-type: none"><li>For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers): [edit forwarding-options <b>storm-control-profiles</b> <i>profile-name</i> all]</li><li>For platforms without ELS: [edit ethernet-switching-options storm-control interface (all   <i>interface-name</i>)]</li></ul>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>Hierarchy level [edit forwarding-options] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software”</a> on page 3 for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p>
<b>Description</b>	Disable storm control for broadcast traffic for the specified interface or for all interfaces.
<b>Default</b>	<ul style="list-style-type: none"><li>On EX2200, EX3200, EX3300, and EX4200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams. You can selectively disable storm control on broadcast, multicast, or unknown-unicast traffic.</li><li>On EX4300 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. You can selectively disable storm control on any type of traffic.</li><li>On EX4500 and EX8200 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. On EX8200 switches, you can selectively disable storm control on registered multicast traffic, on unregistered multicast traffic, or on both types of multicast traffic.</li><li>On EX6200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams. You can selectively disable storm control for each type of traffic individually.</li><li>On EX9200 switches—Storm control is not enabled by default.</li><li>On MX Series routers—Storm control is not enabled by default.</li></ul>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

**Related  
Documentation**

- *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*
- [Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches on page 2201](#)
- *Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers*
- *Disabling or Enabling Storm Control (CLI Procedure)*
- [Configuring or Disabling Storm Control \(CLI Procedure\) on page 2207](#)

## no-multicast

---

<b>Syntax</b>	no-multicast;
<b>Hierarchy Level</b>	<ul style="list-style-type: none"><li>For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers): [edit forwarding-options <b>storm-control-profiles</b> <i>profile-name</i> all]</li><li>For platforms without ELS: [edit ethernet-switching-options storm-control interface (all   <i>interface-name</i>)]</li></ul>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Hierarchy level [edit forwarding-options] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software”</a> on page 3 for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p>
<b>Description</b>	Disable storm control for all multicast traffic (both registered multicast and unregistered multicast) for the specified interface or for all interfaces.
<b>Default</b>	<ul style="list-style-type: none"><li>On EX2200, EX3200, EX3300, and EX4200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams.</li><li>On EX4300 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. You can selectively disable storm control on any type of traffic.</li><li>On EX4500 and EX8200 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. On EX8200 switches, you can selectively disable storm control on registered multicast traffic, on unregistered multicast traffic, or on both types of multicast traffic.</li><li>On EX6200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams. You can selectively disable storm control for each type of traffic individually.</li><li>On EX9200 switches—Storm control is not enabled by default.</li><li>On MX Series routers—Storm control is not enabled by default.</li></ul>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

- Related Documentation**
- [no-registered-multicast on page 2224](#)
  - [no-unregistered-multicast on page 2227](#)
  - *Disabling or Enabling Storm Control (CLI Procedure)*
  - [Configuring or Disabling Storm Control \(CLI Procedure\) on page 2207](#)

## no-registered-multicast

---

<b>Syntax</b>	no-registered-multicast;
<b>Hierarchy Level</b>	<ul style="list-style-type: none"><li>For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers): [edit forwarding-options <b>storm-control-profiles</b> <i>profile-name</i> all]</li><li>For platforms without ELS: [edit ethernet-switching-options storm-control interface (all   <i>interface-name</i>)]</li></ul>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Hierarchy level [edit forwarding-options] introduced in Junos OS Release 13.2X50-D10. (See “Getting Started with Enhanced Layer 2 Software” on page 3 for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p>
<b>Description</b>	<p>(EX8200 switches only) Disable storm control for registered multicast traffic for the specified interface or for all interfaces.</p> <p>(EX4300 and EX9200 switches only) Exclude storm control for registered multicast traffic from a storm control profile.</p> <p>(MX Series routers only) Exclude storm control for registered multicast traffic from a storm control profile.</p>
<b>Default</b>	<p>EX4300 and EX8200 switches—Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic. The default storm control level is 80 percent of the available bandwidth used by the combined applicable traffic streams.</p> <p>EX9200 switches—Storm control is not enabled by default.</p> <p>MX Series routers—Storm control is not enabled by default.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li><a href="#">no-multicast on page 2222</a></li><li><a href="#">no-unregistered-multicast on page 2227</a></li><li><a href="#">Understanding Storm Control on EX Series Switches</a></li><li><a href="#">Understanding Storm Control on Switching Devices on page 2197</a></li></ul>



## no-unknown-unicast

<b>Syntax</b>	no-unknown-unicast;
<b>Hierarchy Level</b>	<ul style="list-style-type: none"> <li>For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers): [edit forwarding-options <b>storm-control-profiles</b> <i>profile-name</i> all]</li> <li>For platforms without ELS: [edit ethernet-switching-options storm-control interface (all   <i>interface-name</i>)]</li> </ul>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>Hierarchy level [edit forwarding-options] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software”</a> on page 3 for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p>
<b>Description</b>	Disable storm control for unknown unicast traffic for the specified interface or for all interfaces.
<b>Default</b>	<ul style="list-style-type: none"> <li>On EX2200, EX3200, EX3300, and EX4200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams. You can selectively disable storm control on broadcast, multicast, or unknown-unicast traffic.</li> <li>On EX4300 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. You can selectively disable storm control on any type of traffic.</li> <li>On EX4500 and EX8200 switches—The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams. On EX8200 switches, you can selectively disable storm control on registered multicast traffic, on unregistered multicast traffic, or on both types of multicast traffic.</li> <li>On EX6200 switches—Storm control does not apply to multicast traffic by default. The factory default configuration enables storm control on all interfaces at 80 percent of the available bandwidth used by the combined unknown unicast and broadcast traffic streams. You can selectively disable storm control for each type of traffic individually.</li> <li>On EX9200 switches—Storm control is not enabled by default.</li> <li>MX Series routers—Storm control is not enabled by default.</li> </ul>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>


**Related  
Documentation**

- *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*
- [Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches on page 2201](#)
- *Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers*
- *Disabling or Enabling Storm Control (CLI Procedure)*
- [Configuring or Disabling Storm Control \(CLI Procedure\) on page 2207](#)

## no-unregistered-multicast

<b>Syntax</b>	no-unregistered-multicast;
<b>Hierarchy Level</b>	<ul style="list-style-type: none"> <li>For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers): [edit forwarding-options <b>storm-control-profiles</b> <i>profile-name</i> all]</li> <li>For platforms without ELS: [edit ethernet-switching-options storm-control interface (all   <i>interface-name</i>)],</li> </ul>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Hierarchy level [edit forwarding-options] introduced in Junos OS Release 13.2X50-D10. (See “Getting Started with Enhanced Layer 2 Software” on page 3 for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p>
<b>Description</b>	<p>(EX8200 switches only) Disable storm control for unregistered multicast traffic for the specified interface or for all interfaces.</p> <p>(EX4300 and EX9200 switches only) Exclude storm control for unregistered multicast traffic from a storm control profile.</p> <p>(MX Series routers) Exclude storm control for unregistered multicast traffic from a storm control profile.</p>
<b>Default</b>	<p>EX4300 and EX8200 switches—Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic. The default storm control level is 80 percent of the available bandwidth used by the combined applicable traffic streams.</p> <p>EX9200 switches—Storm control is not enabled by default.</p> <p>MX Series routers—Storm control is not enabled by default.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">no-multicast on page 2222</a></li> <li><a href="#">no-registered-multicast on page 2224</a></li> <li><a href="#">Understanding Storm Control on EX Series Switches</a></li> <li><a href="#">Understanding Storm Control on Switching Devices on page 2197</a></li> </ul>

## recovery-timeout

<b>Syntax</b>	<code>recovery-timeout seconds;</code>
<b>Hierarchy Level (EX Series and QFX Series)</b>	[edit interfaces <i>interface-name</i> unit 0 family ethernet-switching]
<b>Hierarchy Level (MX Series)</b>	[edit interfaces <i>interface-name</i> unit 0 family bridge]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1 for the MX Series routers.
<b>Description</b>	<p>Disable rather than block an interface when enforcing MAC limiting, MAC move limiting, or rate-limiting configuration options for shutting down the interface, and allow the interface to recover automatically from the error condition after the specified period of time:</p> <ul style="list-style-type: none"> <li>• If you have enabled MAC limiting with the <b>shutdown</b> option and you enable <b>recovery-timeout</b>, the switch disables (rather than shuts down) the interface when the MAC address limit is reached.</li> <li>• If you have enabled MAC move limiting (not supported on EX9200) with the <b>shutdown</b> option and you enable <b>recovery-timeout</b>, the switch disables (rather than shuts down) the interface when the maximum number of moves to a new interface is reached.</li> <li>• If you have enabled storm control with the <b>action-shutdown</b> option and you enable <b>recovery-timeout</b>, the switch disables (rather than shuts down) the interface when applicable traffic exceeds the specified levels. Depending upon the configuration, applicable traffic could include broadcast, unknown unicast, and multicast traffic.</li> </ul>
	<p> <b>NOTE:</b> The <b>recovery-timeout</b> configuration does not apply to pre-existing error conditions. It impacts only error conditions that are detected after the <b>recovery-timeout</b> statement has been enabled and committed. To clear a pre-existing error condition and restore the interface to service, use the operational mode command <b>clear ethernet-switching recovery-timeout</b> for EX Series and QFX Series and <b>clear bridge recovery-timeout</b> for MX Series routers.</p>
<b>Default</b>	Not enabled.
<b>Options</b>	<p><b>seconds</b>— Number of seconds that the interface remains in a disabled state due to a port error prior to automatic recovery.</p> <p><b>Range:</b> 10 through 3600</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [action-shutdown on page 2215](#)
  - [Configuring MAC Limiting \(CLI Procedure\) on page 2361](#)
  - [Configuring MAC Move Limiting \(CLI Procedure\) on page 4568](#)
  - [Configuring or Disabling Storm Control \(CLI Procedure\) on page 2207](#)

## storm-control

<b>Syntax</b>	<code>storm-control <i>storm-control-profile</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>number</i> family ethernet-switching], [edit interfaces <i>interface-name</i> unit <i>number</i> family bridge]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. Statement introduced in Junos OS Release 14.1 for the MX Series routers.
<b>Description</b>	Bind a storm control profile to a logical interface.  On switches running ELS software, storm control is enabled by default on all switch interfaces at a level of 80 percent of the combined broadcast and unknown unicast streams. (For the equivalent statement for platforms running non-ELS software, see <i>storm-control</i> .)
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches on page 2201</a></li> <li>• <a href="#">Understanding Storm Control on Switching Devices on page 2197</a></li> </ul>

## storm-control-profiles

---

**Syntax**    storm-control-profiles *profile-name* {  
              action-shutdown;  
              all {  
                  bandwidth-level;  
                  bandwidth-percentage;  
                  no-broadcast;  
                  no-multicast;  
                  no-registered-multicast;  
                  no-unknown-unicast;  
                  no-unregistered-multicast;  
              }  
              }

**Hierarchy Level**    [edit forwarding-options]

**Release Information**    Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.  
                              Statement introduced in Junos OS Release 13.2 for the QFX Series.  
                              Statement introduced in Junos OS Release 14.1 for MX Series routers.

**Description**    Configure a storm control profile on a switch or router.  
  
                      The remaining statements are explained separately.

**Required Privilege Level**    system—To view this statement in the configuration.  
                                  system-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches on page 2201](#)
- [Understanding Storm Control on Switching Devices on page 2197](#)

## unknown-unicast-forwarding

<b>Syntax</b>	<pre>unknown-unicast-forwarding {   vlan <i>vlan-name</i> {     interface <i>interface-name</i>;   } }</pre>
<b>Hierarchy Level</b>	<ul style="list-style-type: none"> <li>For platforms with ELS: [edit <a href="#">switch-options</a> on page 2212]</li> <li>For platforms without ELS: [edit <a href="#">ethernet-switching-options</a>]</li> </ul>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Hierarchy level [edit <a href="#">switch-options</a>] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software”</a> on page 3 for information about ELS.)</p>
<b>Description</b>	<p>Configure the switch to forward all unknown unicast packets in a VLAN or on all VLANs to a particular interface.</p> <p>The remaining statements are explained separately.</p>
<b>Default</b>	Unknown unicast packets are flooded to all interfaces that belong to the same VLAN.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">show vlans</a></li> <li><a href="#">show ethernet-switching table</a> on page 2238</li> <li><a href="#">Configuring Unknown Unicast Forwarding (CLI Procedure)</a></li> <li><a href="#">Configuring Unknown Unicast Forwarding (CLI Procedure)</a> on page 2205</li> <li><a href="#">Understanding Unknown Unicast Forwarding</a> on page 2199</li> </ul>





# Administration

- [Routine Monitoring on page 2233](#)
- [Operational Commands on page 2234](#)

## Routine Monitoring

---

- [Verifying That Unknown Unicast Packets Are Forwarded to a Single Interface on page 2233](#)

### Verifying That Unknown Unicast Packets Are Forwarded to a Single Interface

**Purpose** Verify that a VLAN is forwarding all unknown unicast packets (those with unknown destination MAC addresses) to a single interface instead of flooding unknown unicast packets across all interfaces that are members of the same VLAN.



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, See: *Verifying That Unknown Unicast Packets Are Forwarded to a Trunk Interface*. For ELS details see: “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

**Action** Display the forwarding interface for unknown unicast packets for a VLAN (here, the VLAN name is v1):

```
user@switch> show configuration switch-options

unknown-unicast-forwarding {
  vlan v1 {
    interface ge-0/0/7.0;
  }
}
```

**Meaning** The sample output from the **show configuration switch-options** command shows that the unknown unicast forwarding interface for VLAN v1 is interface ge-0/0/7. The **show ethernet-switching table** command shows that an unknown unicast packet is received on interface ge-0/0/3 with the destination MAC address (DMAC) 00:01:09:00:00:00 and the source MAC address (SMAC) of 00:11:09:00:01:00. This shows that the SMAC

of the packet is learned in the normal way (through the interface ge-0/0/3.0), while the DMAC is learned on interface ge-0/0/7.

**Related  
Documentation**

- [Configuring Unknown Unicast Forwarding \(CLI Procedure\) on page 2205](#)

## Operational Commands

---

- [clear ethernet-switching recovery-timeout](#)
- [clear ethernet-switching table](#)
- [show ethernet-switching table](#)


---

## clear ethernet-switching recovery-timeout

---

<b>Syntax</b>	clear ethernet-switching recovery-timeout
<b>Release Information</b>	Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
<b>Description</b>	Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch, and restore the interfaces to service.
<b>Options</b>	<b>none</b> —Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch and restore these interfaces to service.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)</a> on page 2206</li></ul>
<b>Output Fields</b>	This command produces no output.

## clear ethernet-switching table

<b>Syntax</b>	clear ethernet-switching table <interface <i>interface-name</i> > <mac <i>mac-address</i> > <management-vlan> <persistent-mac < <i>interface</i>   <i>mac-address</i> >> <vlan <i>vlan-name</i> >
<b>Syntax (QFX Series)</b>	clear ethernet-switching table <interface <i>interface-name</i> > <mac <i>mac-address</i> > <persistent-mac < <i>interface</i>   <i>mac-address</i> >> <vlan <i>vlan-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<div>  <p><b>NOTE:</b> On a QFabric system, using this command on an FCoE-enabled VLAN when FCoE sessions are active can cause traffic flooding and FCoE traffic drop. The FCoE sessions are not terminated and the traffic reconverges after a short period of time.</p> </div> <p>Clear learned entries, which are media access control (MAC) addresses, in the Ethernet switching table (also called the forwarding database table).</p>
<b>Options</b>	<p><b>none</b>—Clear learned entries in the Ethernet switching table, except for persistent MAC addresses.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear all learned MAC addresses for the specified interface from the Ethernet switching table.</p> <p><b>mac <i>mac-address</i></b>—(Optional) Clear the specified learned MAC address from the Ethernet switching table.</p> <p><b>management-vlan</b>—(Optional) Clear all MAC addresses learned for the management VLAN from the Ethernet switching table. Note that you do not specify a VLAN name because only one management VLAN exists.</p> <p><b>persistent-mac &lt;<i>interface</i>   <i>mac-address</i>&gt;</b>—(Optional) Clear all MAC addresses, including persistent MAC addresses. Use the <b>interface</b> option to clear all MAC addresses on an interface, or use the <b>mac-address</b> option to clear all entries for a specific MAC address.</p> <p>Use this command whenever you move a device in your network that has a persistent MAC address on the switch. If you move the device to another port on the switch and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address and the device will not be able to connect. If the original port is down when you move the device, then the new port</p>

will learn the MAC address and the device can connect—however, unless you cleared the MAC address on the original port, when the port comes back up, the system reinstalls the persistent MAC address in the forwarding table for that port. If this occurs, the address is removed from the new port and the device loses connectivity.

**vlan *vlan-name***—(Optional) Clear all MAC addresses learned for the specified VLAN from the Ethernet switching table.

**Required Privilege Level**

view

**Related Documentation**

- [show ethernet-switching table on page 2238](#)
- *show ethernet-switching table*
- *Verifying That Persistent MAC Learning Is Working Correctly*

**List of Sample Output** [clear ethernet-switching table on page 2237](#)


**Output Fields** This command produces no output.

## Sample Output

[clear ethernet-switching table](#)

```
user@switch> clear ethernet-switching table
```

## show ethernet-switching table

<b>Syntax</b>	<pre>show ethernet-switching table &lt;brief   detail   extensive   summary&gt; &lt;interface <i>interface-name</i>&gt; &lt;management-vlan&gt; &lt;persistent-mac &lt;interface <i>interface-name</i>&gt;&gt; &lt;sort-by (<i>name</i>   <i>tag</i>)&gt; &lt;vlan <i>vlan-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Options <b>summary</b>, <b>management-vlan</b>, and <b>vlan <i>vlan-name</i></b> introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Option <b>sort-by</b> and field name <b>tag</b> introduced in Junos OS Release 10.1 for EX Series switches.</p> <p>Option <b>persistent-mac</b> introduced in Junos OS Release 11.4 for EX Series switches.</p>
<b>Description</b>	<p> <b>NOTE:</b> If your EX Series switch CLI displays different options for the <b>show ethernet-switching table</b> command than the options shown in this document, see <a href="#">show ethernet-switching table</a>.</p> <p>Display the Ethernet switching table.</p>
<b>Options</b>	<p><b>none</b>—(Optional) Display brief information about the Ethernet switching table.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display the Ethernet switching table for a specific interface.</p> <p><b>management-vlan</b>—(Optional) Display the Ethernet switching table for a management VLAN.</p> <p><b>persistent-mac &lt;interface <i>interface-name</i>&gt;</b>—(Optional) Display the persistent MAC addresses learned for all interfaces or a specified interface. You can use this command to view entries that you want to clear for an interface that you intentionally disabled.</p> <p><b>sort-by (<i>name</i>   <i>tag</i>)</b>—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.</p> <p><b>vlan <i>vlan-name</i></b>—(Optional) Display the Ethernet switching table for a specific VLAN.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear ethernet-switching table on page 2236</a></li> <li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch</a></li> </ul>

- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- *Example: Setting Up Q-in-Q Tunneling on EX Series Switches*

**List of Sample Output** [show ethernet-switching table on page 2240](#)  
[show ethernet-switching table brief on page 2240](#)  
[show ethernet-switching table detail on page 2241](#)  
[show ethernet-switching table extensive on page 2241](#)  
[show ethernet-switching table persistent-mac on page 2242](#)  
[show ethernet-switching table persistent-mac interface ge-0/0/16.0 on page 2242](#)

**Output Fields** [Table 228 on page 2239](#) lists the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

**Table 228: show ethernet-switching table Output Fields**

Field Name	Field Description	Level of Output
<b>VLAN</b>	The name of a VLAN.	All levels
<b>Tag</b>	The VLAN ID tag name or number.	<b>extensive</b>
<b>MAC or MAC address</b>	The MAC address associated with the VLAN.	All levels
<b>Type</b>	The type of MAC address. Values are: <ul style="list-style-type: none"> <li>• <b>static</b>—The MAC address is manually created.</li> <li>• <b>learn</b>—The MAC address is learned dynamically from a packet's source MAC address.</li> <li>• <b>flood</b>—The MAC address is unknown and flooded to all members.</li> <li>• <b>persistent</b>—The learned MAC addresses that will persist across restarts of the switch or interface-down events.</li> </ul>	All levels except <b>persistent-mac</b>
<b>Type</b>	The type of MAC address. Values are: <ul style="list-style-type: none"> <li>• <b>installed</b>—addresses that are in the Ethernet switching table.</li> <li>• <b>uninstalled</b>—addresses that could not be installed in the table or were uninstalled in an interface-down event and will be reinstalled in the table when the interface comes back up.</li> </ul>	<b>persistent-mac</b>
<b>Age</b>	The time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels
<b>Interfaces</b>	Interface associated with learned MAC addresses or <b>All-members</b> (flood entry).	All levels
<b>Learned</b>	For learned entries, the time which the entry was added to the Ethernet switching table.	<b>detail, extensive</b>
<b>Nexthop index</b>	The next-hop index number.	<b>detail, extensive</b>

Table 228: show ethernet-switching table Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>persistent-mac</b>	<b>installed</b> indicates MAC addresses that are in the Ethernet switching table and <b>uninstalled</b> indicates MAC addresses that could not be installed in the table or were uninstalled in an interface-down event (and will be reinstalled in the table when the interface comes back up).	

## Sample Output

### show ethernet-switching table

```

user@switch> show ethernet-switching table
Ethernet-switching table: 57 entries, 15 learned, 2 persistent
VLAN      MAC address      Type      Age Interfaces
F2         *                Flood     - All-members
F2         00:00:05:00:00:03 Learn     0 ge-0/0/44.0
F2         00:19:e2:50:7d:e0 Static    - Router
Linux      *                Flood     - All-members
Linux      00:19:e2:50:7d:e0 Static    - Router
Linux      00:30:48:90:54:89 Learn     0 ge-0/0/47.0
T1         *                Flood     - All-members
T1         00:00:05:00:00:01 Persistent 0 ge-0/0/46.0
T1         00:00:5e:00:01:00 Static    - Router
T1         00:19:e2:50:63:e0 Persistent 0 ge-0/0/46.0
T1         00:19:e2:50:7d:e0 Static    - Router
T10        *                Flood     - All-members
T10        00:00:5e:00:01:09 Static    - Router
T10        00:19:e2:50:63:e0 Learn     0 ge-0/0/46.0
T10        00:19:e2:50:7d:e0 Static    - Router
T111       *                Flood     - All-members
T111       00:19:e2:50:63:e0 Learn     0 ge-0/0/15.0
T111       00:19:e2:50:7d:e0 Static    - Router
T111       00:19:e2:50:ac:00 Learn     0 ge-0/0/15.0
T2         *                Flood     - All-members
T2         00:00:5e:00:01:01 Static    - Router
T2         00:19:e2:50:63:e0 Learn     0 ge-0/0/46.0
T2         00:19:e2:50:7d:e0 Static    - Router
T3         *                Flood     - All-members
T3         00:00:5e:00:01:02 Static    - Router
T3         00:19:e2:50:63:e0 Learn     0 ge-0/0/46.0
T3         00:19:e2:50:7d:e0 Static    - Router
T4         *                Flood     - All-members
T4         00:00:5e:00:01:03 Static    - Router
T4         00:19:e2:50:63:e0 Learn     0 ge-0/0/46.0
[output truncated]

```

### show ethernet-switching table brief

```

user@switch> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 15 learned, 2 persistent entries
VLAN      MAC address      Type      Age Interfaces
F2         *                Flood     - All-members
F2         00:00:05:00:00:03 Learn     0 ge-0/0/44.0
F2         00:19:e2:50:7d:e0 Static    - Router
Linux      *                Flood     - All-members
Linux      00:19:e2:50:7d:e0 Static    - Router
Linux      00:30:48:90:54:89 Learn     0 ge-0/0/47.0
T1         *                Flood     - All-members

```



```

T1          00:00:05:00:00:01 Persistent    0 ge-0/0/46.0
T1          00:00:5e:00:01:00 Static        - Router
T1          00:19:e2:50:63:e0 Persistent    0 ge-0/0/46.0
T1          00:19:e2:50:7d:e0 Static        - Router
T10         *                          Flood  - All-members
T10         00:00:5e:00:01:09 Static        - Router
T10         00:19:e2:50:63:e0 Learn        0 ge-0/0/46.0
T10         00:19:e2:50:7d:e0 Static        - Router
T111        *                          Flood  - All-members
T111        00:19:e2:50:63:e0 Learn        0 ge-0/0/15.0
T111        00:19:e2:50:7d:e0 Static        - Router
T111        00:19:e2:50:ac:00 Learn        0 ge-0/0/15.0
T2          *                          Flood  - All-members
T2          00:00:5e:00:01:01 Static        - Router
T2          00:19:e2:50:63:e0 Learn        0 ge-0/0/46.0
T2          00:19:e2:50:7d:e0 Static        - Router
T3          *                          Flood  - All-members
T3          00:00:5e:00:01:02 Static        - Router
T3          00:19:e2:50:63:e0 Learn        0 ge-0/0/46.0
T3          00:19:e2:50:7d:e0 Static        - Router
T4          *                          Flood  - All-members
T4          00:00:5e:00:01:03 Static        - Router
T4          00:19:e2:50:63:e0 Learn        0 ge-0/0/46.0
[output truncated]

```

### show ethernet-switching table detail

```

user@switch> show ethernet-switching table detail
Ethernet-switching table: 5 entries, 2 learned entries
VLAN: default, Tag: 0, MAC: *, Interface: All-members
  Interfaces:
    ge-0/0/11.0, ge-0/0/20.0, ge-0/0/30.0, ge-0/0/36.0, ge-0/0/3.0
  Type: Flood
  Nexthop index: 1307

VLAN: default, Tag: 0, MAC: 00:1f:12:30:b8:83, Interface: ge-0/0/3.0
  Type: Learn, Age: 0, Learned: 20:09:26
  Nexthop index: 1315

VLAN: v1, Tag: 101, MAC: *, Interface: All-members
  Interfaces:
    ge-0/0/31.0
  Type: Flood
  Nexthop index: 1313

VLAN: v1, Tag: 101, MAC: 00:1f:12:30:b8:89, Interface: ge-0/0/31.0
  Type: Learn, Age: 0, Learned: 20:09:25
  Nexthop index: 1312

VLAN: v2, Tag: 102, MAC: *, Interface: All-members
  Interfaces:
    ae0.0
  Type: Flood
  Nexthop index: 1317

```

### show ethernet-switching table extensive

```

user@switch> show ethernet-switching table extensive
Ethernet-switching table: 3 entries, 1 learned, 5 persistent entries

VLAN: v1, Tag: 10, MAC: *, Interface: All-members

```

## Interfaces:

ge-0/0/14.0, ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0, ge-0/0/4.0,  
ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0, ge-0/0/10.0,  
ge-0/0/0.0

Type: Flood

Nexthop index: 567

VLAN: v1, Tag: 10, MAC: 00:21:59:c6:93:22, Interface: Router

Type: Static

Nexthop index: 0

VLAN: v1, Tag: 10, MAC: 00:21:59:c9:9a:4e, Interface: ge-0/0/14.0

Type: Learn, Age: 0, Learned: 18:40:50

Nexthop index: 564

**show ethernet-switching table persistent-mac**

user@switch&gt; show ethernet-switching table persistent-mac

VLAN	MAC address	Type	Interface
default	00:10:94:00:00:02	installed	ge-0/0/42.0
default	00:10:94:00:00:03	installed	ge-0/0/42.0
default	00:10:94:00:00:04	installed	ge-0/0/42.0
default	00:10:94:00:00:05	installed	ge-0/0/42.0
default	00:10:94:00:00:06	installed	ge-0/0/42.0
default	00:10:94:00:05:02	uninstalled	ge-0/0/16.0
default	00:10:94:00:06:03	uninstalled	ge-0/0/16.0
default	00:10:94:00:07:04	uninstalled	ge-0/0/16.0

**show ethernet-switching table persistent-mac interface ge-0/0/16.0**

VLAN	MAC address	Type	Interface
default	00:10:94:00:05:02	uninstalled	ge-0/0/16.0
default	00:10:94:00:06:03	uninstalled	ge-0/0/16.0
default	00:10:94:00:07:04	uninstalled	ge-0/0/16.0

## PART 12

# Ethernet Switching

- [Overview on page 2245](#)
- [Configuration on page 2281](#)
- [Administration on page 2435](#)



## CHAPTER 37

# Overview

- [Bridging and VLANs on page 2245](#)
- [Q-in-Q Tunneling on page 2269](#)
- [Layer 2 Protocol Tunneling on page 2272](#)
- [Redundant Trunk Groups on page 2275](#)
- [Proxy ARP on page 2278](#)

## Bridging and VLANs

---

- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)
- [Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches on page 2254](#)
- [Understanding Private VLANs on page 2258](#)
- [Understanding Virtual Routing Instances on EX Series Switches on page 2263](#)
- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on EX Series Switches on page 2264](#)
- [Understanding MAC Address Aging on page 2267](#)

## Understanding Bridging and VLANs on EX Series Switches

Network switches use Layer 2 bridging protocols to discover the topology of their LAN and to forward traffic toward destinations on the LAN. This topic explains the following concepts regarding bridging and VLANs on Juniper Networks EX Series Ethernet Switches:

- [History of VLANs on page 2246](#)
- [How Bridging of VLAN Traffic Works on page 2246](#)
- [Packets Are Either Tagged or Untagged on page 2247](#)
- [Switch Interface Modes—Access, Trunk, or Tagged Access on page 2248](#)
- [Additional Advantages of Using VLANs on page 2250](#)
- [Maximum VLANs and VLAN Members Per Switch on page 2250](#)
- [A Default VLAN Is Configured on Most Switches on page 2251](#)
- [Assigning Traffic to VLANs on page 2252](#)

- [Forwarding VLAN Traffic on page 2252](#)
- [VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces on page 2252](#)

---

## History of VLANs

Ethernet LANs were originally designed for small, simple networks that primarily carried text. However, over time, the type of data carried by LANs grew to include voice, graphics, and video. This more complex data, when combined with the ever-increasing speed of transmission, eventually became too much of a load for the original Ethernet LAN design. Multiple packet collisions were significantly slowing down the larger LANs.

The IEEE 802.1D-2004 standard helped evolve Ethernet LANs to cope with the higher data and transmission requirements by defining the concept of *transparent bridging* (generally called simply *bridging*). Bridging divides a single physical LAN (now called a single *broadcast domain*) into two or more virtual LANs, or VLANs. Each VLAN is a collection of some of the LAN nodes grouped together to form individual broadcast domains.

When VLANs are grouped logically by function or organization, a significant percentage of data traffic stays within the VLAN. This relieves the load on the LAN because all traffic no longer has to be forwarded to all nodes on the LAN. A VLAN first transmits packets within the VLAN, thereby reducing the number of packets transmitted on the entire LAN. Because packets whose origin and destination are in the same VLAN are forwarded only within the local VLAN, packets that are not destined for the local VLAN are the only ones forwarded to other broadcast domains. This way, bridging and VLANs limit the amount of traffic flowing across the entire LAN by reducing the possible number of collisions and packet retransmissions within VLANs and on the LAN as a whole.

---

## How Bridging of VLAN Traffic Works

Because the objective of the IEEE 802.1D-2004 standard was to reduce traffic and therefore reduce potential transmission collisions for Ethernet, a system was implemented to reuse information. Instead of having a switch go through a location process every time a frame is sent to a node, the transparent bridging protocol allows a switch to record the location of known nodes. When packets are sent to nodes, those destination node locations are stored in address-lookup tables called *Ethernet switching tables*. Before sending a packet, a switch using bridging first consults the switching tables to see if that node has already been located. If the location of a node is known, the frame is sent directly to that node.

Transparent bridging uses five mechanisms to create and maintain Ethernet switching tables on the switch:

- Learning
- Forwarding
- Flooding
- Filtering
- Aging

The key bridging mechanism used by LANs and VLANs is *learning*. When a switch is first connected to an Ethernet LAN or VLAN, it has no information about other nodes on the network. As packets are sent, the switch learns the embedded MAC addresses of the sending nodes and stores them in the Ethernet switching table, along with two other pieces of information—the interface (or port) on which the traffic was received on the destination node and the time the address was learned.

Learning allows switches to then do *forwarding*. By consulting the Ethernet switching table to see whether the table already contains the frame's destination MAC address, switches save time and resources when forwarding packets to the known MAC addresses. If the Ethernet switching table does not contain an entry for an address, the switch uses flooding to learn that address.

*Flooding* finds a particular destination MAC address without using the Ethernet switching table. When traffic originates on the switch and the Ethernet switching table does not yet contain the destination MAC address, the switch first floods the traffic to all other interfaces within the VLAN. When the destination node receives the flooded traffic, it can send an acknowledgment packet back to the switch, allowing it to learn the MAC address of the node and add the address to its Ethernet switching table.

*Filtering*, the fourth bridging mechanism, is how broadcast traffic is limited to the local VLAN whenever possible. As the number of entries in the Ethernet switching table grows, the switch pieces together an increasingly complete picture of the VLAN and the larger LAN—it learns which nodes are in the local VLAN and which are on other network segments. The switch uses this information to filter traffic. Specifically, for traffic whose source and destination MAC addresses are in the local VLAN, filtering prevents the switch from forwarding this traffic to other network segments.

To keep entries in the Ethernet switching table current, the switch uses a fifth bridging mechanism, *aging*. Aging is the reason that the Ethernet switching table entries include timestamps. Each time the switch detects traffic from a MAC address, it updates the timestamp. A timer on the switch periodically checks the timestamp, and if it is older than a user-configured value, the switch removes the node's MAC address from the Ethernet switching table. This aging process eventually flushes unavailable network nodes out of the Ethernet switching table.

### Packets Are Either Tagged or Untagged

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q ID. The VLAN IDs 1 through 4094 can be assigned to VLANs, while VLAN IDs 0 and 4095 are reserved by Junos OS and cannot be assigned.

Ethernet packets include a tag protocol identifier (TPID) EtherType field, which identifies the protocol being transported. When a device within a VLAN generates a packet, this field includes a value of 0x8100, which indicates that the packet is a VLAN-tagged packet. The packet also has a VLAN ID field that includes the unique 802.1Q ID, which identifies the VLAN to which the packet belongs.

In addition to the TPID EtherType value of 0x8100, EX Series switches that run Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style also

support values of 0x88a8 (Provider Bridging and Shortest Path Bridging) and 0x9100 (Q-inQ).

For a simple network that has only a single VLAN, all packets include a default 802.1Q tag, which is the only VLAN membership that does not mark the packet as tagged. These packets are untagged packets.

### Switch Interface Modes—Access, Trunk, or Tagged Access

---

Ports, or interfaces, on a switch operate in one of three modes:

- Access mode
- Trunk mode
- Tagged-access mode

#### **Access Mode**

An interface in access mode connects a switch to a single network device, such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. Access interfaces accept only untagged packets.

By default, when you boot an EX Series switch that runs Junos OS that does not support ELS and use the factory default configuration, or when you boot such a switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode and accept only untagged packets from the VLAN named **default**. You can optionally configure another VLAN and use that VLAN instead of **default**.

On an EX Series switch that runs Junos OS that supports ELS, the VLAN named **default** is not supported. Therefore, on such switches, you must explicitly configure at least one VLAN, even if your network is simple and you want only one broadcast domain to exist. After you assign an interface to a VLAN, the interface functions in access mode.

For EX Series switches that run either type of software, you can also configure a trunk port or interface to accept untagged packets from a user-configured VLAN. For details about this concept (native VLAN), see [“Trunk Mode and Native VLAN” on page 2249](#).

#### **Trunk Mode**

Trunk mode interfaces are generally used to connect switches to one another. Traffic sent between switches can then consist of packets from multiple VLANs, with those packets multiplexed so that they can be sent over the same physical connection. Trunk interfaces usually accept only tagged packets and use the VLAN ID tag to determine both the packets' VLAN origin and VLAN destination.

On an EX Series switch that runs software that does not support ELS, an untagged packet is not recognized on a trunk port unless you configure additional settings on that port.

On an EX Series switch that runs Junos OS that supports ELS, a trunk port recognizes untagged control packets for protocols such as the Link Aggregation Control Protocol (LACP) and the Link Layer Discovery Protocol (LLDP). However, the trunk port does not recognize untagged data packets unless you configure additional settings on that port.



In the rare case where you want untagged packets to be recognized by a trunk port on EX Series switches that run either type of software, you must configure the single VLAN on a trunk port as a *native VLAN*. For more information about native VLANs, see [“Trunk Mode and Native VLAN” on page 2249](#).

### **Trunk Mode and Native VLAN**

On an EX Series switch that runs Junos OS that does not support ELS, a trunk port does not recognize packets that do not include VLAN tags, which are also known as untagged packets. On an EX Series switch that runs Junos OS that supports ELS, a trunk port recognizes untagged control packets, but it does not recognize untagged data packets. With native VLAN configured, untagged packets that a trunk port normally does not recognize are sent over the trunk interface. In a situation where packets pass from a device, such as an IP phone or printer, to a switch in access mode, and you want those packets sent from the switch over a trunk port, use native VLAN mode. Create a native VLAN by configuring a VLAN ID for it, and specify that the trunk port is a member of the native VLAN.

The switch's trunk port will then treat those packets differently than the other tagged packets. For example, if a trunk port has three VLANs, 10, 20, and 30, assigned to it with VLAN 10 being the native VLAN, packets on VLAN 10 that leave the trunk port on the other end have no 802.1Q header (tag).

There is another native VLAN option for EX Series switches that do not support ELS. You can have the switch add and remove tags for untagged packets. To do this, you first configure the single VLAN as a native VLAN on a port attached to a device on the edge. Then, assign a VLAN ID tag to the single native VLAN on the port connected to a device. Last, add the VLAN ID to the trunk port. Now, when the switch receives the untagged packet, it adds the ID you specified and sends and receives the tagged packets on the trunk port configured to accept that VLAN.

### **Tagged-Access Mode**

Only EX Series switches that run Junos OS that does not use the ELS configuration style support tagged-access mode.

Tagged-access mode accommodates cloud computing scenarios, specifically deployments including servers that adhere to the edge virtual bridging (EVB) standard (IEEE 803.1Qbg). See *Understanding Edge Virtual Bridging for Use with VEPA Technology*.

Because several virtual computers can be included on one physical server, the packets generated by one server can contain an aggregation of VLAN packets from different virtual machines on that server. To accommodate this situation, tagged-access mode reflects packets back to the physical server on the same downstream port when the destination address of the packet was learned on that downstream port. Packets are also reflected back to the physical server on the downstream port when the destination has not yet been learned. Therefore, the third interface mode, tagged access, has some characteristics of access mode and some characteristics of trunk mode:

- Like access mode, tagged-access mode connects the switch to an access layer device. Unlike access mode, tagged-access mode is capable of accepting VLAN tagged packets.

- Like trunk mode, tagged-access mode accepts VLAN tagged packets from multiple VLANs. Unlike trunk port interfaces, which are connected at the core/distribution layer, tagged-access port interfaces connect devices at the access layer.

Like trunk mode, tagged-access mode also supports native VLAN.



**NOTE:** Control packets are never reflected back on the downstream port.

---

### Additional Advantages of Using VLANs

In addition to reducing traffic and thereby speeding up the network, VLANs have the following advantages:

- VLANs provide segmentation services traditionally provided by routers in LAN configurations, thereby reducing hardware equipment costs.
- Packets coupled to a VLAN can be reliably identified and sorted into different domains. You can contain broadcasts within parts of the network, thereby freeing up network resources. For example, when a DHCP server is plugged into a switch and starts broadcasting its presence, you can prevent some hosts from accessing it by using VLANs to split up the network.
- For security issues, VLANs provide granular control of the network because each VLAN is identified by a single IP subnetwork. All packets passing in and out of a VLAN are consistently tagged with the VLAN ID of that VLAN, thereby providing easy identification, because a VLAN ID on a packet cannot be altered. (For an EX Series switch that runs Junos OS that does not support ELS, we recommend that you avoid using 1 as a VLAN ID, because that ID is a default value.)
- VLANs react quickly to host relocation—this is also due to the persistent VLAN tag on packets.
- On an Ethernet LAN, all network nodes must be physically connected to the same network. In VLANs, the physical location of nodes is not important—you can group network devices in any way that makes sense for your organization, such as by department or business function, types of network nodes, or physical location.

---

### Maximum VLANs and VLAN Members Per Switch

The number of VLANs supported per switch varies for each switch. Use the configuration-mode command **set vlans *vlan-name* *vlan-id* ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because you have to assign a specific ID number when you create a VLAN—you could overwrite one of the numbers, but you cannot exceed the limit.

You can, however, exceed the recommended VLAN member maximum for a switch.

On an EX Series switch that runs Junos OS that does not support the ELS configuration style, the maximum number of VLAN members allowed on the switch is eight times the maximum number of VLANs that the switch supports ( $\text{vmember limit} = \text{vlan max} * 8$ ). If the configuration of the switch exceeds the recommended VLAN member maximum,

a warning message appears when you commit the configuration. If you commit the configuration despite the warning, the commit succeeds, but there is a risk of the Ethernet switching process (eswd) failing as a result of memory allocation failure.

On an EX Series switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs that the switch supports ( $\text{vmember limit} = \text{vlan max} * 24$ ). If the configuration of the switch exceeds the recommended VLAN member maximum, a warning message appears in the system log (syslog).

### A Default VLAN Is Configured on Most Switches



**NOTE:** EX Series switches that run Junos OS with the ELS configuration style do not support a default VLAN.

Some EX Series switches that run Junos OS that does not support the ELS configuration style are preconfigured with a VLAN named **default** that does not tag packets and operates only with untagged packets. On these switches, each interface already belongs to the VLAN named **default** and all traffic uses this VLAN until you configure more VLANs and assign traffic to those VLANs.

The following EX Series switches that run Junos OS that does not support the ELS are not preconfigured to belong to **default** or any other VLAN:.

- Modular switches, such as the EX8200 switches and EX6200 switches
- Switches that are part of a Virtual Chassis

The reason that these switches are not preconfigured is that the physical configuration in both situations is flexible. There is no way of knowing which line cards have been inserted in either the EX8200 switch or EX6200 switch. There is also no way of knowing which switches are included in the Virtual Chassis. Switch interfaces in these two cases must first be defined as Ethernet switching interfaces. After an interface is defined as an Ethernet switching interface, the default VLAN appears in the output from the ? help and other commands.



**NOTE:** When a Juniper Networks EX4500 Ethernet Switch, EX4200 Ethernet Switch, or EX3300 Ethernet Switch is interconnected with other switches in a Virtual Chassis configuration, each individual switch that is included as a member of the configuration is identified with a member ID. The member ID functions as an FPC slot number. When you are configuring interfaces for a Virtual Chassis configuration, you specify the appropriate member ID (0 through 9) as the slot element of the interface name. The default factory settings for a Virtual Chassis configuration include FPC 0 as a member of the default VLAN because FPC 0 is configured as part of the ethernet-switching family. In order to include FPC 1 through FPC 9 in the default VLAN, add the ethernet-switching family to the configurations for those interfaces.

## Assigning Traffic to VLANs

---

You can assign traffic on any switch to a particular VLAN by referencing either the interface port of the traffic or the MAC addresses of devices sending traffic.

### ***Assign VLAN Traffic According to the Interface Port Source***

This method is most commonly used to assign traffic to VLANs. In this case, you specify that all traffic received on a particular switch interface is assigned to a specific VLAN. You configure this VLAN assignment when you configure the switch, by using either the VLAN number (called a VLAN ID) or by using the VLAN name, which the switch then translates into a numeric VLAN ID. This method is referred to simply as creating a VLAN because it is the most commonly used method.

### ***Assign VLAN Traffic According to the Source MAC Address***

In this case, all traffic received from a specific MAC address is forwarded to a specific egress interface (next hop) on the switch. MAC-based VLANs are either static (named MAC addresses configured one at a time) or dynamic (configured using a RADIUS server).

To configure a static MAC-based VLAN on an EX Series switch that supports ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*. To configure a static MAC-based VLAN on an EX Series switch that does not support ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*.

For information about using 802.1X authentication to authenticate end devices and allow access to dynamic VLANs configured on a RADIUS server, see [“Understanding Dynamic VLANs for 802.1X on EX Series Switches” on page 1837](#). You can optionally implement this feature to offload the manual assignment of VLAN traffic to automated RADIUS server databases.

## Forwarding VLAN Traffic

---

To pass traffic within a VLAN, the switch uses Layer 2 forwarding protocols, including IEEE 802.1Q spanning-tree protocols and Multiple VLAN Registration Protocol (MVRP).

To pass traffic between two VLANs, the switch uses standard Layer 3 routing protocols, such as static routing, OSPF, and RIP. On EX Series switches, the same interfaces that support Layer 2 bridging protocols also support Layer 3 routing protocols, providing multilayer switching.

To pass traffic from a single device on an access port to a switch and then pass those packets on a trunk port, use the native mode configuration previously discussed under [“Trunk Mode” on page 2248](#).

## VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces

---

Traditionally, switches sent traffic to hosts that were part of the same broadcast domain (VLAN) but routers were needed to route traffic from one broadcast domain to another. Also, only routers performed other Layer 3 functions such as traffic engineering.

EX Series switches that run Junos OS that supports the ELS configuration style perform inter-VLAN routing functions using an integrated routing and bridging (IRB) interface named `irb`, while EX Series switches that run Junos OS that does not support ELS perform these functions using a routed VLAN interface (RVI) named `vlan`. These interfaces detect both MAC addresses and IP addresses and route data to Layer 3 interfaces, thereby frequently eliminating the need to have both a switch and a router.

**Related  
Documentation**

- *Understanding Private VLANs on EX Series Switches*
- *Understanding Layer 2 Protocol Tunneling on EX Series Switches*
- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on EX Series Switches on page 2264](#)
- [Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches on page 2254](#)
- *Understanding Edge Virtual Bridging for Use with VEPA Technology*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 2281](#)
- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- *Example: Connecting an Access Switch to a Distribution Switch*
- [Example: Connecting Access Switches to a Distribution Switch on page 2291](#)

## Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches

Virtual LANs (VLANs), by definition, divide a LAN's broadcast environment into isolated virtual broadcast domains, thereby limiting the amount of traffic flowing across the entire LAN and reducing the possible number of collisions and packet retransmissions within the LAN. For example, you might want to create a VLAN that includes the employees in a department and the resources that they use often, such as printers, servers, and so on.

Of course, you also want to allow these employees to communicate with people and resources in other VLANs. To forward packets between VLANs, you traditionally needed a router that connected the VLANs. However, you can also accomplish this forwarding with a switch by configuring one of the following features:

- On Juniper Networks EX Series Ethernet Switches that run Juniper Networks Junos operating system (Junos OS) that supports the Enhanced Layer 2 Software (ELS) configuration style, configure an integrated routing and bridging (IRB) interface.
- On EX Series switches that run Junos OS that does not support ELS, configure a routed VLAN interface (RVI).



**NOTE:** IRB interfaces and RVIs provide the same functionality. Where the functionality for both features is the same, this topic uses the term *these interfaces* to refer collectively to both IRB interfaces and RVIs. Where differences exist between the two features, this topic calls out the IRB interfaces and RVIs separately.


Configuring a switch to route traffic between VLANs reduces complexity and eliminates the costs associated with purchasing, installing, managing, powering, and cooling a router.

These interfaces route only VLAN traffic and work by logically dividing a switch into multiple virtual routing instances, thereby isolating VLAN traffic traveling across the network into virtual segments. These interfaces allow switches to recognize which packets are being sent to another VLAN's MAC addresses—then, packets are bridged (switched) whenever the destination is within the same VLAN and are routed through these interfaces only when necessary. Whenever packets can be switched instead of routed, several layers of processing are eliminated. The switches rely on their Layer 3 capabilities to provide this basic routing between VLANs:

- Two VLANs on the same switch
- Two VLANs on different switches (routing is provided by an intermediary third switch.)

Figure 27 on page 2255 illustrates a switch routing VLAN traffic between two access layer switches using one of these interfaces.

**Figure 27: An IRB Interface or RVI on a Switch Providing Routing Between Two Access Switches**



This topic describes:

- [When Should I Use an IRB Interface or RVI? on page 2255](#)
- [How Does an IRB Interface or RVI Work? on page 2255](#)
- [Creating an IRB Interface or RVI on page 2256](#)
- [Viewing IRB Interface and RVI Statistics on page 2257](#)
- [IRB Interfaces and RVI Functions and Other Technologies on page 2257](#)

### When Should I Use an IRB Interface or RVI?

Configure an IRB interface or an RVI for a VLAN if you need to:

- Allow traffic to be routed between VLANs.
- Provide Layer 3 IP connectivity to the switch.
- Monitor individual VLANs for billing purposes. Service providers often need to monitor traffic for this purpose, but this capability can be useful for enterprises where various groups share the cost of the network.

### How Does an IRB Interface or RVI Work?

For an IRB interface, the switch provides the name `irb`, and for an RVI, the switch provides the name `vlan`. Like all Layer 3 interfaces, these interfaces require a logical unit number with an IP address assigned to it. In fact, to be useful, the implementation of these interfaces in an enterprise with multiple VLANs requires at least two logical units and two IP addresses—you must create units with addresses in each of the subnets associated with the VLANs between which you want traffic to be routed. That is, if you have two

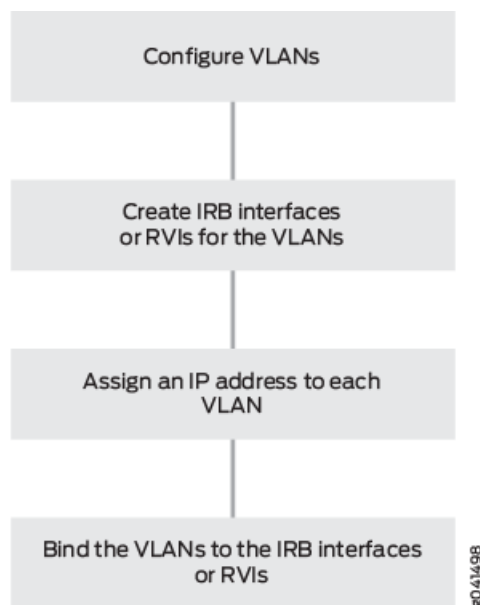
VLANs (for example, VLAN **red** and VLAN **blue**) with corresponding subnets, your interfaces must have a logical unit with an address in the subnet for **red** and a logical unit with an address in the subnet for **blue**. The switch automatically creates direct routes to these subnets and uses these routes to forward traffic between VLANs.

The interface on the switch detects both MAC addresses and IP addresses, then routes data to other Layer 3 interfaces on routers or other switches. These interfaces detect both IPv4 and IPv6 unicast and multicast virtual routing and forwarding (VRF) traffic. Each logical interface can belong to only one routing instance and is further subdivided into logical interfaces, each with a logical interface number appended as a suffix to the names `irb` and `vlan`—for example, `irb.10` and `vlan.10`.

### Creating an IRB Interface or RVI

There are four basic steps in creating an IRB interface or RVI as shown in [Figure 28 on page 2256](#).

**Figure 28: Creating an IRB Interface or RVI**



The following explanations correspond to the four steps for creating a VLAN, as depicted in [Figure 28 on page 2256](#).

- **Configure VLANs**—Virtual LANs are groups of hosts that communicate as if they were attached to the same broadcast stream. VLANs are created with software and do not require a physical router to forward traffic. VLANs are Layer 2 constructs.
- **Create IRB interfaces or RVIs for the VLANs**—The switch's IRB interfaces and RVIs use Layer 3 logical interfaces (unlike routers, which can use either physical or logical interfaces).



- Assign an IP address to each VLAN—An IRB interface or RVI cannot be activated unless it is associated with a physical interface.
- Bind the VLANs to the logical interfaces—There is a one-to-one mapping between a VLAN and an IRB interface or RVI, which means that only one of these interfaces can be mapped to a VLAN.

For specific instructions for creating an IRB interface, see [“Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\)” on page 2340](#), and for an RVI, see [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#).

### Viewing IRB Interface and RVI Statistics

Some switches automatically track IRB interface and RVI traffic statistics. Other switches allow you to configure tracking. [Table 229 on page 2257](#) illustrates the IRB interface- and RVI-tracking capability on various switches.

**Table 229: Tracking IRB Interface and RVI Usage**

Switch	Input (ingress)	Output (Egress)
EX4300	Automatic	Automatic
EX3200, EX4200	Automatic	–
EX8200	Configurable	Automatic
EX2200, EX3300, EX4500, EX6200	–	–

You can view input (ingress) and output (egress) totals with the following commands:

- For IRB interfaces, use the **show interfaces irb extensive** command. Look at the input and output values in the Transit Statistics field for IRB interface activity values.
- For RVI, use the **show interfaces vlan extensive** command. Look at the input and output values in the Logical Interface Transit Statistics field for RVI activity values.

### IRB Interfaces and RVI Functions and Other Technologies

IRB interfaces and RVIs are similar to switch virtual interfaces (SVIs) and bridge-group virtual interfaces (BVI), which are supported on other vendors' devices. They can also be combined with other functions:

- VRF is often used in conjunction with Layer 3 subinterfaces, allowing traffic on a single physical interface to be differentiated and associated with multiple virtual routers. For more information about VRF, see [“Understanding Virtual Routing Instances on EX Series Switches” on page 2263](#).
- For redundancy, you can combine an IRB interface or RVI with implementations of the Virtual Router Redundancy Protocol (VRRP) in both bridging and virtual private LAN service (VPLS) environments. For more information about VRRP, see [“Understanding VRRP on EX Series Switches” on page 2501](#).

- Related Documentation**
- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)

## Understanding Private VLANs



**NOTE:** This topic describes support for private VLANs on Junos OS switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see *Understanding Private VLANs on EX Series Switches*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

VLANs limit broadcasts to specified users. Private VLANs (PVLANS) take this concept a step further by limiting communication within a VLAN. PVLANS accomplish this by restricting traffic flows through their member switch ports (which are called *private ports*) so that these ports communicate only with a specified uplink trunk port or with specified ports within the same VLAN. The uplink trunk port is usually connected to a router, firewall, server, or provider network.

PVLANS provide Layer 2 isolation between ports within a VLAN, splitting a broadcast domain into multiple discrete broadcast subdomains by creating secondary VLANs (community VLANs and an isolated VLAN) inside a primary VLAN. Ports within the same community VLAN can communicate with each other. Ports within an isolated VLAN can communicate *only* with a single uplink port.

PVLANS are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between hosts. The need to segregate a single VLAN is particularly useful in the following deployment scenarios:

- **Server farms**—A typical Internet service provider uses a server farm to provide Web hosting for numerous customers. Locating the various servers within a single server farm provides ease of management. Security concerns arise if all servers are in the same VLAN because Layer 2 broadcasts go to all servers in the VLAN.
- **Metropolitan Ethernet networks**—A metro service provider offers Layer 2 Ethernet access to assorted homes, rental communities, and businesses. The traditional solution of deploying one VLAN per customer is not scalable and is difficult to manage, leading to potential waste of IP addresses. PVLANS provide a more secure and more efficient solution.

This topic explains the following concepts regarding PVLANS:

- [Typical Structure and Primary Application of PVLANS on page 2259](#)
- [Routing Between Isolated and Community VLANs on page 2261](#)
- [PVLANS Use 802.1Q Tags to Identify Packets on page 2261](#)
- [PVLANS Use IP Addresses Efficiently on page 2262](#)
- [PVLAN Port Types and Forwarding Rules on page 2262](#)
- [Guidelines and Restrictions for PVLANS on page 2262](#)

### Typical Structure and Primary Application of PVLANS

A PVLAN can be configured on a single switch or can be configured to span multiple switches. The types of domains and ports are:

- **Primary VLAN**—The primary VLAN of the PVLAN is defined with an 802.1Q tag (VLAN ID) for the complete PVLAN. The primary PVLAN can contain multiple secondary VLANs (one isolated VLAN and multiple community VLANs).
- **Isolated VLAN**—A primary VLAN can contain only one isolated VLAN. An interface within an isolated VLAN can forward packets only to a promiscuous port or the Inter-Switch Link (ISL) port. An isolated interface cannot forward packets to another isolated interface; and an isolated interface cannot receive packets from another isolated interface. If a customer device needs to have access *only* to a gateway router, the device must be attached to an isolated trunk port.
- **Community VLAN**—You can configure multiple community VLANs within a single PVLAN. An interface within a specific community VLAN can establish Layer 2 communications with any other interface that belongs to the same community VLAN. An interface within a community VLAN can also communicate with a promiscuous port or the ISL port.
- **Promiscuous port**—A promiscuous port has Layer 2 communications with all the interfaces that are in the PVLAN, regardless of whether the interface belongs to an isolated VLAN or a community VLAN. A promiscuous port is a member of the primary VLAN, but is not included within any of the secondary subdomains. Layer 3 gateways, DHCP servers, and other trusted devices that need to communicate with endpoint devices are typically connected to a promiscuous port.
- **Inter-Switch Link (ISL)**—An ISL is a trunk port that connects multiple switches in a PVLAN and is required only when a PVLAN spans multiple switches..



**NOTE:** The primary VLAN of the PVLAN is defined with an 802.1Q tag (VLAN ID).

For information on configuring an isolated VLAN or community VLANs, see [“Creating a Private VLAN on a Single Switch \(CLI Procedure\)” on page 2343](#) or [“Creating a Private VLAN Spanning Multiple Switches \(CLI Procedure\)” on page 2345](#).

[Figure 29 on page 2260](#) shows a PVLAN on a single switch in which the primary VLAN (100) contains two community VLANs and one isolated VLAN.

Figure 29: Private VLAN on a Single Switch





Figure 30 on page 2261 shows a PVLAN spanning multiple switches in which the primary VLAN (100) contains two community VLANs (300) and (400) and one isolated VLAN (200).

Figure 30: PVLAN Spanning Multiple Switches



### Routing Between Isolated and Community VLANs

To route Layer 3 traffic between isolated and community VLANs, you must connect an external router or switch to a trunk port of the primary VLAN. The trunk port of the primary VLAN is a *promiscuous* port; therefore, it can communicate with *all* the ports in the PVLAN.

### PVLANS Use 802.1Q Tags to Identify Packets

You must specify an 802.1Q tag on the primary VLAN and on each of the secondary VLANs, regardless of whether the PVLAN is configured on a single switch or across

multiple switches. (You can also configure a VLAN name for the primary VLAN and for each of the secondary VLANs. This is optional.)

### PVLANS Use IP Addresses Efficiently

PVLANS provide IP address conservation and efficient allocation of IP addresses. In a typical network, VLANs usually correspond to a single IP subnet. In PVLANS, the hosts in all secondary VLANs belong to the same IP subnet because the subnet is allocated to the primary VLAN. Hosts within the secondary VLAN are assigned IP addresses based on IP subnets associated with the primary VLAN, and their IP subnet masking information reflects that of the primary VLAN subnet.

### PVLAN Port Types and Forwarding Rules

Table 230 on page 2262 summarizes Layer 2 connectivity between the different types of ports within a PVLAN.

**Table 230: PVLAN Ports and Layer 2 Forwarding**

From Port Type	To Isolated Ports?	To Promiscuous Ports?	To Community Ports?	To Inter-switch Link Port?
Isolated	Deny	Permit	Deny	Permit
Promiscuous	Permit	Permit	Permit	Permit
Community 1	Deny	Permit	Permit	Permit

### Guidelines and Restrictions for PVLANS

When configuring PVLANS, adhere to the following guidelines:

- An access interface can belong to only one PVLAN domain, that is, it cannot participate in two different primary VLANs..
- A trunk interface can be a member of two secondary VLANs as long as the secondary VLANs are in two *different* primary VLANs. A trunk interface cannot be a member of two secondary VLANs that are in the *same* primary VLAN.
- A single region of Multiple Spanning Tree Protocol (MSTP) must be configured on all VLANs that are included within the PVLAN.
- VLAN Spanning Tree Protocol (VSTP) is not supported.

Some configuration statements cannot be specified on a secondary VLAN. You can configure the following statements at the **[edit vlans vlan-name switch-options]** hierarchy level *only* on the primary PVLAN:

- **mac-table-size**
- **no-mac-learning**
- **mac-statistics**
- **interface-mac-limit**

The following features are *not* supported on PVLANS on Junos switches with support for the ELS configuration style:

- DHCP security features (DHCP snooping, dynamic ARP inspection, IP source guard)
- Egress VLAN firewall filters
- Ethernet ring protection (ERP)
- Flexible VLAN tagging
- *global-mac-statistics*
- Integrated routing and bridging (IRB) interfaces/routed VLAN interfaces (RVIs)
- Multicast snooping or IGMP snooping
- Multichassis link aggregation groups (MC-LAGs)
- Port mirroring
- Q-in-Q tunneling
- VLAN Spanning Tree Protocol (VSTP)
- Voice over IP (VoIP)

**Related  
Documentation**

- [Example: Configuring a Private VLAN on a Single Switch on page 2302](#)
- [Creating a Private VLAN on a Single Switch \(CLI Procedure\) on page 2343](#)
- [Creating a Private VLAN Spanning Multiple Switches \(CLI Procedure\) on page 2345](#)

## Understanding Virtual Routing Instances on EX Series Switches

Virtual routing instances allow administrators to divide a Juniper Networks EX Series Ethernet Switch into multiple independent virtual routers, each with its own routing table. Splitting a device into many virtual routing instances isolates traffic traveling across the network without requiring multiple devices to segment the network.

You can use virtual routing instances to isolate customer traffic on your network and to bind customer-specific instances to customer-owned interfaces.

Virtual routing and forwarding (VRF) is often used in conjunction with Layer 3 subinterfaces, allowing traffic on a single physical interface to be differentiated and associated with multiple virtual routers. Each logical Layer 3 subinterface can belong to only one routing instance.

EX Series switches support IPv4 and IPv6 unicast and multicast VRF traffic. See *EX Series Switch Software Features Overview* for details on VRF support by switch per Junos OS release.

**Related  
Documentation**

- [Understanding Layer 3 Subinterfaces on page 2591](#)
- [Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 2306](#)
- [Configuring Virtual Routing Instances \(CLI Procedure\) on page 2347](#)

## Understanding Multiple VLAN Registration Protocol (MVRP) on EX Series Switches

Multiple VLAN Registration Protocol (MVRP) is a Layer 2 messaging protocol that manages the addition, deletion, and renaming of active virtual LANs, thereby reducing network administrators' time spent on these tasks. Use MVRP on Juniper Networks EX Series Ethernet Switches to dynamically register and unregister active VLANs on trunk interfaces. Using MVRP means that you do not have to manually register VLANs on all connections—that is, you do not need to explicitly bind a VLAN to each trunk interface. With MVRP, you configure a VLAN on one switch interface and the VLAN configuration is distributed through all active switches in the domain.

MVRP is an application protocol of the Multiple Registration Protocol (MRP) and is defined in the IEEE 802.1ak standard. MRP and MVRP replace Generic Attribute Registration Protocol (GARP) and GARP VLAN Registration Protocol (GVRP) and overcome GARP and GVRP limitations.

This topic describes:

- [How MVRP Updates, Creates, and Deletes VLANs on the Switches on page 2264](#)
- [MVRP Is Disabled by Default on the Switches on page 2264](#)
- [MRP Timers Control MVRP Updates on page 2265](#)
- [MVRP Uses MRP Messages to Transmit Switch and VLAN States on page 2265](#)
- [Compatibility Issues with Junos OS Releases of MVRP on page 2266](#)

### How MVRP Updates, Creates, and Deletes VLANs on the Switches

---

When any MVRP-member VLAN is changed, that VLAN sends a protocol data unit (PDU) to all other MVRP-member active VLANs. The PDU informs the other VLANs which switches and interfaces currently belong to the sending VLAN. This way, all MVRP-member VLANs are always updated with the current VLAN state of all other MVRP-member VLANs. Timers dictate when PDUs can be sent and when switches receiving MVRP PDUs can update their MVRP VLAN information.

In addition to sending PDU updates, MVRP dynamically creates VLANs on member interfaces when a new VLAN is added to any one interface. This way, VLANs created on one member switch are propagated to other member switches as part of the MVRP message exchange process.

To keep VLAN membership information current, MVRP removes switches and interfaces when they become unavailable. Pruning VLAN information has these benefits:

- Limits the network VLAN configuration to active participants, thereby reducing network overhead.
- Limits broadcast, unknown unicast, and multicast (BUM) traffic to interested devices.

### MVRP Is Disabled by Default on the Switches

---

MVRP is disabled by default on the switches and, when enabled, affects only trunk interfaces. Once you enable MVRP, all VLAN interfaces on the switch belong to MVRP



(the default **normal** registration mode) and those interfaces accept PDU messages and send their own PDU messages. To prevent one or more interfaces from participating in MVRP, you can specifically configure an interface to **forbidden** registration mode instead of the default **normal** mode.

VLAN updating, dynamic VLAN configuration through MVRP, and VLAN pruning are all active on trunk interfaces when MVRP is enabled.

### MRP Timers Control MVRP Updates

MVRP registration and updates are controlled by timers that are part of the MRP. The timers define when MVRP PDUs can be sent and when MVRP information can be updated on a switch.

The timers are set on a per-interface basis, and on EX Series switches that use Juniper Networks Junos operating system (Junos OS) with support for the Enhanced Layer 2 Software (ELS) configuration style, the timers are also set on a per-switch basis.

On an EX Series switch that uses Junos OS with support for ELS, if the timer value set on an interface level is different from the value set on a switch level, the value on the interface level takes precedence.

The following MRP timers are used to control the operation of MVRP:

- Join timer—Controls the interval for the next MVRP PDU transmit opportunity.
- Leave timer—Controls the period of time that an interface on the switch waits in the leave state before changing to the unregistered state.
- LeaveAll timer—Controls the frequency with which the interface generates LeaveAll messages.



**BEST PRACTICE:** Unless there is a compelling reason to change the timer settings, leave the default settings in place. Modifying timers to inappropriate values can cause an imbalance in the operation of MVRP.

### MVRP Uses MRP Messages to Transmit Switch and VLAN States

MVRP uses MRP messages to register and declare MVRP states for a switch or VLAN and to inform the switching network that a switch or VLAN is leaving MVRP. These messages are communicated as part of the PDU sent by any switch interface to the other switches in the network.

The following MRP messages are communicated for MVRP:

- Empty—MVRP information is not declared and no VLAN is registered.
- In—MVRP information is not declared but a VLAN is registered.
- JoinEmpty—MVRP information is declared but no VLAN is registered.
- JoinIn—MVRP information is declared and a VLAN is registered.

- Leave—MVRP information that was previously declared is withdrawn.
- LeaveAll—Unregister all VLANs on the switch. VLANs must re-register to participate in MVRP.
- New—The MVRP information is new and a VLAN might not be registered yet.

### Compatibility Issues with Junos OS Releases of MVRP

Except in Junos OS Releases 11.2 and earlier, MVRP has conformed with IEEE standard 802.1ak and IEEE Draft 802.1Q regarding the inclusion of an extra byte in the protocol data units (PDUs) sent and received by MVRP. [Table 231 on page 2266](#) outlines the MVRP versions and whether or not each version includes the extra byte in the PDU.

[Table 231 on page 2266](#) also labels each MVRP version with a scenario number, which is used throughout the remainder of this discussion for brevity.

**Table 231: Junos OS MVRP Versions and Inclusion of Extra Byte in PDU**

MVRP in Junos OS Releases 11.2 and Earlier For EX Series Switches That Do Not Support Enhanced Layer 2 Software (ELS) Configuration Style	MVRP in Junos OS Releases 11.3 and Later For EX Series Switches That Do Not Support ELS	MVRP in Junos OS Releases 13.2 and Later For EX Series Switches With Support For ELS
Scenario 1	Scenario 2	Scenario 3
Includes extra byte in the PDU	By default, does not include extra byte in the PDU	By default, includes extra byte in the PDU

As a result of the non-conformance of Releases 11.2 and earlier and changes in the standards with regard to the extra byte, a compatibility issue exists between some of the Junos OS versions of MVRP. This issue can result in some versions of MVRP not recognizing PDUs without the extra byte.

To address this compatibility issue, the MVRP versions described in scenarios 2 and 3 include the ability to control whether or not the PDU includes the extra byte. Before using these controls, however, you must understand each scenario that applies to your environment and plan carefully so that you do not inadvertently create an additional compatibility issue between the MVRP versions in scenarios 2 and 3.

[Table 232 on page 2266](#) provides a summary of environments that include the various MVRP scenarios and whether or not a particular environment requires you to take action.

**Table 232: MVRP Environments and Description of Required Actions**

Environment	Action Required?	Action Description
Includes MVRP versions in scenario 1 only	No	—
Includes MVRP versions in scenario 2 only	No	—
Includes MVRP versions in scenario 3 only	No	—

Table 232: MVRP Environments and Description of Required Actions (*continued*)

Environment	Action Required?	Action Description
Includes MVRP versions in scenarios 1 and 2. MVRP version in scenario 2 is in its default state.	Yes	On switches running MVRP version in scenario 2, use the <b>add-attribute-length-in-pdu</b> statement. For more information, see <i>Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)</i> .
Includes MVRP versions in scenarios 1 and 3. MVRP version in scenario 3 is in its default state.	No	—
Includes MVRP versions in scenarios 2 and 3, and both versions are in their respective default states	Yes	Do one of the following: <ul style="list-style-type: none"> <li>On switches running MVRP version in scenario 2, use the <b>add-attribute-length-in-pdu</b> statement. For more information, see <i>Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)</i>.</li> <li>On switches running MVRP version in scenario 3, use the <b>no-attribute-length-in-pdu</b> statement. For more information, see “<i>Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)</i>” on page 2348.</li> </ul>

You can determine whether the switches in your network are running incompatible versions of MVRP by issuing the **show mvrp statistics** command. For more information on diagnosing and correcting this MVRP compatibility situation, see *Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)* or “*Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)*” on page 2348.

#### Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)
- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches](#)
- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 2310](#)

## Understanding MAC Address Aging

Juniper Networks EX Series Ethernet Switches store MAC addresses in the Ethernet switching table, also called the *MAC table*. When the aging time for a MAC address in the table expires, the address is removed.

If your switch runs Juniper Networks Junos operating system (Junos OS) for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, you can configure the MAC table aging time on all VLANs on the switch. If your switch runs Junos OS that does not support ELS, you can configure the MAC table aging time on all VLANs on the switch or on specified VLANs, as well as configure aging time to be unlimited, either on all VLANs or on specified VLANs, so that MAC addresses never age out of the table.

To learn MAC addresses, the switch reads all packets that it detects on the LAN or on the local VLAN, looking for MAC addresses of sending nodes. It places these addresses

into its Ethernet switching table, along with two other pieces of information—the interface on which the traffic was received and the time when the address was learned.

When the switch receives traffic on an interface, it searches the Ethernet switching table for the MAC address of the destination. If the MAC address is not found, the traffic is flooded out all of the other interfaces associated with the VLAN. For example, if traffic is received on an interface that is associated with VLAN v-10 and there is no entry in the Ethernet switching table for VLAN v-10 (the Ethernet switching table is organized by VLAN), then the traffic is flooded to all access and trunk interfaces that are members of VLAN v-10.

Flooding allows the switch to learn about destinations that are not yet in its Ethernet switching table. If a particular destination MAC address is not in the Ethernet switching table, the switch floods the traffic to all interfaces except the interface on which it was received. When the destination node receives the flooded traffic, it sends an acknowledgment packet back to the switch, allowing the switch to learn the MAC address of the node and to add the address to its Ethernet switching table.

The switch uses a mechanism called aging to keep the Ethernet switching table current. For each MAC address in the Ethernet switching table, the switch records a timestamp of when the information about the network node was learned. Each time the switch detects traffic from a MAC address that is in its Ethernet switching table, it updates the timestamp of that MAC address. A timer on the switch periodically checks the timestamp, and if the MAC address of a node is older than the value set, the switch removes that MAC address from the Ethernet switching table. This aging process ensures that the switch tracks only active MAC addresses on the network and that it is able to flush out from the Ethernet switching table MAC addresses that are no longer available.

You configure how long MAC addresses remain in the Ethernet switching table by:

- (On switches that run Junos OS with support for the ELS configuration style) Using the **global-mac-table-aging-time** statement in the **[edit protocols l2-learning]** hierarchy.
- (On switches that run Junos OS that does not support ELS) Using the **mac-table-aging-time** statement in either the **[edit ethernet-switching-options]** or the **[edit vlans]** hierarchy, depending on whether you want to configure it for the entire switch or only for specific VLANs.

For example, in a topology with EX switches that run Junos OS that does not support ELS, if you have a printer VLAN, you might choose to configure the aging time for that VLAN to be considerably longer than for other VLANs so that MAC addresses of printers on this VLAN age out less frequently. Because the MAC addresses remain in the table, even if a printer has been idle for some time before traffic arrives for it, the switch still finds the MAC address and does not need to flood the traffic to all other interfaces.

Similarly, in a data center environment where the list of servers connected to the switch is fairly stable, you might choose to increase MAC address aging time, or even set it to unlimited, to increase the efficiency of the utilization of network bandwidth by reducing flooding.

## Related Documentation

- [Configuring MAC Table Aging \(CLI Procedure\)](#)
- [Configuring MAC Table Aging \(CLI Procedure\) on page 2341](#)
- [Controlling Authentication Session Timeouts \(CLI Procedure\) on page 1930](#)

## Q-in-Q Tunneling

- [Understanding Q-in-Q Tunneling on EX Series Switches on page 2269](#)

### Understanding Q-in-Q Tunneling on EX Series Switches



**NOTE:** This topic applies to Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Understanding Q-in-Q Tunneling on EX Series Switches*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

Q-in-Q tunneling enables service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q VLAN tags are prepended by the service-provider VLAN (S-VLAN) tag. The Juniper Networks Junos operating system (Junos OS) implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.

This topic describes:

- [How Q-in-Q Tunneling Works on page 2269](#)
- [Sending and Receiving Untagged Packets on page 2270](#)
- [Disabling MAC Address Learning on page 2270](#)
- [Mapping C-VLANs to S-VLANs on page 2271](#)
- [Limitations for Q-in-Q Tunneling on page 2272](#)

#### How Q-in-Q Tunneling Works

In Q-in-Q tunneling, as a packet travels from a customer VLAN (C-VLAN) to an S-VLAN, a service-provider-specific 802.1Q tag is added to the packet. This additional tag is used to segregate traffic into S-VLANs. The original customer 802.1Q tag of the packet is retained and is transmitted transparently, passing through the service provider's network. As the packet leaves the S-VLAN in the downstream direction, the additional 802.1Q tag is removed.

When Q-in-Q tunneling is configured on Juniper Networks EX Series Ethernet Switches, trunk interfaces are assumed to be part of the service-provider network and access interfaces are assumed to be part of the customer network. Therefore, this topic also refers to trunk interfaces as S-VLAN interfaces (network-to-network interfaces [NNI]),

and to access interfaces as C-VLAN interfaces (user-network interfaces [UNI]). An access interface can receive both tagged and untagged frames in this case.

An interface can be a member of multiple S-VLANs. You can map one C-VLAN to one S-VLAN (1:1) or many C-VLANs to many S-VLANs (N:N). Customer packets that traverse an S-VLAN are double-tagged for an additional layer of segregating or bundling of C-VLANs. C-VLAN and S-VLAN tags are unique—for instance, you can have both a C-VLAN tag of 101 and an S-VLAN tag of 101. You can limit the set of accepted customer tags to a range of tags or to discrete values. Class-of-service (CoS) values of C-VLANs are unchanged in the downstream direction. You may, optionally, copy ingress priority and CoS settings to the S-VLAN.

C-VLAN and S-VLAN interfaces accept priority-tagged packets without any configuration.



**NOTE:** On an EX4300 switch, you can configure multiple logical interfaces on the same Ethernet port, but each logical interface supports only single-tagged packets and that tag must include a different VLAN ID than those supported by the other logical interfaces. Given this situation, you cannot enable Q-in-Q tunneling on Ethernet ports with multiple logical subinterfaces.

---

### **Sending and Receiving Untagged Packets**

---

To enable a C-VLAN or S-VLAN interface to send and receive untagged packets, you must configure a native VLAN for the interface, then specify a VLAN ID for the native VLAN. After performing this configuration, when a C-VLAN or S-VLAN interface receives an untagged packet, it adds the VLAN ID of the native VLAN to the packet and sends the newly tagged packet to the mapped interface.

To specify a native VLAN ID, use the **native-vlan-id** statement at the **[edit interfaces interface-name]** hierarchy level. When specifying a native VLAN ID on a C-VLAN or S-VLAN physical interface, the value must match the VLAN ID or be included in the VLAN ID list specified on the C-VLAN or S-VLAN logical interface.

For example, on a logical interface for a C-VLAN interface, you specify a C-VLAN ID list of 100-200. Then, on the C-VLAN physical interface, you specify a native VLAN ID of 150. This configuration will work because the native VLAN of 150 is included in the C-VLAN ID list of 100-200.

We recommend configuring a native VLAN when using any of the approaches to map C-VLANs to S-VLANs. If you do not configure a native VLAN on an interface, untagged packets received by the interface are discarded. See the Mapping C-VLANs to S-VLANs section in this topic for information about the methods of mapping C-VLANs to S-VLANs.

---

### **Disabling MAC Address Learning**

---

In a Q-in-Q deployment, customer packets from downstream interfaces are transported without any changes to source and destination MAC addresses. You can disable MAC address learning at the global, interface, and VLAN levels:

At the global level, you disable MAC address learning for the switch.

At the interface level, you disable MAC address learning for all VLANs of which the specified interface is a member.

At the VLAN level, you disable MAC address learning for a specified VLAN. MAC addresses that have already been learned for the VLAN are flushed.

### Mapping C-VLANs to S-VLANs

There are three ways to map C-VLANs to S-VLANs:

- [All-in-One Bundling on page 2271](#)
- [Many-to-Many Bundling on page 2271](#)
- [Mapping a Specific Interface on page 2271](#)

If you configure multiple mapping methods, the switch gives priority to mapping a specific interface, then to many-to-many bundling, and last to all-in-one bundling. However, for a particular mapping method, setting up overlapping rules for the same C-VLAN is not supported.

#### ***All-in-One Bundling***

All-in-one bundling maps all packets from all C-VLAN interfaces to an S-VLAN.

The C-VLAN interface accepts untagged and single-tagged packets. An S-VLAN 802.1Q tag is then added to these packets, and the packets are sent to the S-VLAN interface, which accepts untagged, single-tagged, and double-tagged packets.



**NOTE:** The C-VLAN and S-VLAN interfaces accept untagged packets provided that the `native-vlan-id` statement is configured on these interfaces.

#### ***Many-to-Many Bundling***

Many-to-many bundling is used to specify which C-VLANs are mapped to which S-VLANs.

Many-to-many bundling is used when you want a subset of the C-VLANs on the access switch to be part of multiple S-VLANs. With many-to-many bundling, the C-VLAN interfaces accept untagged and single-tagged packets. An S-VLAN 802.1Q tag is then added to these packets, and the packets are sent to the S-VLAN interfaces, which accept untagged, single-tagged, and double-tagged packets.



**NOTE:** The C-VLAN and S-VLAN interfaces accept untagged packets provided that the `native-vlan-id` statement is configured on these interfaces.

#### ***Mapping a Specific Interface***

Use specific interface mapping when you want to assign an S-VLAN to a specific C-VLAN on an interface. The configuration applies only to the specific interface, not to all access

interfaces as in the cases of the all-in-one bundling and many-to-many bundling approaches.

Specific interface mapping has two suboptions for treatment of traffic: push and swap. When traffic that is mapped to a specific interface is pushed, the packet retains its tag as it moves from the C-VLAN to the S-VLAN, then an additional S-VLAN tag is added to the packet. When traffic that is mapped to a specific interface is swapped, the incoming tag is replaced with a new VLAN tag, which is also referred to as VLAN rewrite.

It might be useful to have S-VLANs that provide service to multiple customers. Each customer typically has its own S-VLAN plus access to one or more S-VLANs that are used by multiple customers. A specific tag on the customer side is mapped to an S-VLAN. Typically, this functionality is used to keep data from different customers separate or to provide individualized treatment of the packets on a certain interface.

When using specific interface mapping, the C-VLAN interfaces accept untagged and single-tagged packets, while the S-VLAN interfaces accept untagged, single-tagged, and double-tagged packets.



**NOTE:** The C-VLAN and S-VLAN interfaces accept untagged packets provided that the `native-vlan-id` statement is configured on these interfaces.

---

### Limitations for Q-in-Q Tunneling

---

Q-in-Q tunneling does not support most access port security features. There is no per-VLAN (customer) policing or per-VLAN (outgoing) shaping and limiting with Q-in-Q tunneling unless you configure these security features by using firewall filters.

#### Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)
- [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 2351](#)

## Layer 2 Protocol Tunneling

---

- [Understanding Layer 2 Protocol Tunneling on EX Series Switches on page 2272](#)

### Understanding Layer 2 Protocol Tunneling on EX Series Switches

---



**NOTE:** This topic describes Layer 2 protocol tunneling (L2PT) on Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Understanding Layer 2 Protocol Tunneling on EX Series Switches*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

---

L2PT enables service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain. This feature is useful when you want to run



Layer 2 protocols on a network that includes switches located at remote sites that are connected across a service provider network.

This topic includes:

- [Layer 2 Protocols Supported by L2PT on EX Series Switches on page 2273](#)
- [How L2PT Works on page 2274](#)
- [L2PT and Q-in-Q Tunneling on EX Series Switches on page 2275](#)

### Layer 2 Protocols Supported by L2PT on EX Series Switches

L2PT on EX Series switches supports the following Layer 2 protocols:

- 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM)



**NOTE:** If you enable L2PT for untagged OAM LFM packets, do not configure LFM on the corresponding access interface.

- Cisco Discovery Protocol (CDP)
- Ethernet Local Management Interface (E-LMI)
- MVRP VLAN Registration Protocol (MVRP)
- Link Aggregation Control Protocol (LACP)



**NOTE:** If you enable L2PT for untagged LACP packets, do not configure LACP on the corresponding access interface.

- Link Layer Discovery Protocol (LLDP)
- Multiple MAC Registration Protocol (MMRP)
- Multiple VLAN Registration Protocol (MVRP)
- Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP)
- Unidirectional Link Detection (UDLD)
- VLAN Spanning Tree Protocol (VSTP)
- VLAN Trunking Protocol (VTP)

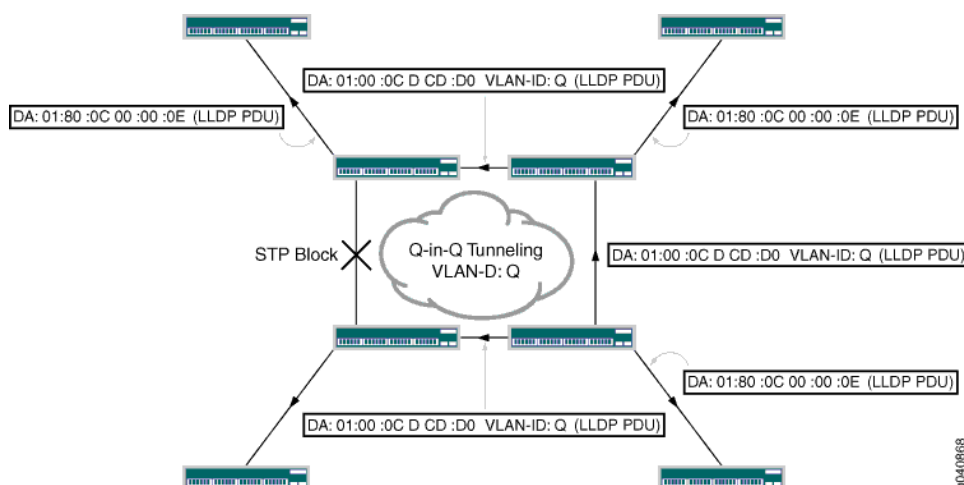


**NOTE:** CDP, UDLD, and VTP cannot be configured on EX Series switches. L2PT does, however, tunnel CDP, UDLD, and VTP PDUs.

## How L2PT Works

L2PT works by encapsulating Layer 2 PDUs, tunneling them across a service provider network, and decapsulating them for delivery to their destination switches. L2PT encapsulates Layer 2 PDUs by enabling the ingress provider edge (PE) device to rewrite the PDUs' destination media access control (MAC) addresses before forwarding them to the service provider network. The devices in the service provider network treat these encapsulated PDUs as multicast Ethernet packets. On receiving these PDUs, the egress PE devices decapsulate them by replacing the destination MAC addresses with the address of the Layer 2 protocol that is being tunneled before forwarding the PDUs to their destination switches. This process is illustrated in [Figure 31 on page 2274](#).

Figure 31: L2PT Example



When a PE device receives a Layer 2 control PDU from any of the customer PE devices, it changes the destination MAC address to 01:00:0C:CD:CD:D0. The modified packet is then sent to the provider network. All devices on the provider network treat these packets as multicast Ethernet packets and deliver them to all PE devices for the customer. The egress PE devices receive all the control PDUs with the same MAC address (01:00:0C:CD:CD:D0). Then they identify the packet type by doing deeper packet inspection and replace the destination MAC address 01:00:0C:CD:CD:D0 with the appropriate destination address. The modified PDUs are sent out to the customer PE devices, thus ensuring the Layer 2 control PDUs are delivered, in their original state, across the provider network. The L2PT protocol is valid for all types of packets (untagged, tagged, and Q-in-Q tagged).

The switch tunnels all Layer 2 PDUs regardless of the speed at which they are received, although the number of packets tunneled per second might be limited by other factors.

L2PT supports tunneling of STP, LLDP, CDP, and VTP control PDUs across the service provider network. The PE device identifies the Layer 2 control protocols by their encapsulated MAC address. The destination MAC address used by different protocols is listed in [Table 233 on page 2275](#):

Table 233: Protocol Destination MAC Addresses

Protocol	Ethernet Encapsulation	MAC Address
802.3ah	Ether-II	01:80:C2:00:00:02
CDP	SNAP	01:00:0C:CC:CC:CC
E-LMI	Ether-II	01:80:C2:00:00:07
MVRP	Ether-II	01:80C2:00:00:21
LACP	Ether-II	01:80:C2:00:00:02
STP, RSTP, and MSTP	SNAP	01:80:C2:00:00:21
LLDP	Ether-II	01:80:0C:00:00:0E
MMRP	Ether-II	01:80:C2:00:00:0E
UDLD	SNAP	01:00:0C:CC:CC:CC
VSTP	SNAP	01:00:0C:CC:CC:CD
VTP	SNAP	01:00:0C:CC:CC:CC

### L2PT and Q-in-Q Tunneling on EX Series Switches

You must enable Q-in-Q tunneling before you can configure L2PT. For information about Q-in-Q tunneling, see [“Understanding Q-in-Q Tunneling on EX Series Switches” on page 2269](#).

#### **Related Documentation**

- [Configuring Layer 2 Protocol Tunneling on EX Series Switches \(CLI Procedure\) on page 2357](#)
- [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 2351](#)

### Redundant Trunk Groups

- [Understanding Redundant Trunk Links on page 2276](#)

## Understanding Redundant Trunk Links

In a typical enterprise network composed of distribution and access layers, a redundant trunk link provides a simple solution for network recovery when a trunk port on a switch goes down. In that case, traffic is routed to another trunk port, keeping network convergence time to a minimum.

To configure a redundant trunk link, create a redundant trunk group. The redundant trunk group is configured on the access switch and contains two links: a primary or active link, and a secondary link. If the active link fails, the secondary link automatically starts forwarding data traffic without waiting for normal spanning-tree protocol convergence.

Data traffic is forwarded only on the active link. Data traffic on the secondary link is dropped and shown as dropped packets when you issue the operational mode command **show interfaces *interface-name* extensive**.

While data traffic is blocked on the secondary link, Layer 2 control traffic is still permitted. For example, an LLDP session can be run between two switches on the secondary link.

Rapid Spanning Tree Protocol (RSTP) is enabled by default on the switches to create a loop-free topology, but an interface is not allowed to be in both a redundant trunk group and in a spanning-tree protocol topology at the same time. You must disable RSTP on an interface if a redundant trunk group is configured on that interface. For example, in [Figure 32 on page 2277](#), in addition to disabling RSTP on the Switch 3 interfaces, you must also disable RSTP on the Switch 1 and Switch 2 interfaces connected to Switch 3. Spanning-tree protocols can, however, continue operating on other interfaces on those switches—for example on the link between Switch 1 and Switch 2.

[Figure 32 on page 2277](#) shows three switches in a basic topology for redundant trunk links. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer. Switch 3 is connected to the distribution layer through trunk ports ge-0/0/9.0 (Link 1) and ge-0/0/10.0 (Link 2). Link 1 and Link 2 are in a redundant trunk group called group1. Link 1 is designated as the primary link. Traffic flows between Switch 3 in the access layer and Switch 1 in the distribution layer through Link 1. While Link 1 is active, Link 2 blocks traffic.

Figure 32: Redundant Trunk Group, Link 1 Active





Figure 33 on page 2277 illustrates how the redundant trunk link topology works when the primary link goes down.

Figure 33: Redundant Trunk Group, Link 2 Active



When Link 1 between Switch 1 and Switch 3 goes down, Link 2 takes over as the active link. Traffic between the access layer and the distribution layer is then automatically switched to Link 2 between Switch 1 and Switch 2.

**Related  
Documentation**

- [Example: Configuring Redundant Trunk Links for Faster Recovery](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 2326](#)

---

## Proxy ARP

- [Understanding Proxy ARP on EX Series Switches on page 2278](#)

### Understanding Proxy ARP on EX Series Switches

You can configure proxy Address Resolution Protocol (ARP) on your Juniper Networks EX Series Ethernet Switch to enable the switch to respond to ARP queries for network addresses by offering its own Ethernet media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

Proxy ARP is useful in situations where hosts are on different physical networks and you do not want to use subnet masking. Because ARP broadcasts are not propagated between hosts on different physical networks, hosts will not receive a response to their ARP request if the destination is on a different subnet. Enabling the switch to act as an ARP proxy allows the hosts to transparently communicate with each other through the switch. Proxy ARP can help hosts on a subnet reach remote subnets without your having to configure routing or a default gateway.

- [What Is ARP? on page 2278](#)
- [Proxy ARP Overview on page 2278](#)
- [Best Practices for Proxy ARP on EX Series Switches on page 2279](#)

---

#### What Is ARP?

Ethernet LANs use ARP to map Ethernet MAC addresses to IP addresses. Each device maintains a cache containing a mapping of MAC addresses to IP addresses. The switch maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

---

#### Proxy ARP Overview

When proxy ARP is enabled, if the switch receives an ARP request for which it has a route to the target (destination) IP address, the switch responds by sending a proxy ARP reply packet containing its own MAC address. The host that sent the ARP request then sends its packets to the switch, which forwards them to the intended host.



**NOTE:** For security reasons, the source address in an ARP request must be on the same subnet as the interface on which the ARP request is received.

---

You can configure proxy ARP for each interface. You can also configure proxy ARP for an integrated routing and bridging (IRB) interface named `irb` or a routed VLAN interface (RVI) named `vlan`. (On EX Series switches that use Juniper Networks Junos operating system (Junos OS) with support for the Enhanced Layer 2 Software (ELS) configuration style, the feature is known as an IRB interface. On EX Series switches that use Junos OS that does not support ELS, the feature is known as an RVI.)

EX Series switches support two modes of proxy ARP, restricted and unrestricted. Both modes require that the switch have an active route to the destination address of the ARP request.

- **Restricted**—The switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are on the same subnet. In this mode, hosts on the same subnet communicate without proxy ARP. We recommend that you use this mode on the switch.
- **Unrestricted**—The switch responds to all ARP requests for which it has a route to the destination. This is the default mode (because it is the default mode in Juniper Networks Junos operating system (Junos OS) configurations other than those on the switch). We recommend using restricted mode on the switch.

### Best Practices for Proxy ARP on EX Series Switches

We recommend these best practices for configuring proxy ARP on the switches:

- Set proxy ARP on the interfaces that you want, including IRB interfaces or RVIs, to restricted mode.
- If you set proxy ARP to unrestricted, disable gratuitous ARP requests on each interface enabled for proxy ARP.

#### Related Documentation

- [Example: Configuring Proxy ARP on an EX Series Switch on page 2331](#)
- [Configuring Proxy ARP \(CLI Procedure\) on page 2360](#)





## CHAPTER 38

# Configuration

- Configuration Examples on page 2281
- Configuration Tasks on page 2333
- Configuration Statements on page 2364

### Configuration Examples

---

- Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 2281
- Example: Connecting Access Switches to a Distribution Switch on page 2291
- Example: Configuring a Private VLAN on a Single Switch on page 2302
- Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 2306
- Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 2310
- Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 2322
- Example: Configuring Redundant Trunk Links for Faster Recovery on page 2326
- Example: Configuring Proxy ARP on an EX Series Switch on page 2331

### Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch

---



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs Junos OS that does not support ELS, see *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

---

EX Series switches use bridging and virtual LANs (VLANs) to connect network devices in a LAN—desktop computers or laptops, IP telephones, printers, file servers, wireless access points, and others—and to segment the LAN into smaller broadcast domains.

This example describes how to configure basic bridging and a VLAN on an EX Series switch:

- [Requirements on page 2282](#)
- [Overview and Topology on page 2282](#)
- [Configuration on page 2283](#)
- [Verification on page 2287](#)

---

## Requirements

This example uses the following hardware and software components:

- One EX Series switch
- Junos OS Release 13.2X50-D10 or later for EX Series switches

Before you set up bridging and a VLAN, be sure you have:

- Installed your EX Series switch. See the installation instructions for your switch.
- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (CLI Procedure)*.

---

## Overview and Topology

EX Series switches connect network devices in an office LAN or a data center LAN to provide sharing of common resources such as printers and file servers and to enable wireless devices to connect to the LAN through wireless access points. Without bridging and VLANs, all devices on the Ethernet LAN are in a single broadcast domain, and all the devices detect all the packets on the LAN. Bridging creates separate broadcast domains on the LAN, creating VLANs, which are independent logical networks that group together related devices into separate network segments. The grouping of devices on a VLAN is independent of where the devices are physically located in the LAN.

To use an EX Series switch to connect network devices on a LAN, you must, at a minimum, explicitly configure at least one VLAN, even if your network is simple and you want only one broadcast domain to exist, as is the case with this example. You must also assign all needed interfaces to the VLAN, after which the interfaces function in access mode. After the VLAN is configured, you can plug access devices—such as desktop or laptop computers, IP telephones, file servers, printers, and wireless access points—into the switch, and they are joined immediately into the VLAN, and the LAN is up and running.

The topology used in this example consists of one EX4300-24P switch, which has a total of 24 ports. All ports support Power over Ethernet (PoE), which means they provide both network connectivity and electric power for the device connecting to the port. To these ports, you can plug in devices requiring PoE, such as Avaya VoIP telephones, wireless access points, and some IP cameras. (Avaya phones have a built-in hub that allows you to connect a desktop PC to the phone, so the desktop and phone in a single office require only one port on the switch.) [Table 234 on page 2283](#) details the topology used in this configuration example.

Table 234: Components of the Basic Bridging Configuration Topology

Property	Settings
Switch hardware	EX4300-24P switch, with 24 Gigabit Ethernet ports: in this example, 8 ports are used as PoE ports (ge-0/0/0 through ge-0/0/7 ) and 16 ports used as non-PoE ports (ge-0/0/8 through ge-0/0/23)
VLAN name	employee-vlan
VLAN ID	10
Connection to wireless access point (requires PoE)	ge-0/0/0
Connections to Avaya IP telephone—with integrated hub, to connect phone and desktop PC to a single port (requires PoE)	ge-0/0/1 through ge-0/0/7
Direct connections to desktop PCs and laptops (no PoE required)	ge-0/0/8 through ge-0/0/12
Connections to file servers (no PoE required)	ge-0/0/17 and ge-0/0/18
Connections to integrated printer/fax/copier machines (no PoE required)	ge-0/0/19 through ge-0/0/20
Unused ports (for future expansion)	ge-0/0/13 through ge-0/0/16, and ge-0/0/21 through ge-0/0/23

### Configuration

To set up basic bridging and a VLAN:

#### CLI Quick Configuration

To quickly configure a VLAN, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans employee-vlan vlan-id 10
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/4 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/5 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/6 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/7 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/8 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/12 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/17 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/19 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces ge-0/0/20 unit 0 family ethernet-switching vlan members employee-vlan
```

You must then plug the wireless access point into PoE-enabled port **ge-0/0/0** and the Avaya IP phones into the PoE-enabled ports **ge-0/0/1** through **ge-0/0/7**. Also, plug the PCs, file servers, and printers into ports **ge-0/0/8** through **ge-0/0/12** and **ge-0/0/17** through **ge-0/0/20**.

#### Step-by-Step Procedure

To set up basic bridging and a VLAN:

1. Create a VLAN named **employee-vlan** and specify the VLAN ID of 10 for it:  

```
[edit vlans]
user@switch# set employee-vlan vlan-id 10
```
2. Assign interfaces **ge-0/0/0** through **ge-0/0/12**, and **ge-0/0/17** through **ge-0/0/20** to the **employee-vlan** VLAN:  

```
[edit interface]
user@switch# set ge-0/0/0 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/1 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/2 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/3 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/4 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/5 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/6 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/7 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/8 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/9 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/11 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/12 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/17 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/18 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/19 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set ge-0/0/20 unit 0 family ethernet-switching vlan members employee-vlan
```
3. Connect the wireless access point to switch port **ge-0/0/0**.
4. Connect the seven Avaya phones to switch ports **ge-0/0/1** through **ge-0/0/7**.
5. Connect the five PCs to ports **ge-0/0/8** through **ge-0/0/12**.
6. Connect the two file servers to ports **ge-0/0/17** and **ge-0/0/18**.
7. Connect the two printers to ports **ge-0/0/19** and **ge-0/0/20**.

**Results** Check the results of the configuration:

```
user@switch> show configuration
ge-0/0/0 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
ge-0/0/2 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
ge-0/0/3 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
ge-0/0/4 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
ge-0/0/5 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
ge-0/0/6 {
  unit 0 {
    family ethernet-switching {
```

```
        vlan {
            members employee-vlan;
        }
    }
}
ge-0/0/7 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
ge-0/0/8 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
ge-0/0/9 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
ge-0/0/10 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
ge-0/0/11 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
ge-0/0/12 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
```

```

ge-0/0/17 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
ge-0/0/18 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
ge-0/0/19 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
ge-0/0/20 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}

```

### Verification

To verify that switching is operational and that **employee-vlan** has been created, perform these tasks:

- [Verifying That the VLAN Has Been Created on page 2287](#)
- [Verifying That Interfaces Are Associated with the Proper VLANs on page 2288](#)

#### *Verifying That the VLAN Has Been Created*

**Purpose** Verify that the VLAN named **employee-vlan** has been created on the switch.

**Action** List all VLANs configured on the switch:

```
user@switch> show vlans
Routing instance      VLAN name      Tag      Interfaces
default-switch        employee-vlan   10
                      ge-0/0/0.0
                      ge-0/0/1.0
                      ge-0/0/2.0
                      ge-0/0/3.0
                      ge-0/0/4.0
                      ge-0/0/5.0
                      ge-0/0/6.0
                      ge-0/0/7.0
                      ge-0/0/8.0
                      ge-0/0/9.0
                      ge-0/0/10.0
                      ge-0/0/11.0
                      ge-0/0/12.0
                      ge-0/0/17.0
                      ge-0/0/18.0
                      ge-0/0/19.0
                      ge-0/0/20.0
...
```

**Meaning** The `show vlans` command lists the VLANs configured on the switch. This output shows that the VLAN `employee-vlan` has been created.

***Verifying That Interfaces Are Associated with the Proper VLANs***

**Purpose** Verify that Ethernet switching is enabled on switch interfaces and that all interfaces are included in the VLAN.



**Action** List all interfaces on which switching is enabled:

```

user@switch> show ethernet-switching interfaces
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
ge-0/0/0.0                65535                untagged
                        employee-vlan 10
                        65535      Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
ge-0/0/1.0                65535                untagged
                        employee-vlan 10
                        65535      Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
ge-0/0/2.0                65535                untagged
                        employee-vlan 10
                        65535      Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
ge-0/0/3.0                65535                untagged
                        employee-vlan 10
                        65535      Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
ge-0/0/4.0                65535                untagged
                        employee-vlan 10
                        65535      Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
ge-0/0/5.0                65535                untagged
                        employee-vlan 10
                        65535      Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
ge-0/0/6.0                65535                untagged
                        employee-vlan 10
                        65535      Discarding
Routing Instance Name : default-switch

```

```

Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state interface flags
ge-0/0/7.0
    employee-vlan 10
                        65535
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state interface flags
ge-0/0/8.0
    employee-vlan 10
                        65535
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state interface flags
ge-0/0/9.0
    employee-vlan 10
                        65535
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state interface flags
ge-0/0/10.0
    employee-vlan 10
                        65535
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state interface flags
ge-0/0/11.0
    employee-vlan 10
                        65535
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state interface flags
ge-0/0/12.0
    employee-vlan 10
                        65535
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state interface flags
ge-0/0/17.0
    employee-vlan 10
                        65535
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state interface flags
ge-0/0/18.0
    employee-vlan 10
                        65535
                        65535 Discarding

```

```

        employee-vlan 10
                               65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members limit state interface flags
ge-0/0/19.0      65535      untagged
        employee-vlan 10
                               65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members limit state interface flags
ge-0/0/20.0      65535      untagged
        employee-vlan 10
                               65535    Discarding
...

```

**Meaning** The **show ethernet-switching interfaces** command lists all interfaces on which switching is enabled (in the **Logical interface** column), along with the VLANs that are active on the interfaces (in the **VLAN members** column). The output in this example shows all the connected interfaces, ge-0/0/0 through ge-0/0/12 and ge-0/0/17 through ge-0/0/20 and that they are all part of VLAN **employee-vlan**. Notice that the interfaces listed are the logical interfaces, not the physical interfaces. For example, the output shows ge-0/0/0.0 instead of ge-0/0/0. This is because Junos OS creates VLANs on logical interfaces, not directly on physical interfaces.

**Related Documentation**

- [Configuring VLANs for EX Series Switches \(CLI Procedure\) on page 2337](#)
- [Example: Connecting Access Switches to a Distribution Switch on page 2291](#)
- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)

## Example: Connecting Access Switches to a Distribution Switch



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Connecting an Access Switch to a Distribution Switch*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

In large local area networks (LANs), you commonly need to aggregate traffic from a number of access switches into a distribution switch.

This example describes how to connect access switches to a distribution switch:

- [Requirements on page 2292](#)
- [Overview and Topology on page 2292](#)
- [Configuring the Access Switch on page 2294](#)

- [Configuring the Distribution Switch on page 2298](#)
- [Verification on page 2300](#)

## Requirements

---

This example uses the following hardware and software components:

- Three EX Series access switches.
- One EX Series distribution switch.



**NOTE:** In an access switch-distribution switch topology, you can connect EX Series switches that run a version of Junos OS that supports ELS with EX Series switches that do not run a version of Junos OS that supports ELS. However, this example uses switches running ELS only to show how to configure this topology using the ELS CLI.

---

- Junos OS Release 12.3R2 or later that supports ELS for EX Series switches.

Before you connect an access switch to a distribution switch, be sure you have:

- Installed the switches. See the installation instructions for your switch.
- Performed the initial software configuration on both switches. For information about the initial software configuration for all EX Series switches except the EX9200 Series switches, see *Connecting and Configuring an EX Series Switch (CLI Procedure)*. For information about the initial software configuration for the EX9200 Series switches, see *Connecting and Configuring an EX9200 Switch (CLI Procedure)*.

## Overview and Topology

---

In a large office that is spread across several floors or buildings, or in a data center, you commonly aggregate traffic from a number of access switches into a distribution switch. This configuration example shows a simple topology to illustrate how to connect three access switches to a distribution switch.

In the topology, the LAN is segmented into two VLANs, one for the sales department and the second for the support team. One 1-Gigabit Ethernet port on one of the access switch's uplink modules connects to the distribution switch, to one 1-Gigabit Ethernet port on the distribution switch.

[Figure 34 on page 2293](#) shows an EX9200 distribution switch that is connected to three EX4300 access switches.

Figure 34: Sample Access Switch-Distribution Switch Topology

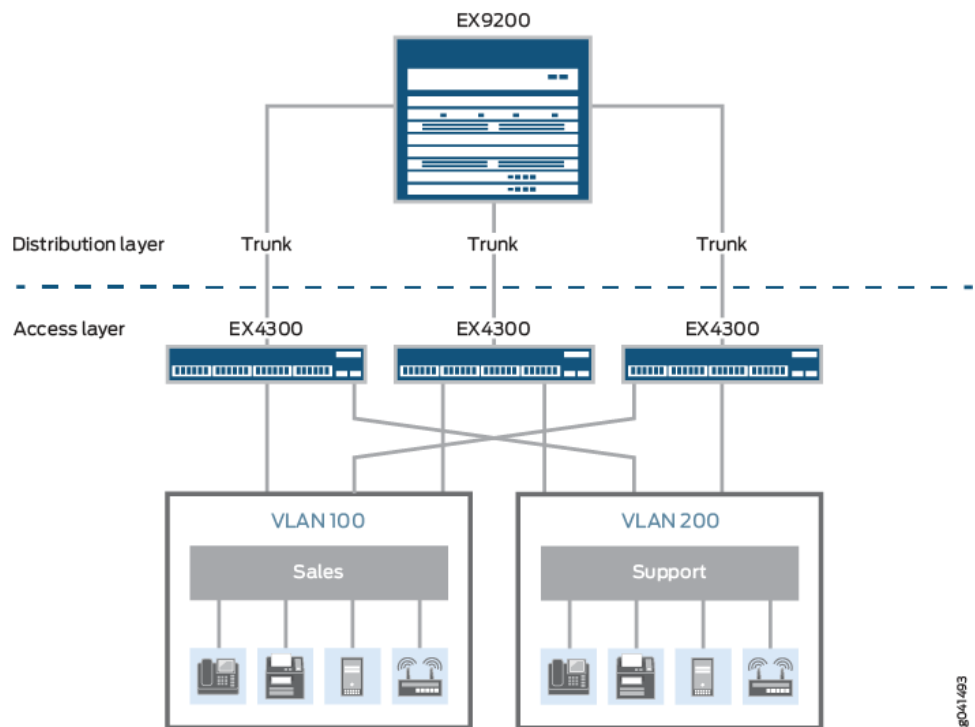


Table 235 on page 2293 describes the components of the example topology. The example shows how to configure one of the three access switches. The other access switches could be configured in the same manner.

**Table 235: Components of the Topology for Connecting an Access Switch to a Distribution Switch**

Property	Settings
Access switch hardware	Three EX4300 switches, each with an uplink module with 1-Gigabit Ethernet ports..
Distribution switch hardware	One EX9208 with up to three EX9200-40T line cards installed, which at full duplex, can provide up to 240 1-Gigabit ports.
VLAN names and tag IDs	<b>sales</b> , tag <b>100</b> <b>support</b> , tag <b>200</b>
VLAN subnets	<b>sales</b> : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) <b>support</b> : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)
Trunk port interfaces	On the access switch: ge-0/2/0 On the distribution switch: ge-0/0/0
Access port interfaces in VLAN <b>sales</b> (on access switch)	Avaya IP telephones: ge-0/0/3 through ge-0/0/19 Wireless access points: ge-0/0/0 and ge-0/0/1 Printers: ge-0/0/22 and ge-0/0/23 File servers: ge-0/0/20 and ge-0/0/21

Table 235: Components of the Topology for Connecting an Access Switch to a Distribution Switch (*continued*)

Property	Settings
Access port interfaces in VLAN <b>support</b> (on access switch)	Avaya IP telephones: ge-0/0/25 through ge-0/0/43 Wireless access points: ge-0/0/24 Printers: ge-0/0/44 and ge-0/0/45 File servers: ge-0/0/46 and ge-0/0/47

### Configuring the Access Switch

To configure the access switch:

#### CLI Quick Configuration

To quickly configure the access switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 unit 0 description "Sales wireless access point port"
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/3 unit 0 description "Sales phone port"
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/22 unit 0 description "Sales printer port"
set interfaces ge-0/0/22 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/20 unit 0 description "Sales file server port"
set interfaces ge-0/0/20 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/24 unit 0 description "Support wireless access point port"
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/26 unit 0 description "Support phone port"
set interfaces ge-0/0/26 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/44 unit 0 description "Support printer port"
set interfaces ge-0/0/44 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/46 unit 0 description "Support file server port"
set interfaces ge-0/0/46 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/2/0 unit 0 description "Uplink module port connection to distribution switch"
set interfaces ge-0/2/0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/2/0 native-vlan-id 1
set interfaces ge-0/2/0 unit 0 family ethernet-switching vlan members [sales support]
set interfaces ge-0/2/0 unit 0 family ethernet-switching vlan members 1
set interfaces irb unit 0 family inet address 192.0.2.1/25
set interfaces irb unit 1 family inet address 192.0.2.129/25
set vlans sales description "Sales VLAN"
set vlans sales l3-interface irb.0
set vlans sales vlan-id 100
set vlans support description "Support VLAN"
set vlans support vlan-id 200
set vlans support l3-interface irb.1
```

#### Step-by-Step Procedure

To configure the access switch:

1. Configure the 1-Gigabit Ethernet interface on the uplink module to be the trunk port that connects to the distribution switch:
 

```
[edit interfaces]
user@access-switch# set ge-0/2/0 unit 0 description "Uplink module port connection to distribution switch"
user@access-switch# set ge-0/2/0 unit 0 family ethernet-switching interface-mode trunk
```
2. Specify the VLANs to be aggregated on the trunk port:

- ```
[edit interfaces]
user@access-switch# set ge-0/2/0 unit 0 family ethernet-switching vlan members [ sales support ]
```
3. To handle untagged packets that are received on the trunk port, create a native VLAN by configuring a VLAN ID and specifying that the trunk port is a member of the native VLAN:
 

```
[edit interfaces]
user@access-switch# set ge-0/2/0 native-vlan-id 1
user@access-switch# set ge-0/2/0 unit 0 family ethernet-switching vlan members 1
```
  4. Configure the sales VLAN:
 

```
[edit vlans]
user@access-switch# set sales description "Sales VLAN"
user@access-switch# set sales vlan-id 100
user@access-switch# set sales l3-interface irb.0
```
  5. Configure the support VLAN:
 

```
[edit vlans]
user@access-switch# set support description "Support VLAN"
user@access-switch# set support vlan-id 200
user@access-switch# set support l3-interface irb.1
```
  6. Create the subnet for the sales VLAN:
 

```
[edit interfaces]
user@access-switch# set irb unit 0 family inet address 192.0.2.1/25
```
  7. Create the subnet for the support VLAN:
 

```
[edit interfaces]
user@access-switch# set irb unit 1 family inet address 192.0.2.129/25
```
  8. Configure the interfaces in the sales VLAN:
 

```
[edit interfaces]
user@access-switch# set ge-0/0/0 unit 0 description "Sales wireless access point port"
user@access-switch# set ge-0/0/0 unit 0 family ethernet-switching vlan members sales
user@access-switch# set ge-0/0/3 unit 0 description "Sales phone port"
user@access-switch# set ge-0/0/3 unit 0 family ethernet-switching vlan members sales
user@access-switch# set ge-0/0/20 unit 0 description "Sales file server port"
user@access-switch# set ge-0/0/20 unit 0 family ethernet-switching vlan members sales
user@access-switch# set ge-0/0/22 unit 0 description "Sales printer port"
user@access-switch# set ge-0/0/22 unit 0 family ethernet-switching vlan members sales
```
  9. Configure the interfaces in the support VLAN:
 

```
[edit interfaces]
user@access-switch# set ge-0/0/24 unit 0 description "Support wireless access point port"
user@access-switch# set ge-0/0/24 unit 0 family ethernet-switching vlan members support
user@access-switch# set ge-0/0/26 unit 0 description "Support phone port"
user@access-switch# set ge-0/0/26 unit 0 family ethernet-switching vlan members support
user@access-switch# set ge-0/0/44 unit 0 description "Support printer port"
user@access-switch# set ge-0/0/44 unit 0 family ethernet-switching vlan members support
user@access-switch# set ge-0/0/46 unit 0 description "Support file server port"
user@access-switch# set ge-0/0/46 unit 0 family ethernet-switching vlan members support
```

**Results** Display the results of the configuration:

```
user@access-switch> show configuration
interfaces {
  ge-0/0/0 {
    unit 0 {
      description "Sales wireless access point port";
      family ethernet-switching {
        vlan {
          members sales;
        }
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      description "Sales phone port";
      family ethernet-switching {
        vlan {
          members sales;
        }
      }
    }
  }
  ge-0/0/20 {
    unit 0 {
      description "Sales file server port";
      family ethernet-switching {
        vlan {
          members sales;
        }
      }
    }
  }
  ge-0/0/22 {
    unit 0 {
      description "Sales printer port";
      family ethernet-switching {
        vlan {
          members sales;
        }
      }
    }
  }
  ge-0/0/24 {
    unit 0 {
      description "Support wireless access point port";
      family ethernet-switching {
        vlan {
          members support;
        }
      }
    }
  }
  ge-0/0/26 {
    unit 0 {
      description "Support phone port";
      family ethernet-switching {
        vlan {
          members support;
        }
      }
    }
  }
}
```



```

    }
  }
}
ge-0/0/44 {
  unit 0 {
    description "Support printer port";
    family ethernet-switching {
      vlan {
        members support;
      }
    }
  }
}
ge-0/0/46 {
  unit 0 {
    description "Support file server port";
    family ethernet-switching {
      vlan {
        members support;
      }
    }
  }
}
ge-0/2/0 {
  native-vlan-id 1;
  unit 0 {
    description "Uplinking module connection to distribution switch";
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members [ 1 sales support ];
      }
    }
  }
}
irb {
  unit 0 {
    family inet {
      address 192.0.2.1/25;
    }
  }
  unit 1 {
    family inet {
      address 192.0.2.129/25;
    }
  }
}
}
vpls {
  sales {
    description "Sales VLAN";
    vlan-id 100;
    l3-interface irb.0;
  }
  support {
    description "Support VLAN";
    vlan-id 200;
    l3-interface irb.1;
  }
}
}

```



**TIP:** To quickly configure the access switch, issue the load merge terminal command, then copy the hierarchy and paste it into the switch terminal window.

## Configuring the Distribution Switch

To configure the distribution switch:

### CLI Quick Configuration

To quickly configure the distribution switch, copy the following commands and paste them into the switch terminal window:

```
set interfaces ge-0/0/0 unit 0 description "Connection to access switch"
set interfaces ge-0/0/0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members [ sales support ]
set interfaces ge-0/0/0 native-vlan-id 1
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members 1
set interfaces irb unit 0 family inet address 192.0.2.2/25
set interfaces irb unit 1 family inet address 192.0.2.130/25
set vlans sales description "Sales VLAN"
set vlans sales vlan-id 100
set vlans sales l3-interface irb.0
set vlans support description "Support VLAN"
set vlans support vlan-id 200
set vlans support l3-interface irb.1
```

### Step-by-Step Procedure

To configure the distribution switch:

1. Configure the interface on the switch to be the trunk port that connects to the access switch:  

```
[edit interfaces]
user@distribution-switch# set ge-0/0/0 unit 0 description "Connection to access switch"
user@distribution-switch# set ge-0/0/0 unit 0 family ethernet-switching interface-mode trunk
```
2. Specify the VLANs to be aggregated on the trunk port:  

```
[edit interfaces]
user@distribution-switch# set ge-0/0/0 unit 0 family ethernet-switching vlan members [ sales support ]
```
3. To handle untagged packets that are received on the trunk port, create a native VLAN by configuring a VLAN ID and specifying that the trunk port is a member of the native VLAN:  

```
[edit interfaces]
user@distribution-switch# set ge-0/0/0 native-vlan-id 1
user@distribution-switch# set ge-0/0/0 unit 0 family ethernet-switching vlan members 1
```
4. Configure the sales VLAN:  

```
[edit vlans]
user@distribution-switch# set sales description "Sales VLAN"
user@distribution-switch# set sales vlan-id 100
user@distribution-switch# set sales l3-interface irb.0
```

The VLAN configuration for the distribution switch includes the **set l3-interface irb.0** command to route traffic between the sales and support VLANs. The VLAN configuration for the access switch does not include this statement because the access switch is not monitoring IP addresses. Instead, the access switch is passing the IP addresses to the distribution switch for interpretation.

5. Configure the support VLAN:

```
[edit vlans]
user@distribution-switch# set support description "Support VLAN"
user@distribution-switch# set support vlan-id 200
user@distribution-switch# set support l3-interface irb.1
```

The VLAN configuration for the distribution switch includes the **set l3-interface irb.1** command to route traffic between the sales and support VLANs. The VLAN configuration for the access switch does not include this statement because the access switch is not monitoring IP addresses. Instead, the access switch is passing the IP addresses to the distribution switch for interpretation.

6. Create the subnet for the sales VLAN:

```
[edit interfaces]
user@distribution-switch# set irb unit 0 family inet address 192.0.2.2/25
```

7. Create the subnet for the support VLAN:

```
[edit interfaces]
user@distribution-switch# set irb unit 1 family inet address 192.0.2.130/25
```

**Results** Display the results of the configuration:

```
user@distribution-switch> show configuration
interfaces {
  ge-0/0/0 {
    native-vlan-id 1;
    unit 0 {
      description "Connection to access switch";
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members [ 1 sales support ];
        }
      }
    }
  }
  irb {
    unit 0 {
      family inet {
        address 192.0.2.2/25;
      }
    }
    unit 1 {
      family inet {
        address 192.0.2.130/25;
      }
    }
  }
}
vlans {
  sales {
    description "Sales VLAN";
    vlan-id 100;
    l3-interface irb.0;
  }
  support {
    description "Support VLAN";
    vlan-id 200;
    l3-interface irb.1;
  }
}
```



**TIP:** To quickly configure the distribution switch, issue the load merge terminal command, then copy the hierarchy and paste it into the switch terminal window.

---

### Verification

---

To confirm that the configuration is working properly, perform these tasks:

- [Verifying the VLAN Members and Interfaces on the Access Switch on page 2301](#)
- [Verifying the VLAN Members and Interfaces on the Distribution Switch on page 2301](#)

**Verifying the VLAN Members and Interfaces on the Access Switch**

**Purpose** Verify that the **sales** and **support** VLANs have been created on the switch.

**Action** List all VLANs configured on the switch:

```
user@access-switch> show vlans
```

| Routing instance | VLAN name | Tag | Interfaces                                                                |
|------------------|-----------|-----|---------------------------------------------------------------------------|
| default-switch   | sales     | 100 | ge-0/0/20.0<br>ge-0/0/22.0<br>ge-0/0/3.0*<br>ge-0/0/0.0*<br>ge-0/2/0.0*   |
| default-switch   | support   | 200 | ge-0/0/24.0<br>ge-0/0/26.0<br>ge-0/0/44.0*<br>ge-0/0/46.0*<br>ge-0/2/0.0* |

**Meaning** The output shows the **sales** and **support** VLANs and the interfaces that are configured as members of the respective VLANs.

**Verifying the VLAN Members and Interfaces on the Distribution Switch**

**Purpose** Verify that the **sales** and **support** VLANs have been created on the switch.

**Action** List all VLANs configured on the switch:

```
user@distribution-switch> show vlans
```

| Routing instance | VLAN name | Tag | Interfaces  |
|------------------|-----------|-----|-------------|
| default-switch   | sales     | 100 | ge-0/0/0.0* |
| default-switch   | support   | 200 | ge-0/0/0.0* |

**Meaning** The output shows the **sales** and **support** VLANs and the interface (ge-0/0/0.0) that is configured as a member of both VLANs. Interface ge-0/0/0.0 is also the trunk interface connected to the access switch.

- Related Documentation**
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 2281](#)
  - [Configuring VLANs for EX Series Switches \(CLI Procedure\) on page 2337](#)
  - [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)

## Example: Configuring a Private VLAN on a Single Switch



**NOTE:** This example uses Junos OS for switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX switch runs software that does not support ELS, see *Example: Configuring a Private VLAN on a Single EX Series Switch*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic or limit the communication between known hosts. Private VLANs (PVLANS) enable you to split a broadcast domain (primary VLAN) into multiple isolated broadcast subdomains (secondary VLANs), essentially putting a VLAN inside a VLAN.

This example describes how to create a PVLAN on a single switch:

- [Requirements on page 2302](#)
- [Overview and Topology on page 2302](#)
- [Configuration on page 2304](#)
- [Verification on page 2305](#)

### Requirements

This example uses the following hardware and software components:

- One Junos OS switch
- Junos OS Release 14.1X53-D10 or later for EX Series switches  
Junos OS Release 14.1X53-D15 or later for QFX Series switches

### Overview and Topology

You can isolate groups of subscribers for improved security and efficiency. This configuration example uses a simple topology to illustrate how to create a PVLAN with one primary VLAN and three secondary VLANs (one isolated VLAN, and two community VLANs).

[Table 236 on page 2302](#) lists the interfaces of the topology used in the example.

**Table 236: Interfaces of the Topology for Configuring a PVLAN**

| Interface            | Description                         |
|----------------------|-------------------------------------|
| ge-0/0/0             | Promiscuous member ports            |
| ge-1/0/0             |                                     |
| ge-0/0/11, ge-0/0/12 | HR community VLAN member ports      |
| ge-0/0/13, ge-0/0/14 | Finance community VLAN member ports |

Table 236: Interfaces of the Topology for Configuring a PVLAN (*continued*)

| Interface            | Description           |
|----------------------|-----------------------|
| ge-0/0/15, ge-0/0/16 | Isolated member ports |


Table 237 on page 2303 lists the VLAN IDs of the topology used in the example.

Table 237: VLAN IDs in the Topology for Configuring a PVLAN

| VLAN ID | Description            |
|---------|------------------------|
| 100     | Primary VLAN           |
| 200     | HR community VLAN      |
| 300     | Finance community VLAN |
| 400     | Isolated VLAN          |

Figure 35 on page 2303 shows the topology for this example.

Figure 35: Topology of a Private VLAN on a Single EX Series Switch



## Configuration

---

You can use an existing VLAN as the basis for your private PVLAN and create subdomains within it. This example creates a primary VLAN—using the VLAN name **vlan-pri**—as part of the procedure.

To configure a PVLAN, perform these tasks:

**CLI Quick Configuration** To quickly create and configure a PVLAN, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans vlan-pri vlan-id 100
set vlans vlan-iso private-vlan isolated vlan-id 400
set vlans vlan-hr private-vlan community vlan-id 200
set vlans vlan-finance private-vlan community vlan-id 300
set vlans vlan-pri vlan-id 100 isolated-vlan vlan-iso community-vlan vlan-hr community-vlan
vlan-finance
set interface ge-0/0/15 unit 0 family ethernet-switching interface-mode access vlan members
vlan-iso
set interface ge-0/0/16 unit 0 family ethernet-switching interface-mode access vlan members
vlan-iso
set interface ge-0/0/11 unit 0 family ethernet-switching interface-mode access vlan members
vlan-hr
set interface ge-0/0/12 unit 0 family ethernet-switching interface-mode trunk vlan members
vlan-hr
set interface ge-0/0/13 unit 0 family ethernet-switching interface-mode access vlan members
vlan-finance
set interface ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk vlan members
vlan-finance
set interface ge-0/0/0 unit 0 family ethernet-switching interface-mode trunk vlan members
vlan-pri
set interface ge-1/0/0 unit 0 family ethernet-switching interface-mode trunk vlan members
vlan-pri
```

**Step-by-Step Procedure** To configure the PVLAN:

1. Create the primary VLAN (in this example, the name is **vlan-pri**) of the private VLAN:  

```
[edit vlans]
user@switch# set vlan-pri vlan-id 100
```
2. Create an isolated VLAN and assign it a VLAN ID:  

```
[edit vlans]
user@switch# set vlan-iso private-vlan isolated vlan-id 400
```
3. Create the HR community VLAN and assign it a VLAN ID:  

```
[edit vlans]
user@switch# set vlan-hr private-vlan community vlan-id 200
```
4. Create the finance community VLAN and assign it a VLAN ID:  

```
[edit vlans]
user@switch# set vlan-finance private-vlan community vlan-id 300
```
5. Associate the secondary VLANs with the primary VLAN:  

```
[edit vlans]
user@switch# set vlan-pri vlan-id 100 isolated-vlan vlan-iso community-vlan vlan-hr
community-vlan vlan-finance
```
6. Set the interfaces to the appropriate interface modes:



[edit interfaces]

```
user@switch# set ge-0/0/15 unit 0 family ethernet-switching interface-mode access vlan members vlan-iso
```

```
user@switch# set ge-0/0/16 unit 0 family ethernet-switching interface-mode trunk vlan members vlan-iso
```

```
user@switch# set ge-0/0/11 unit 0 family ethernet-switching interface-mode access vlan members vlan-hr
```

```
user@switch# set ge-0/0/12 unit 0 family ethernet-switching interface-mode access vlan members vlan-hr
```

```
user@switch# set ge-0/0/13 unit 0 family ethernet-switching interface-mode access vlan members vlan-finance
```

```
user@switch# set ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk vlan members vlan-finance
```

7. Configure a promiscuous trunk interface of the primary VLAN. This interface is used by the primary VLAN to communicate with the secondary VLANs.

```
user@switch# set ge-0/0/0 unit 0 family ethernet-switching interface-mode trunk vlan members vlan-pri
```

8. Configure another trunk interface (it is also a promiscuous interface) of the primary VLAN, connecting the PVLAN to the router.

```
user@switch# set ge-1/0/0 unit 0 family ethernet-switching interface-mode trunk vlan members vlan-pri
```

### Results

Check the results of the configuration:

```
user@switch> show configuration
```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Private VLAN and Secondary VLANs Were Created on page 2305](#)

#### *Verifying That the Private VLAN and Secondary VLANs Were Created*

**Purpose** Verify that the primary VLAN and secondary VLANs were properly created on the switch.

**Action** Use the `show vlans` command:

```
user@switch> show vlans extensive
```

**Meaning** The output shows that the primary VLAN was created and identifies the interfaces and secondary VLANs associated with it.

#### Related Documentation

- [Understanding Private VLANs on page 2258](#)
- [Creating a Private VLAN on a Single Switch \(CLI Procedure\) on page 2343](#)
- [Creating a Private VLAN Spanning Multiple Switches \(CLI Procedure\) on page 2345](#)

## Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches

Virtual routing instances enable an EX Series switch to have multiple routing tables. With virtual routing instances, you can segment your network to isolate traffic without setting up additional devices.

In a large office, you might need multiple VLANs to properly manage your traffic. This example describes how to create a virtual routing instance associated with each VLAN for this large office:

- [Requirements on page 2306](#)
- [Overview and Topology on page 2306](#)
- [Configuration on page 2307](#)
- [Verification on page 2309](#)

---

### Requirements

This example uses the following hardware and software components:

- One EX Series switch
- Junos OS Release 9.2 or later for EX Series switches

Before you create the virtual routing instances, make sure you have:

- Configured the necessary VLANs. See *Configuring VLANs for EX Series Switches (CLI Procedure)*, “[Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)” on page 2337, or “[Configuring VLANs for EX Series Switches \(J-Web Procedure\)](#)” on page 2334.

---

### Overview and Topology

This configuration example shows a simple large office topology wherein a LAN is segmented into two VLANs, each of which is associated with an interface and a virtual routing instance.

For example, VLAN 1031 is associated with interface ge-0/0/3.1 and virtual routing instance **r1**. VLAN 1032 is associated with interface ge-0/0/3.2 and virtual routing instance **r2**.

This example also includes interfaces ge-0/0/1.0 and ge-0/0/2.0. Although these interfaces are not part of either VLAN, interface ge-0/0/1.0 is associated with virtual routing instance **r1**, while interface ge-0/0/2.0 is associated with virtual routing instance **r2**.

This example also shows how to use policy statements to import routes from virtual routing instance **r1** to **r2** and from virtual routing instance **r2** to **r1**.



**NOTE:** On EX Series switches, importing directly connected routes from one virtual routing instance to another is not supported. (A directly connected route is a route that is created by specifying an IP address on one of the switch interfaces or for a VLAN.) For more information, including a workaround for this situation, see [KB23027 - Exchanging \(leaking\) directly connected routes across routing instances is not supported](#).

## Configuration

### CLI Quick Configuration

To quickly create and configure virtual routing instances, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/3 vlan-tagging
set interfaces ge-0/0/3 unit 1 vlan-id 1031 family inet address 10.3.1.1/24
set interfaces ge-0/0/3 unit 2 vlan-id 1032 family inet address 10.4.1.1/24
set interfaces ge-0/0/1 unit 0 family inet address 10.1.1.1/24
set interfaces ge-0/0/2 unit 0 family inet address 10.2.1.1/24
set routing-instances r1 instance-type virtual-router
set routing-instances r1 interface ge-0/0/1.0
set routing-instances r1 interface ge-0/0/3.1
set routing-instances r1 routing-options instance-import import-from-r2
set routing-instances r2 instance-type virtual-router
set routing-instances r2 interface ge-0/0/2.0
set routing-instances r2 interface ge-0/0/3.2
set routing-instances r2 routing-options instance-import import-from-r1
set policy-options policy-statement import-from-r1 term 1 from instance r1
set policy-options policy-statement import-from-r1 term 1 then accept
set policy-options policy-statement import-from-r2 term 1 from instance r2
set policy-options policy-statement import-from-r2 term 1 then accept
```

### Step-by-Step Procedure

To configure virtual routing instances:

1. Create a VLAN-tagged interface:

```
[edit]
user@switch# set interfaces ge-0/0/3 vlan-tagging
```

2. Create one or more logical interfaces on the interfaces to be included in each virtual routing instance:

```
[edit]
user@switch# set interfaces ge-0/0/3 unit 1 vlan-id 1031 family inet address 10.3.1.1/24
user@switch# set interfaces ge-0/0/3 unit 2 vlan-id 1032 family inet address 10.4.1.1/24
user@switch# set interfaces ge-0/0/1 unit 0 family inet address 10.1.1.1/24
user@switch# set interfaces ge-0/0/2 unit 0 family inet address 10.2.1.1/24
```

3. Create two virtual routing instances:

```
[edit]
user@switch# set routing-instances r1 instance-type virtual-router
user@switch# set routing-instances r2 instance-type virtual-router
```

4. Set the interfaces for the virtual routing instances:

```
[edit]
user@switch# set routing-instances r1 interface ge-0/0/1.0
user@switch# set routing-instances r1 interface ge-0/0/3.1
user@switch# set routing-instances r2 interface ge-0/0/2.0
user@switch# set routing-instances r2 interface ge-0/0/3.2
```

5. Apply a policy to routes being imported into each of the virtual routing instances:

```
[edit]
```

```
user@switch# set routing-instances r1 routing-options instance-import import-from-r2
```

```
user@switch# set routing-instances r2 routing-options instance-import import-from-r1
```

6. Create a policy for importing routes from virtual routing instance r1 to r2 and another policy for importing routes from virtual routing instance r2 to r1:

```
[edit]
```

```
user@switch# set policy-options policy-statement import-from-r1 term 1 from instance r1
```

```
user@switch# set policy-options policy-statement import-from-r1 term 1 then accept
```

```
user@switch# set policy-options policy-statement import-from-r2 term 1 from instance r2
```

```
user@switch# set policy-options policy-statement import-from-r2 term 1 then accept
```

**Results** Check the results of the configuration:

```
user@switch> show configuration
```

```
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.1.1/24;
      }
    }
  }
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 10.2.1.1/24;
      }
    }
  }
  ge-1/0/3 {
    vlan-tagging;
    unit 1 {
      vlan-id 1031;
      family inet {
        address 10.3.1.1/24;
      }
    }
    unit 2 {
      vlan-id 1032;
      family inet {
        address 10.4.1.1/24;
      }
    }
  }
}
policy-options {
  policy-statement import-from-r1 {
    term 1 {
      from instance r1;
      then accept;
    }
  }
  policy-statement import-from-r2 {
    term 1 {
      from instance r2;
      then accept;
    }
  }
}
```

```

    }
  }
  routing-instances {
    r1 {
      instance-type virtual-router;
      interface ge-0/0/1.0;
      interface ge-0/0/3.1;
      routing-options {
        instance-import import-from-r2;
      }
    }
    r2 {
      instance-type virtual-router;
      interface ge-0/0/2.0;
      interface ge-0/0/3.2;
      routing-options {
        instance-import import-from-r1;
      }
    }
  }
}

```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Routing Instances Were Created on page 2309](#)

#### *Verifying That the Routing Instances Were Created*

**Purpose** Verify that the virtual routing instances were properly created on the switch.

**Action** Use the `show route instance` command:

```

user@switch> show route instance
Instance      Primary RIB      Type      Active/holdown/hidden
master
  inet.0      forwarding      6/0/0
  iso.0       1/0/0
  inet6.0     2/0/0
...
r1
  r1.inet.0   virtual-router  7/0/0
r2
  r2.inet.0   virtual-router  7/0/0

```

**Meaning** Each routing instance created is displayed, along with its type, information about whether it is active or not, and its primary routing table.

**Related Documentation**

- [Configuring Virtual Routing Instances \(CLI Procedure\) on page 2347](#)

## Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches

---



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

---

As a network expands and the number of clients and VLANs increases, VLAN administration becomes complex and the task of efficiently configuring VLANs on multiple EX Series switches becomes increasingly difficult. However, you can automate VLAN administration by enabling Multiple VLAN Registration Protocol (MVRP) on the network.

MVRP also dynamically creates VLANs, further simplifying the network overhead required to statically configure VLANs.

---



**NOTE:** Only trunk interfaces can be enabled for MVRP.

---

This example describes how to use MVRP to automate administration of VLAN membership changes within your network and how to use MVRP to dynamically create VLANs:

- [Requirements on page 2310](#)
- [Overview and Topology on page 2311](#)
- [Configuring VLANs and MVRP on Access Switch A on page 2313](#)
- [Configuring VLANs and MVRP on Access Switch B on page 2315](#)
- [Configuring VLANs and MVRP on Distribution Switch C on page 2317](#)
- [Verification on page 2318](#)

### Requirements

---

This example uses the following hardware and software components:

- Two EX Series access switches
- One EX Series distribution switch
- Junos OS Release 13.2X50-D10 or later for EX Series switches

Before you configure MVRP on an interface, you must enable one of the following spanning tree protocols on that interface:

- Rapid Spanning-Tree Protocol (RSTP). For more information about RSTP, see [“Understanding RSTP for EX Series Switches” on page 4906](#).
- Multiple Spanning-Tree Protocol (MSTP). For more information about MSTP, see [“Understanding MSTP for EX Series Switches” on page 4904](#).

## Overview and Topology

MVRP is used to manage dynamic VLAN registration in a LAN. It can also be used to dynamically create VLANs.

This example uses MVRP to dynamically create VLANs on the switching network. Alternatively, you can disable dynamic VLAN creation and create VLANs statically. Enabling MVRP on the trunk interface of each switch in your switching network ensures that the active VLAN information for the switches in the network is propagated to each switch through the trunk interfaces, assuming dynamic VLAN creation is enabled for MVRP.

MVRP ensures that the VLAN membership information on the trunk interface is updated as the switch's access interfaces become active or inactive in the configured VLANs in a static or dynamic VLAN creation setup.

You do not need to explicitly bind a VLAN to the trunk interface. When MVRP is enabled, the trunk interface advertises all the VLANs that are active (bound to access interfaces) on that switch. An MVRP-enabled trunk interface does not advertise VLANs that are configured on the switch but are not currently bound to an access interface. Thus, MVRP provides the benefit of reducing network overhead—by limiting the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

When VLAN access interfaces become active or inactive, MVRP ensures that the updated information is advertised on the trunk interface. Thus, in this example, distribution Switch C does not forward traffic to inactive VLANs.



**NOTE:** This example shows a network with three VLANs: **finance**, **sales**, and **lab**. All three VLANs are running the same version of Junos OS. If switches in this network were running a mix of Junos OS releases that included Release 11.3, additional configuration would be necessary—see *Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)* for details.

Figure 36 on page 2312 shows MVRP configured on two access switches and one distribution switch.

Figure 36: MVRP Configured on Two Access Switches and One Distribution Switch for Automatic VLAN Administration




Table 238 on page 2312 explains the components of the example topology.

Table 238: Components of the Network Topology

| Settings               | Settings                                                                                                                  |
|------------------------|---------------------------------------------------------------------------------------------------------------------------|
| Switch hardware        | <ul style="list-style-type: none"> <li>Access Switch A</li> <li>Access Switch B</li> <li>Distribution Switch C</li> </ul> |
| VLAN names and tag IDs | <b>finance</b> , tag 100<br><b>lab</b> , tag 200<br><b>sales</b> , tag 300                                                |



Table 238: Components of the Network Topology (*continued*)

| Settings   | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interfaces | <p>Access Switch A interfaces:</p> <ul style="list-style-type: none"> <li>• ge-0/0/1—Connects PC1 to access Switch A.</li> <li>• ge-0/0/2—Connects PC2 to access Switch A.</li> <li>• ge-0/0/3—Connects PC3 to access Switch A.</li> <li>• xe-0/1/1—Connects access Switch A to distribution Switch C (trunk).</li> </ul> <p>Access Switch B interfaces:</p> <ul style="list-style-type: none"> <li>• ge-0/0/0—Connects PC4 to access Switch B.</li> <li>• ge-0/0/1—Connects PC5 to access Switch B.</li> <li>• ge-0/0/2—Reserved for future use,</li> <li>• xe-0/1/0—Connects access Switch B to distribution Switch C. (trunk)</li> </ul> <p>Distribution Switch C interfaces:</p> <ul style="list-style-type: none"> <li>• xe-0/1/1—Connects distribution Switch C to access Switch A. (trunk)</li> <li>• xe-0/1/0—Connects distribution Switch C to access Switch B. (trunk)</li> </ul> |

### Configuring VLANs and MVRP on Access Switch A

To configure VLANs on the switch, bind access interfaces to the VLANs, and enable MVRP on the trunk interface of access Switch A, perform these tasks:

#### CLI Quick Configuration

To quickly configure access Switch A for MVRP, copy the following commands and paste them into the switch terminal window of Switch A:

```
[edit]
set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members finance
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members lab
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/1/1 unit 0 family ethernet-switching interface-mode trunk
set protocols mvrp interface xe-0/1/1
```



**NOTE:** This example uses default MVRP timers. The default values associated with each MVRP timer are: 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms (10 seconds) for the leaveall timer. We recommend retaining the use of default timer values as modifying timers to inappropriate values might cause an imbalance in the operation of MVRP. However, if you choose to change the default settings, keep in mind that on an EX Series switch that uses Junos OS with support for ELS, if the timer value set on an interface level is different from the value set on a switch level, then the value on the interface level takes precedence.

**Step-by-Step  
Procedure**

To configure access Switch A for MVRP:

1. Configure the finance VLAN:  

```
[edit]
user@Access-Switch-A# set vlans finance vlan-id 100
```
2. Configure the lab VLAN:  

```
[edit]
user@Access-Switch-A# set vlans lab vlan-id 200
```
3. Configure the sales VLAN:  

```
[edit]
user@Access-Switch-A# set vlans sales vlan-id 300
```
4. Configure an Ethernet interface as a member of the finance VLAN:  

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan
members finance
```
5. Configure an Ethernet interface as a member of the lab VLAN:  

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan
members lab
```
6. Configure an Ethernet interface as a member of the sales VLAN:  

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan
members sales
```
7. Configure a trunk interface:  

```
[edit]
user@Access-Switch-A# set interfaces xe-0/1/1 unit 0 family ethernet-switching
interface-mode trunk
```
8. Enable MVRP on the trunk interface:  

```
[edit]
user@Access-Switch-A# set protocols mvrp interface xe-0/1/1
```

**Results** Check the results of the configuration on Switch A:

```
[edit]
user@Access-Switch-A# show
interfaces {
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members finance;
        }
      }
    }
  }
  ge-0/0/2 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members lab;
        }
      }
    }
  }
}
```

```

    }
  }
  ge-0/0/3 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members sales;
        }
      }
    }
  }
  xe-0/1/1 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
      }
    }
  }
}
protocols {
  mvrp {
    interface xe-0/1/1;
  }
}
vlangs {
  finance {
    vlan-id 100;
  }
  lab {
    vlan-id 200;
  }
  sales {
    vlan-id 300;
  }
}

```

### Configuring VLANs and MVRP on Access Switch B

To configure three VLANs on the switch, bind access interfaces for PC4 and PC5 to the VLANs, and enable MVRP on the trunk interface of access Switch B, perform these tasks:

#### CLI Quick Configuration

To quickly configure Access Switch B for MVRP, copy the following commands and paste them into the switch terminal window of Switch B:

```

[edit]
set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members finance
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members lab
set interfaces xe-0/1/0 unit 0 family ethernet-switching interface-mode trunk
set protocols mvrp interface xe-0/1/0

```

**Step-by-Step  
Procedure**

To configure access Switch B for MVRP:

1. Configure the finance VLAN:  

```
[edit]  
user@Access-Switch-B# set vlans finance vlan-id 100
```
2. Configure the lab VLAN:  

```
[edit]  
user@Access-Switch-B# set vlans lab vlan-id 200
```
3. Configure the sales VLAN:  

```
[edit]  
user@Access-Switch-B# set vlans sales vlan-id 300
```
4. Configure an Ethernet interface as a member of the finance VLAN:  

```
[edit]  
user@Access-Switch-B# set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan  
members finance
```
5. Configure an Ethernet interface as a member of the lab VLAN:  

```
[edit]  
user@Access-Switch-B# set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan  
members lab
```
6. Configure a trunk interface:  

```
user@Access-Switch-B# set interfaces xe-0/1/0 unit 0 family ethernet-switching  
interface-mode trunk
```
7. Enable MVRP on the trunk interface:  

```
[edit]  
user@Access-Switch-B# set protocols mvrp xe-0/1/0
```



**NOTE:** This example uses default MVRP timers. The default values associated with each MVRP timer are: 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms (10 seconds) for the leaveall timer. We recommend retaining the use of default timer values as modifying timers to inappropriate values might cause an imbalance in the operation of MVRP. However, if you choose to change the default values, keep in mind that on an EX Series switch that uses Junos OS with support for ELS, if the timer value set on an interface level is different from the value set on a switch level, then the value on the interface level takes precedence.

**Results** Check the results of the configuration for Switch B:

```
[edit]  
user@Access-Switch-B# show  
interfaces {  
  ge-0/0/0 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members finance;
```

```

    }
  }
}
ge-0/0/1 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members lab;
      }
    }
  }
}
xe-0/1/0 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
    }
  }
}
}
protocols {
  mvrp {
    interface xe-0/1/0;
  }
}
vlands {
  finance {
    vlan-id 100;
  }
  lab {
    vlan-id 200;
  }
  sales {
    vlan-id 300;
  }
}
}

```

### Configuring VLANS and MVRP on Distribution Switch C

**CLI Quick Configuration** To quickly configure distribution Switch C for MVRP, copy the following commands and paste them into the switch terminal window of distribution Switch C:

```

[edit]
set interfaces xe-0/1/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/1/0 unit 0 family ethernet-switching interface-mode trunk
set protocols mvrp interface xe-0/1/1
set protocols mvrp interface xe-0/1/0

```

**Step-by-Step  
Procedure**

To configure distribution Switch C for MVRP:

1. Configure the trunk interface to access Switch A:  

```
[edit]
user@Distribution-Switch-C# set interfaces xe-0/1/1 unit 0 family ethernet-switching
interface-mode trunk
```
2. Configure the trunk interface to access Switch B:  

```
[edit]
user@Distribution-Switch-C# set interfaces xe-0/1/0 unit 0 family ethernet-switching
interface-mode trunk
```
3. Enable MVRP on the trunk interface for xe-0/1/1 :  

```
[edit]
user@Distribution-Switch-C# set protocols mvrp interface xe-0/1/1
```
4. Enable MVRP on the trunk interface for xe-0/1/0 :  

```
[edit]
user@Distribution-Switch-C# set protocols mvrp interface xe-0/1/0
```

**Results**

Check the results of the configuration for Switch C:

```
[edit]
user@Distribution Switch-C# show
interfaces {
  xe-0/1/0 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
      }
    }
  }
  xe-0/1/1 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
      }
    }
  }
}
protocols {
  mvrp {
    interface xe-0/1/0;
    interface xe-0/1/1;
  }
}
```

---

**Verification**

To confirm that the configuration is updating VLAN membership, perform these tasks:

- [Verifying That MVRP Is Enabled on Access Switch A on page 2319](#)
- [Verifying That MVRP Is Updating VLAN Membership on Access Switch A on page 2319](#)
- [Verifying That MVRP Is Enabled on Access Switch B on page 2320](#)
- [Verifying That MVRP Is Updating VLAN Membership on Access Switch B on page 2320](#)

- [Verifying That MVRP Is Enabled on Distribution Switch C on page 2321](#)
- [Verifying That MVRP Is Updating VLAN Membership on Distribution Switch C on page 2321](#)

### *Verifying That MVRP Is Enabled on Access Switch A*

**Purpose** Verify that MVRP is enabled on the switch.

**Action** Show the MVRP configuration:

```
user@Access-Switch-A> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface   Join   Leave  LeaveAll
  xe-0/1/1    200   1000   10000
```

**Meaning** The results show that MVRP is enabled on the trunk interface of Switch A and that the default timers are used.

### *Verifying That MVRP Is Updating VLAN Membership on Access Switch A*

**Purpose** Verify that MVRP is updating VLAN membership by displaying the Ethernet switching interfaces and associated VLANs that are active on Switch A.

**Action** List Ethernet switching interfaces on the switch:

```
user@Access-Switch-A> show ethernet-switching interface
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members limit state interface flags
ge-0/0/1.0
      finance  100
                        65535   Forwarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members limit state interface flags
ge-0/0/2.0
      lab      200
                        65535   Forwarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members limit state interface flags
ge-0/0/3.0
      sales    300
                        65535   Forwarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members limit state interface flags
xe-0/1/1.0
                        65535   tagged
```

|         |     |       |            |
|---------|-----|-------|------------|
| finance | 100 | 65535 | Forwarding |
| lab     | 200 | 65535 | Forwarding |

**Meaning** MVRP has automatically added **finance** and **lab** as VLAN members on the trunk interface because they are being advertised by access Switch B.

#### *Verifying That MVRP Is Enabled on Access Switch B*

**Purpose** Verify that MVRP is enabled on the switch.

**Action** Show the MVRP configuration:

```
user@Access-Switch-B> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface    Join    Leave    LeaveAll
  xe-0/1/0     200    1000    10000
```

**Meaning** The results show that MVRP is enabled on the trunk interface of Switch B and that the default timers are used.

#### *Verifying That MVRP Is Updating VLAN Membership on Access Switch B*

**Purpose** Verify that MVRP is updating VLAN membership by displaying the Ethernet switching interfaces and associated VLANs that are active on Switch B.

**Action** List Ethernet switching interfaces on the switch:

```
user@Access-Switch-B> show ethernet-switching interface
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
ge-0/0/0.0
      finance  100
                        65535 Forwarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
ge-0/0/1.0
      lab      200
                        65535 Forwarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/1/0.0
      finance  100
                        65535 Forwarding
      lab      200
```



|       |     |       |            |
|-------|-----|-------|------------|
|       |     | 65535 | Forwarding |
| sales | 300 |       |            |
|       |     | 65535 | Forwarding |

**Meaning** MVRP has automatically added **finance**, **lab**, and **sales** as VLAN members on the trunk interface because they are being advertised by access Switch A.

#### *Verifying That MVRP Is Enabled on Distribution Switch C*

**Purpose** Verify that MVRP is enabled on the switch.

**Action** Show the MVRP configuration:

```
user@Distribution-Switch-C> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface    Join    Leave  LeaveAll
  xe-0/1/1     200    1000   10000
  xe-0/1/0     200    1000   10000
```

**Meaning** The results show that MVRP is enabled on the trunk interfaces of Switch C and that the default timers are used.

#### *Verifying That MVRP Is Updating VLAN Membership on Distribution Switch C*

**Purpose** Verify that MVRP is updating VLAN membership on distribution Switch C by displaying the Ethernet switching interfaces and associated VLANs on distribution Switch C.

**Action** List the Ethernet switching interfaces on the switch:

```
user@Distribution-Switch-C> show ethernet-switching interface
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members limit state interface flags
xe-0/1/1.0
  mvrp_100
    65535 Forwarding
  mvrp_200
    65535 Forwarding
  mvrp_300
    65535 Forwarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members limit state interface flags
xe-0/1/0.0
  mvrp_100
    65535 Forwarding
  mvrp_200
    65535 Forwarding
```

List the VLANs that were created dynamically using MVRP on the switch:

```
user@Distribution-Switch-C> show mvrp dynamic-vlan-memberships
```

```
MVRP dynamic vlans for routing instance 'default-switch'
```

```
(s) static vlan, (f) fixed registration
```

| VLAN ID | Interfaces |
|---------|------------|
| 100     | xe-0/1/1.0 |
|         | xe-0/1/0.0 |
| 200     | xe-0/1/1.0 |
|         | xe-0/1/0.0 |
| 300     | xe-0/1/1.0 |

Note that this scenario does not have any fixed registration, which is typical when MVRP is enabled.

**Meaning** Distribution Switch C has two trunk interfaces. Interface **xe-0/1/1.0** connects Distribution Switch C to Access Switch A and is, therefore, updated to show that it is a member of all the VLANs that are active on Switch A. Any traffic for those VLANs will be passed on from Switch C to Switch A, through interface **xe-0/1/1.0**. Interface **xe-0/1/0.0** connects Switch C to Switch B and is updated to show that it is a member of the two VLANs that are active on Switch B. Thus, Switch C sends traffic for **finance** and **lab** to both Switch A and Switch B. But Switch C sends traffic for **sales** only to Switch A.

Switch C also has three dynamic VLANs created using MVRP: **mvrp\_100**, **mvrp\_200**, and **mvrp\_300**. The dynamically created VLANs **mvrp\_100** and **mvrp\_200** are active on interfaces **xe-0/1/1.0** and **xe-0/1/0.0**, and dynamically created VLAN **mvrp\_300** is active on interface **xe-0/1/1.0**.

- Related Documentation**
- [Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure \) on page 2348](#)
  - [Understanding Multiple VLAN Registration Protocol \(MVRP\) on EX Series Switches on page 2264](#)

## Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches



**NOTE:** This topic applies to Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider’s cloud and deliver them to EX Series switches that are not part of the local broadcast domain. This feature is useful when you want to run Layer 2 protocols on a network that includes switches located at remote sites that are connected across a service provider network.

This example describes how to configure L2PT on EX Series switches that are running Junos OS with support for ELS:

- [Requirements on page 2323](#)
- [Overview and Topology on page 2323](#)
- [Configuration on page 2325](#)
- [Verification on page 2325](#)

### Requirements

---

This example uses the following hardware and software components:

- Six EX Series switches, with three each at two customer sites, and with one of the switches at each site designated as the provider edge (PE) device
- Junos OS Release 14.1X53-D10 or later for EX Series switches

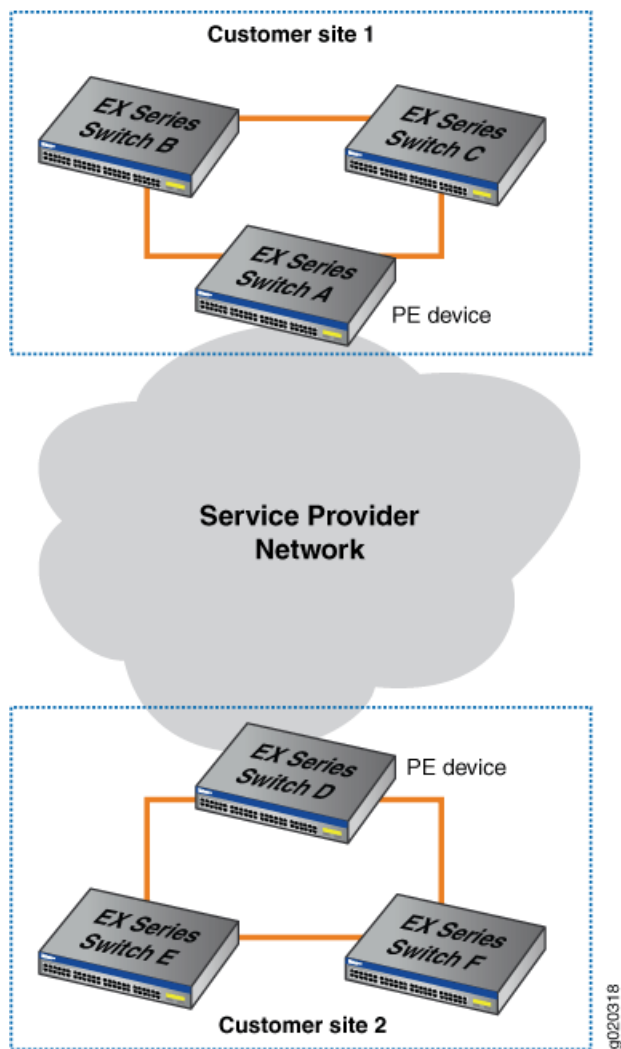
### Overview and Topology

---

L2PT enables you to send Layer 2 PDUs across a service provider network and deliver them to EX Series switches that are not part of the local broadcast domain.

[Figure 37 on page 2324](#) shows a customer network that includes two sites that are connected across a service provider network. Site 1 contains three switches connected in a Layer 2 network, with Switch A designated as a provider edge (PE) device in the service provider network. Site 2 contains a Layer 2 network with a similar topology to that of Site 1, with Switch D designated as a PE device.

Figure 37: L2PT Topology



Q-in-Q tunneling must be enabled before you configure L2PT. L2PT tunneling can be configured *only* on interfaces that have already been configured to support Q-in-Q tunneling.

Q-in-Q tunneling ensures that Switches A, B, C, D, E, and F are part of the same broadcast domain.

This example uses Spanning Tree Protocol (STP) as the Layer 2 protocol being tunneled, but you could substitute any of the supported protocols for STP.



**NOTE:** The implementation of L2PT on EX4300 switches does not support the drop-threshold or the shutdown-threshold configuration statements.

## Configuration

Before you begin, you must configure Q-in-Q tunneling. See [“Configuring Q-in-Q Tunneling \(CLI Procedure\)” on page 2351](#) for a step-by-step procedure for configuring Q-in-Q tunneling.

To configure L2PT, perform these tasks:

**CLI Quick Configuration** To quickly configure L2PT on a Q-in-Q interface, copy the following command and paste it into the switch terminal window of each PE device, and commit the configuration.

```
[edit protocols]
user@switch# set layer2-control mac-rewrite interface ge-0/0/0 protocol STP
```

**Step-by-Step Procedure** To configure L2PT, perform these tasks on each PE device (in [Figure 37 on page 2324](#), Switch A and Switch D are the PE devices):

1. Configure L2PT on a specified Q-in-Q interface by enabling MAC address rewriting for Layer 2 protocol tunneling and selecting the Layer 2 protocol to be tunneled:



**NOTE:** You can select only one layer 2 protocol at a time. If you want an interface to support more than one layer 2 protocol, you must enter the `mac-rewrite` statement multiple times to select the desired protocols.

```
[edit protocols]
user@switch# set layer2-control mac-rewrite interface interface-name protocol STP
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
protocols {
  layer2-control {
    mac-rewrite {
      interface ge-0/0/1 {
        protocol {
          stp;
        }
      }
    }
  }
}
```

## Verification

To verify that L2PT is working correctly, perform this task:

- [Verify That L2PT Is Tunneling the Desired Layer 2 Protocol on page 2325](#)

### *Verify That L2PT Is Tunneling the Desired Layer 2 Protocol*

**Purpose** Check to see that L2PT is tunneling the desired protocol on a specified Q-in-Q interface:

**Action** Check to see that L2PT is tunneling STP on an interface:

```
user@switchA> show mac-rewrite interface ge-0/0/0
Interface      Protocols
ge-0/0/0       STP
```

**Meaning** The output shows that STP is being tunneled on the ge-0/0/0 interface.

**Related Documentation**

- [Configuring Layer 2 Protocol Tunneling on EX Series Switches \(CLI Procedure\)](#)
- [Understanding Layer 2 Protocol Tunneling on EX Series Switches](#)

## Example: Configuring Redundant Trunk Links for Faster Recovery



**NOTE:** This example uses Junos OS for EX Series switches or QFX Series with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see *Example: Configuring Redundant Trunk Links for Faster Recovery*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

You can manage network convergence by configuring both a primary link and a secondary link on a switch; this is called a redundant trunk group (RTG). If the primary link in a redundant trunk group fails, it passes its known MAC address locations to the secondary link, which automatically takes over after one minute.

This example describes how to create a redundant trunk group with a primary and a secondary link:

- [Requirements on page 2326](#)
- [Overview and Topology on page 2327](#)
- [Disabling RSTP on Switches 1 and 2 on page 2329](#)
- [Configuring Redundant Trunk Links on Switch 3 on page 2329](#)
- [Verification on page 2330](#)

### Requirements

---

This example uses the following hardware and software components:

- Two EX Series or QFX Series distribution switches
- One EX Series or QFX Series access switch
- The appropriate software release for your platform:
  - For EX Series switches: Junos OS Release 13.2X50-D10 or later
  - For the QFX Series: Junos OS Release 13.2X50-D15 or later

Before you configure the redundant trunk links network on the access and distribution switches, be sure you have:

- Configured interfaces ge-0/0/9 and ge-0/0/10 on the access switch, Switch 3, as trunk interfaces. .
- Configured one trunk interface on each distribution switch, Switch 1 and Switch 2.
- Connected the three switches as shown in the topology for this example (see [Figure 38 on page 2328](#)).

### Overview and Topology

In a typical enterprise network composed of distribution and access layers, a redundant trunk link provides a simple solution for trunk interface network recovery. When a trunk interface fails, data traffic is routed to another trunk interface after one minute, thereby keeping network convergence time to a minimum.

This example shows the configuration of a redundant trunk group that includes one primary link (and its interface) and one unspecified link (and its interface) that serves as the secondary link.

A second type of redundant trunk group, not illustrated in the example, consists of two unspecified links (and their interfaces); in this case, neither of the links is primary. The software selects an active link by comparing the port numbers of the two links and activating the link with the higher port number. For example, if the two link interfaces use interfaces ge-0/1/0 and ge-0/1/1, the software activates ge-0/1/1. (In the interface names, the final number is the port number.)

The two links in a redundant trunk group generally operate the same way, whether they are configured as primary/unspecified or unspecified/unspecified. Data traffic initially passes through the active link but is blocked on the inactive link. While data traffic is blocked on the secondary link, note that Layer 2 control traffic is still permitted if the link is active. For example, an LLDP session can be run between two switches on the secondary link. If the active link either goes down or is disabled administratively, it broadcasts a list of its known MAC addresses for data traffic; the other link immediately picks up and adds the MAC addresses to its address table, becomes active, and begins forwarding traffic.

The one difference in operation between the two types of redundant trunk groups occurs when a primary link is active, goes down, is replaced by the secondary link, and then reactivates. When a primary link is re-enabled like this while the secondary link is active, the primary link waits 2 minutes (you can change the time interval by using the preempt cutover timer to accommodate your network) and then takes over as the active link. In other words, the primary link has priority and is always activated if it is available. This differs from the behavior of two unspecified links, both of which act as equals. Because the unspecified links are equal, the active link remains active until it either goes down or is disabled administratively; this is the only time that the other unspecified link learns the MAC addresses and immediately becomes active.

The example given here illustrates a primary/unspecified configuration for a redundant trunk group because that configuration gives you more control and is more commonly used.



**NOTE:** Rapid Spanning Tree Protocol (RSTP) is enabled by default on the switches to create a loop-free topology, but an interface is not allowed to be in both a redundant trunk group and in a spanning-tree protocol topology at the same time. You will need to disable RSTP on the two distribution switches in the example, Switch 1 and Switch 2. Spanning-tree protocols can, however, continue operating in other parts of the network—for example, between the distribution switches and also in links between distribution switches and the enterprise core.

Figure 38 on page 2328 displays an example topology containing three switches. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer. Switch 3 is connected to the distribution layer through trunk interfaces ge-0/0/9.0 (Link 1) and ge-0/0/10.0 (Link 2).

Table 239 on page 2329 lists the components used in this redundant trunk group.

Because RSTP and RTGs cannot operate simultaneously on a switch, you disable RSTP on Switch 1 and Switch 2 in the first configuration task, and you disable RSTP on Switch 3 in the second task.

The second configuration task creates a redundant trunk group called example 1 on Switch 3. The trunk interfaces ge-0/0/9.0 and ge-0/0/10.0 are the two links configured in the second configuration task. You configure the trunk interface ge-0/0/9.0 as the primary link. You configure the trunk interface ge-0/0/10.0 as an unspecified link, which becomes the secondary link by default.

Figure 38: Topology for Configuring the Redundant Trunk Links

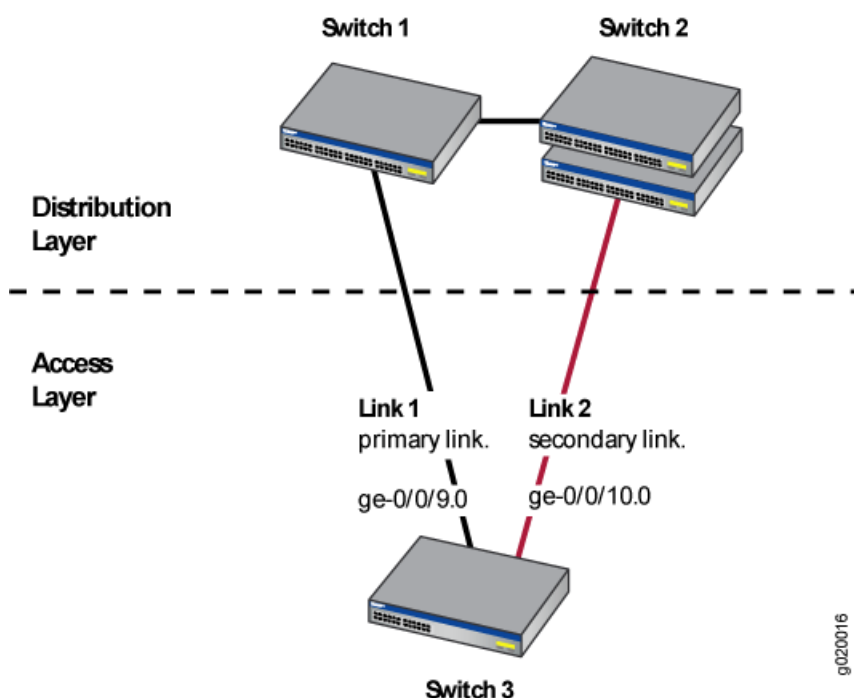




Table 239: Components of the Redundant Trunk Link Topology

| Property              | Settings                                                                                                                                                                                                                           |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch hardware       | <ul style="list-style-type: none"> <li>Switch 1–1 EX Series or QFX Series distribution switch</li> <li>Switch 2–1 EX Series or QFX Series distribution switch</li> <li>Switch 3–1 EX Series or QFX Series access switch</li> </ul> |
| Trunk interfaces      | On Switch 3 (access switch): ge-0/0/9.0 and ge-0/0/10.0                                                                                                                                                                            |
| Redundant trunk group | example1                                                                                                                                                                                                                           |

### Disabling RSTP on Switches 1 and 2

To disable RSTP on Switch 1 and Switch 2, perform this task on each switch:

#### CLI Quick Configuration

To quickly disable RSTP on Switch 1 and Switch 2, copy the following command and paste it into each switch terminal window:

```
[edit]
set protocols rstp disable
```

#### Step-by-Step Procedure

To disable RSTP on Switch 1 and Switch 2:

1. Disable RSTP on Switch 1 and Switch 2:

```
[edit]
user@switch# set protocols rstp disable
```

#### Results

Check the results of the configuration:

```
[edit]
user@switch# show
protocols {
  rstp {
    disable;
  }
}
```

### Configuring Redundant Trunk Links on Switch 3

To configure redundant trunk links on Switch 3, perform this task:

#### CLI Quick Configuration

To quickly configure the redundant trunk group example1 on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols rstp disable
set switch-options redundant-trunk-group group example1 interface ge-0/0/9.0 primary
set switch-options redundant-trunk-group group example1 interface ge-0/0/10.0
set redundant-trunk-group group example1 preempt-cutover-timer 60
```

#### Step-by-Step Procedure

Configure the redundant trunk group example1 on Switch 3.

1. Turn off RSTP:

```
[edit]
user@switch# set protocols rstp disable
```

2. Name the redundant trunk group example1 while configuring trunk interface ge-0/0/9.0 as the primary link and ge-0/0/10 as an unspecified link to serve as the secondary link:

```
[edit switch-options]
user@switch# set redundant-trunk-group group example1 interface ge-0/0/9.0 primary
user@switch# set redundant-trunk-group group example1 interface ge-0/0/10.0
```

3. (Optional) Change the time interval (from the default 120 seconds) that a re-enabled primary link waits to take over for an active secondary link:

```
[edit switch-options]
user@switch# set redundant-trunk-group group example1 preempt-cutover-timer 60
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
switch-options
  redundant-trunk-group {
    group example1 {
      preempt-cutover-timer 60;
      interface ge-0/0/9.0 {
        primary;
      }
      interface ge-0/0/10.0;
    }
  }
protocols {
  rstp {
    disable;
  }
}
```

---

### Verification

To confirm that the configuration is set up correctly, perform this task:

- [Verifying That a Redundant Trunk Group Was Created on page 2330](#)

#### *Verifying That a Redundant Trunk Group Was Created*

**Purpose** Verify that the redundant trunk group example1 has been created on Switch 1 and that trunk interfaces are members of the redundant trunk group.

**Action** List all redundant trunk groups configured on the switch:

```
user@switch> show redundant-trunk-group
```

| Group name | Interface   | State  | Time of last flap | Flap count |
|------------|-------------|--------|-------------------|------------|
| example1   | ge-0/0/9.0  | Up/Pri | Never             | 0          |
|            | ge-0/0/10.0 | Up     | Never             | 0          |

**Meaning** The `show redundant-trunk-group` command lists all redundant trunk groups configured on the switch as well as the interface names and their current states (up or down for an unspecified link, and up or down and primary for a primary link). For this configuration example, the output shows that the redundant trunk group `example1` is configured on the switch. The **Up** beside the interfaces indicates that both link cables are physically connected. The **Pri** beside trunk interface `ge-0/0/9.0` indicates that it is configured as the primary link.

**Related Documentation**

- [Understanding Redundant Trunk Links on page 2276](#)

## Example: Configuring Proxy ARP on an EX Series Switch

You can configure proxy Address Resolution Protocol (ARP) on your EX Series switch to enable the switch to respond to ARP queries for network addresses by offering its own MAC address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

This example shows how to configure proxy ARP on an access switch:

- [Requirements on page 2331](#)
- [Overview and Topology on page 2331](#)
- [Configuration on page 2332](#)
- [Verification on page 2332](#)

### Requirements

This example uses the following hardware and software components:

- One EX Series switch
- Junos OS Release 10.0 or later for EX Series switches

### Overview and Topology

This example shows the configuration of proxy ARP on an interface of an EX Series switch using restricted mode. In restricted mode, the switch does not act as a proxy for hosts on the same subnet.

The topology for this example consists of one EX Series switch. When a host wants to communicate with a host that is not already in its ARP table, it broadcasts an ARP request for the MAC address of the destination host:

- When proxy ARP is not enabled, a host that shares the same IP address replies directly to the ARP request, providing its MAC address, and future transmissions are sent directly to the destination host MAC address.
- When proxy ARP is enabled, the switch responds to ARP requests, providing the switch's MAC address—even when the destination IP address is the same as the source IP address. Thus, communications must be sent through the switch and then routed through the switch to the appropriate destination.

## Configuration

---

To configure proxy ARP, perform the following tasks:

**CLI Quick Configuration** To quickly configure proxy ARP on an interface, copy the following command and paste it into the switch terminal window:

```
[edit]
set interfaces ge-0/0/3 unit 0 proxy-arp restricted
```

**Step-by-Step Procedure** You configure proxy ARP on individual interfaces.

1. To configure proxy ARP on an interface:

```
[edit interfaces]
user@switch# set ge-0/0/3 unit 0 proxy-arp restricted
```



**BEST PRACTICE:** We recommend that you configure proxy ARP in restricted mode. In restricted mode, the switch does not act as a proxy if the source and target IP addresses are on the same subnet. If you use unrestricted mode, disable gratuitous ARP requests on the interface to avoid a situation wherein the switch's response to a gratuitous ARP request appears to the host to be an indication of an IP conflict.

```
[edit interfaces]
user@switch# set ge-0/0/3 no-gratuitous-arp-request
```

**Results** Display the results of the configuration:

```
user@switch> show configuration
interfaces {
  ge-0/0/3 {
    unit 0 {
      proxy-arp restricted;
      family ethernet-switching;
    }
  }
}
```

## Verification

---

To verify that the switch is sending proxy ARP messages, perform these tasks:

- [Verifying That the Switch Is Sending Proxy ARP Messages on page 2332](#)

### *Verifying That the Switch Is Sending Proxy ARP Messages*

**Purpose** Verify that the switch is sending proxy ARP messages.

**Action** List the system statistics for ARP messages:

```
user@switch> show system statistics arp
arp:
    90060 datagrams received
    34 ARP requests received
```

```

610 ARP replies received
2 resolution request received
0 unrestricted proxy requests
0 restricted proxy requests
0 received proxy requests
0 unrestricted proxy requests not proxied
0 restricted proxy requests not proxied
0 datagrams with bogus interface
0 datagrams with incorrect length
0 datagrams for non-IP protocol
0 datagrams with unsupported op code
0 datagrams with bad protocol address length
0 datagrams with bad hardware address length
0 datagrams with multicast source address
0 datagrams with multicast target address
0 datagrams with my own hardware address
0 datagrams for an address not on the interface
0 datagrams with a broadcast source address
294 datagrams with source address duplicate to mine
89113 datagrams which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
309 ARP requests sent
35 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor

```

**Meaning** The statistics show that two proxy ARP requests were received. The **unrestricted proxy requests not proxied** and **restricted proxy requests not proxied** fields indicate that all the unproxied ARP requests received have been proxied by the switch.

**Related Documentation**

- [Configuring Proxy ARP \(CLI Procedure\)](#)
- [Configuring Proxy ARP \(CLI Procedure\) on page 2360](#)
- [Understanding Proxy ARP on EX Series Switches on page 2278](#)

## Configuration Tasks

- [Configuring VLANs for EX Series Switches \(J-Web Procedure\) on page 2334](#)
- [Configuring VLANs for EX Series Switches \(CLI Procedure\) on page 2337](#)
- [Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\) on page 2340](#)
- [Configuring MAC Table Aging \(CLI Procedure\) on page 2341](#)
- [Configuring the Native VLAN Identifier \(CLI Procedure\) on page 2342](#)
- [Creating a Private VLAN on a Single Switch \(CLI Procedure\) on page 2343](#)
- [Creating a Private VLAN Spanning Multiple Switches \(CLI Procedure\) on page 2345](#)

- [Configuring Virtual Routing Instances \(CLI Procedure\) on page 2347](#)
- [Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\) on page 2348](#)
- [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 2351](#)
- [Configuring Layer 2 Protocol Tunneling on EX Series Switches \(CLI Procedure\) on page 2357](#)
- [Configuring Redundant Trunk Groups \(J-Web Procedure\) on page 2358](#)
- [Configuring Proxy ARP \(CLI Procedure\) on page 2360](#)
- [Adding a Static MAC Address Entry to the Ethernet Switching Table \(CLI Procedure\) on page 2361](#)
- [Configuring MAC Limiting \(CLI Procedure\) on page 2361](#)
- [Configuring MAC Notification \(CLI Procedure\) on page 2363](#)

## Configuring VLANs for EX Series Switches (J-Web Procedure)

---



**NOTE:** This topic applies only to the J-Web Application package.

---

You can use the VLAN Configuration page to add a new VLAN or to edit or delete an existing VLAN on an EX Series switch.

To access the VLAN Configuration page:

1. Select **Configure > Switching > VLAN**.

The VLAN Configuration page displays a list of existing VLANs. If you select a specific VLAN, the specific VLAN details are displayed in the Details section.

---



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

---

2. Click one of the following options:

- **Add**—Creates a VLAN.
  - **Edit**—Edits an existing VLAN configuration.
  - **Delete**—Deletes an existing VLAN.
- 



**NOTE:** If you delete a VLAN, the VLAN configuration for all the associated interfaces is also deleted.

---

When you are adding or editing a VLAN, enter information as described in [Table 240 on page 2335](#).

Table 240: VLAN Configuration Details

| Field                                                                                                        | Function                                                                                                                     | Your Action                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| General tab                                                                                                  |                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                      |
| VLAN Name                                                                                                    | Specifies a unique name for the VLAN.                                                                                        | Enter a name.                                                                                                                                                                                                                                                                                                                                                                                                        |
| VLAN ID/Range/VLAN ID/List<br><br><b>NOTE:</b> EX4300 switches support only VLAN ID/List and not VLAN Range. | Specifies the identifier or range for the VLAN.                                                                              | Select one of the following options: <ul style="list-style-type: none"> <li>• <b>VLAN ID</b>—Type a unique identification number from 1 through <b>4094</b>. If no value is specified, the ID defaults to 0.</li> <li>• <b>VLAN Range/List</b>—Type a number range to create VLANs with IDs corresponding to the numbers in the range. For example, the range 2–3 creates two VLANs with the IDs 2 and 3.</li> </ul> |
| Description                                                                                                  | Describes the VLAN.                                                                                                          | Enter a brief description for the VLAN.                                                                                                                                                                                                                                                                                                                                                                              |
| MAC-Table-Aging-Time<br><br><b>NOTE:</b> This option is not supported on EX4300 switches.                    | Specifies the maximum time that an entry can remain in the forwarding table before it <i>ages out</i> .                      | Type the number of seconds from <b>60</b> through <b>1000000</b> .                                                                                                                                                                                                                                                                                                                                                   |
| Input filter                                                                                                 | Specifies the VLAN firewall filter that is applied to incoming packets.                                                      | To apply an input firewall filter, select the firewall filter from the list.                                                                                                                                                                                                                                                                                                                                         |
| Output filter                                                                                                | Specifies the VLAN firewall filter that is applied to outgoing packets.                                                      | To apply an output firewall filter, select the firewall filter from the list.                                                                                                                                                                                                                                                                                                                                        |
| Ports tab                                                                                                    |                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Ports<br><br><b>NOTE:</b> This option is not supported on EX4300 switches.                                   | Specifies the ports (interfaces) to be associated with this VLAN for data traffic. You can also remove the port association. | Click one of the following options: <ul style="list-style-type: none"> <li>• <b>Add</b>—Select the ports from the available list. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.</li> <li>• <b>Remove</b>—Select the port that you do not want associated with the VLAN.</li> </ul>                                                                           |
| IP address tab                                                                                               |                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                      |

Table 240: VLAN Configuration Details (*continued*)

| Field        | Function                                                                                                         | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------|------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPv4 address | Specifies IPv4 address options for the VLAN.                                                                     | <p>Select <b>IPv4 address</b> to enable the IPv4 address options.</p> <p>To configure IPv4:</p> <ol style="list-style-type: none"> <li>1. Enter the IP address.</li> <li>2. Enter the subnet mask—for example, <b>255.255.255.0</b>. You can also specify the address prefix.</li> <li>3. To apply an input firewall filter to an interface, select the firewall filter from the list.</li> <li>4. To apply an output firewall filter to an interface, select the firewall filter from the list.</li> <li>5. Click the <b>ARP/MAC Details</b> button. Enter the static IP address and MAC address in the window that is displayed.</li> </ol> <p><b>NOTE:</b> In EX4300 switches, you also need to select <b>L2 Interface</b> in the window that is displayed.</p> |
| IPv6 address | Specifies IPv6 address options for the VLAN.                                                                     | <p>Select <b>IPv6 address</b> to enable the IPv6 address options.</p> <p>To configure IPv6:</p> <ol style="list-style-type: none"> <li>1. Enter the IP address—for example: <b>2001:ab8:85a3::8a2e:370:7334</b>.</li> <li>2. Specify the subnet mask.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Voip tab     |                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Ports        | Specifies the ports to be associated with this VLAN for voice traffic. You can also remove the port association. | <p>Click one of the following options:</p> <ul style="list-style-type: none"> <li>• <b>Add</b>—Select the ports from the list of available ports. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.</li> <li>• <b>Remove</b>—Select the port that you do not want associated with the VLAN.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                         |

**Related Documentation**

- [Configuring VLANs for EX Series Switches \(CLI Procedure\) on page 2337](#)
- [Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 2281](#)
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch](#)
- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)
- [Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\) on page 2340](#)
- [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#)



## Configuring VLANs for EX Series Switches (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring VLANs for EX Series Switches (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

EX Series switches use VLANs to make logical groupings of network nodes with their own broadcast domains. VLANs limit the traffic flowing across the entire LAN and reduce collisions and packet retransmissions.

- [Why Create a VLAN? on page 2337](#)
- [Creating a VLAN Using the Minimum Procedure on page 2337](#)
- [Creating a VLAN Using All of the Options on page 2338](#)
- [Configuration Guidelines for VLANs on page 2339](#)

### Why Create a VLAN?

For switching to begin, you must explicitly configure at least one VLAN, even if your network is simple and you want only one broadcast domain to exist.

Some reasons to create more than one VLAN are:

- A LAN has more than 200 devices.
- A LAN has a large amount of broadcast traffic.
- A group of clients requires that a higher-than-average level of security be applied to traffic entering or exiting the group's devices.
- A group of clients requires that the group's devices receive less broadcast traffic than they are currently receiving, so that data speed across the group is increased.

### Creating a VLAN Using the Minimum Procedure

These steps are required to create a VLAN:

- Uniquely identify the VLAN. You do this by assigning a name and an ID to the VLAN.
- Assign at least one switch port interface to the VLAN for communication. After assigning one or more interfaces to the VLAN, the interfaces function in access mode. All interfaces in a single VLAN are in a single broadcast domain, even if the interfaces are on different switches. You can assign traffic on any switch to a particular VLAN by referencing either the interface sending traffic or the MAC addresses of devices sending traffic.

The following example creates a VLAN using only a few required steps. The VLAN is created with the name **employee-vlan** and the VLAN ID of **100**. Then, three interfaces are

assigned to that VLAN, and these interfaces, which function in access mode, transmit traffic among themselves.

```
[edit] set vlans employee-vlan
edit] set vlans employee-vlan vlan-id 100
[edit] set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members
employee-vlan
[edit] set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members
employee-vlan
[edit] set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members
employee-vlan
```

In the example, all users connected to the interfaces ge-0/0/1, ge-0/0/2, and ge-0/0/3 can communicate with each other, but not with users on other interfaces in this network. To configure communication between VLANs, you must configure an integrated routing and bridging (IRB) interface. See [“Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\)” on page 2340](#).

---

### Creating a VLAN Using All of the Options

---

To configure a VLAN, follow these steps:

1. Create the VLAN by setting the unique VLAN name:

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

2. Configure the VLAN ID or a VLAN ID list for the VLAN. Using the VLAN ID list option, you can optionally specify a range of VLAN IDs.

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

```
[edit vlans]
user@switch# set vlan-name vlan-id-list [vlan-ids | vlan-id--vlan-id-]
```

3. Assign at least one interface to the VLAN:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family ethernet-switching vlan
members [all | vlan-names | vlan-ids]
```



**NOTE:** You can also specify that a trunk interface is a member of all VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.

4. (Optional) Create a subnet for the VLAN because all computers that belong to a subnet are addressed with a common, identical, most-significant-bit group in their IP address. This makes it easy to identify VLAN members by their IP addresses. To create the subnet for the VLAN:

```
[edit interfaces]
user@switch# set vlan unit logical-unit-number family inet address
ip-address/destination-prefix
```

5. (Optional) Specify the description of the VLAN:

```
[edit vlans]
```

```
user@switch# set vlan-name description text-description
```

6. (Optional) For security purposes, specify a VLAN firewall filter to be applied to incoming or outgoing packets:

```
[edit vlans]
```

```
user@switch# set vlan-name filter (input | output) filter-name
```

### Configuration Guidelines for VLANs

To create a VLAN, you must uniquely identify the VLAN and assign at least one switch port interface to the VLAN for communication. After you assign one or more interfaces to the VLAN, the interfaces function in access mode.

After creating a VLAN, all users connected to interfaces that are assigned to the VLAN can communicate with each other but not with users on other interfaces in the network. To configure communication between VLANs, you must configure an IRB interface. For information about creating an IRB interface, see [“Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\)” on page 2340..](#)

The number of VLANs supported per switch varies. Use the command **set vlans *vlan-name* *vlan-id* ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created.

#### Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)
- [Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches on page 2254](#)
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 2281](#)

## Configuring Integrated Routing and Bridging Interfaces (CLI Procedure)

Integrated routing and bridging (IRB) interfaces allow an EX Series switch to recognize packets that are being sent to local addresses so that they are bridged (switched) whenever possible and are routed only when necessary. Whenever packets can be switched instead of routed, several layers of processing are eliminated.

An interface named `irb` functions as a logical router on which you can configure a Layer 3 logical interface for each virtual LAN (VLAN). For redundancy, you can combine an IRB interface with implementations of the Virtual Router Redundancy Protocol (VRRP) in both bridging and virtual private LAN service (VPLS) environments.

Jumbo frames of up to 9216 bytes are supported on an IRB interface. To route jumbo data packets on the IRB interface, you must configure the jumbo MTU size on the member physical interfaces of the VLAN that you have associated with the IRB interface, as well as on the IRB interface itself (the interface named `irb`).



**CAUTION:** Setting or deleting the jumbo MTU size on the IRB interface (the interface named `irb`) while the switch is transmitting packets might result in dropped packets.

To configure the IRB interface:

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

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

2. Assign an interface to the VLAN by naming the VLAN as a trunk member on the logical interface, thereby making the interface part of the VLAN's broadcast domain:

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

3. Create a logical Layer 3 IRB interface (its name will be `irb.logical-interface-number`, where the value for *logical-interface-number* is the value you supplied for *vlan-id* in Step 1; in the following command, it is the *logical-unit-number*) on a subnet for the VLAN's broadcast domain:

```
[edit]
user@switch# set interfaces irb unit logical-unit-number family inet address inet-address
```

4. Link the Layer 2 VLAN to the logical Layer 3 IRB interface:

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



**NOTE:** Layer 3 interfaces on trunk ports allow the interface to transfer traffic between multiple Layer 2 VLANs. Within a VLAN, traffic is switched, while across VLANs, traffic is routed.

- Related Documentation**
- [Verifying Integrated Routing and Bridging Interface Status and Statistics on page 2436](#)
  - [Understanding Integrated Routing and Bridging Interfaces and Routed VLAN Interfaces on EX Series Switches on page 2254](#)

## Configuring MAC Table Aging (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring MAC Table Aging (CLI Procedure)*.

The Ethernet switching table (or MAC table) aging process ensures that the EX Series switch tracks only active MAC addresses on the network and is able to flush out MAC addresses that are no longer used.

You can configure the MAC table aging time, the maximum time that an entry can remain in the Ethernet Switching table before it *ages out*, on all VLANs on the switch. This setting can influence efficiency of network resource use by affecting the amount of traffic that is flooded to all interfaces because when traffic is received for MAC addresses no longer in the Ethernet switching table, the switch floods the traffic to all interfaces.

To configure the MAC table aging time on all VLANs on the switch:

[edit]

```
user@switch# set protocols l2-learning global-mac-table-aging-time seconds
```

- Related Documentation**
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 2281](#)
  - [Example: Connecting Access Switches to a Distribution Switch on page 2291](#)
  - [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)

## Configuring the Native VLAN Identifier (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches and Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring the Native VLAN Identifier (CLI Procedure)*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 3.

Switches can receive and forward routed or bridged Ethernet frames with 802.1Q VLAN tags. Typically, trunk ports, which connect switches to each other, accept untagged control packets but do not accept untagged data packets. You can enable a trunk port to accept untagged data packets by configuring a native VLAN ID on the interface on which you want the untagged data packets to be received. The logical interface on which untagged packets are to be received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface.

To configure the native VLAN ID by using the command-line interface (CLI):

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

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

2. Configure the native VLAN ID:

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

3. Specify that the logical interface that will receive the untagged data packets is a member of the native VLAN:

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

### Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)
- [Example: Connecting Access Switches to a Distribution Switch on page 2291](#)
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 2281](#)
- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series](#)

## Creating a Private VLAN on a Single Switch (CLI Procedure)



**NOTE:** This task uses Junos OS for switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see *Creating a Private VLAN on a Single EX Series Switch (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic or limit the communication between known hosts. Private VLANs (PVLANS) enable you to split a broadcast domain (primary VLAN) into multiple isolated broadcast subdomains (secondary VLANs), essentially putting a VLAN inside a VLAN. This procedure describes how to create a PVLAN on a single switch.



**NOTE:** You must specify a VLAN ID for each secondary VLAN even if the PVLAN is configured on a single switch.

You do not need to preconfigure the primary VLAN. This topic shows the primary VLAN being configured as part of this PVLAN configuration procedure.

For a list of guidelines on configuring PVLANS, see “[Understanding Private VLANs](#)” on page 2258.

To configure a private VLAN on a single switch:

1. Set the VLAN ID for the primary VLAN:

[edit vlans]

```
user@switch# set primary-vlan-name vlan-id vlan-id-number
```

2. Configure at least one interface within the primary VLAN so that it communicates with all the subdomains of the PVLAN. This interface functions as a *promiscuous* port. It can be either a trunk port or an access port.

[edit interfaces]

```
user@switch# set interface-name unit 0 family ethernet-switching
```

```
user@switch# set interface-name unit 0 family ethernet-switching vlan members  
primary-vlan-name
```

3. Configure another promiscuous interface of the primary VLAN as a trunk port to connect the PVLAN to the external router or switch:

[edit interfaces]

```
user@switch# set interface-name unit 0 family ethernet-switching interface-mode trunk
```

```
user@switch# set interface-name unit 0 family ethernet-switching vlan members  
primary-vlan-name
```

4. Create an isolated VLAN by selecting the **isolated** option for **private-vlan**, and setting a VLAN ID for the isolated VLAN:

[edit vlans]

```
user@switch# set isolated-vlan-name private-vlan isolated vlan-id isolated-vlan-id
```



**NOTE:** You can create only one isolated VLAN within a private VLAN. Setting the VLAN name for the isolated VLAN is optional. Configuring the VLAN ID is required.

5. Create a community VLAN by selecting the **community** option for **private-vlan**, and setting a VLAN ID for this community VLAN:

[edit vlans]

```
user@switch# set community-vlan-name private-vlan community vlan-id community-vlan-id
```



**NOTE:** To create additional community VLANs, repeat this step and specify a different name for the community VLAN. Setting the VLAN name for the community VLAN is optional. Configuring the VLAN ID is required.

6. Associate the isolated VLAN with the primary VLAN:

[edit vlans]

```
user@switch# set primary-vlan-name vlan-id primary-vlan-id isolated-vlan isolated-vlan-name
```

7. Associate each community VLAN with the primary VLAN:

[edit vlans]

```
user@switch# set primary-vlan-name vlan-id primary-vlan-id  
community-vlan community-vlan-name
```

8. If you have not already done so, configure at least one interface of the isolated VLAN.

[edit interfaces]

```
user@switch# set interface-name unit logical-unit-number family ethernet-switching  
interface-mode access vlan members isolated-vlan-name
```

9. If you have not already done so, configure at least one interface of the community VLAN.

[edit interfaces]

```
user@switch# set interface-name unit logical-unit-number family ethernet-switching  
interface-mode access vlan members community-vlan-name
```



**NOTE:** Repeat the same step on other community VLANs that you want to include in the PVLAN.

#### Related Documentation

- [Understanding Private VLANs on page 2258](#)
- [Creating a Private VLAN Spanning Multiple Switches \(CLI Procedure\) on page 2345](#)



## Creating a Private VLAN Spanning Multiple Switches (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic or limit the communication between known hosts. Private VLANs (PVLANS) enable you to split a broadcast domain (primary VLAN) into multiple isolated broadcast subdomains (secondary VLANs), essentially putting a VLAN inside a VLAN. This procedure describes how to configure a PVLAN to span multiple switches.

For a list of guidelines on configuring PVLANS, see “[Understanding Private VLANs](#)” on page 2258.

To configure a PVLAN to span multiple switches, perform the following procedure on all the switches that will participate in the PVLAN::

1. Create the primary VLAN by setting the unique VLAN name and specify an 802.1Q tag for the VLAN:

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

2. On the switch that will connect to a router, configure a promiscuous interface as a trunk port to connect the PVLAN to the router:

```
[edit interfaces]
user@switch# set interface-name unit 0 family ethernet-switching interface-mode trunk
user@switch# set interface-name unit 0 family ethernet-switching vlan members
primary-vlan-name
```

3. On all the switches, configure a trunk interface as the Inter-Switch Link (ISL) that will be used to connect the switches to each other:

```
[edit interfaces]
user@switch# set interface-name unit 0 family ethernet-switching interface-mode trunk
inter-switch-link
user@switch# set interface-name unit 0 family ethernet-switching vlan members
name-of-private-vlan
```

4. Create an isolated VLAN within the primary VLAN by selecting the **isolated** option for **private-vlan**, and setting a VLAN ID for the isolated VLAN:

```
[edit vlans]
user@switch# set isolated-vlan-name private-vlan isolated-vlan-id isolated-vlan-id
```



**NOTE:** You can create only one isolated VLAN within a private VLAN. The isolated VLAN can contain member interfaces from the multiple switches that compose the PVLAN.

Setting the VLAN name for the isolated VLAN is optional. Configuring the VLAN ID is required.

5. Create a community VLAN within the primary VLAN by selecting the **community** option for **private-vlan**, and setting a VLAN ID for this community VLAN::

[edit vlans]

```
user@switch# set community-vlan-name private-vlan community vlan-id community-vlan-id
```

---



**NOTE:** To create additional community VLANs, repeat this step and specify a different name for the community VLAN. Setting the VLAN name for the community VLAN is optional. Configuring the VLAN ID is required.

---

6. Associate the isolated VLAN with the primary VLAN:

[edit vlans *primary-vlan-name* vlan-id *primary-vlan-id*]

```
user@switch# set isolated-vlan isolated-vlan-name
```

7. Associate each community VLAN with the primary VLAN:

[edit vlans *primary-vlan-name* vlan-id *primary-vlan-id*]

```
user@switch# set community-vlan community-vlan-name
```

8. If you have not already done so, configure at least one access interface to be a member of the isolated VLAN.

[edit interface]

```
user@switch# set interface-name unit logical-unit-number family ethernet-switching  
interface-mode access vlan members isolated-vlan-name
```

9. If you have not already done so, configure at least one access interface to be a member of the community VLAN.

[edit interface]

```
user@switch# set interface-name unit logical-unit-number family ethernet-switching  
interface-mode access vlan members community-vlan-name
```

---



**NOTE:** Repeat this step for the other community VLANs that you are including in the PVLAN.

---

#### Related Documentation

- [Understanding Private VLANs on page 2258](#)
- [Example: Configuring a Private VLAN on a Single Switch on page 2302](#)
- [Creating a Private VLAN on a Single Switch \(CLI Procedure\) on page 2343](#)

## Configuring Virtual Routing Instances (CLI Procedure)

Use virtual routing and forwarding (VRF) to divide an EX Series switch into multiple virtual routing instances. VRF allows you to isolate traffic traversing the network without using multiple devices to segment your network. VRF is supported on all Layer 3 interfaces.

Before you begin, make sure to set up your VLANs. See *Configuring VLANs for EX Series Switches (CLI Procedure)*, “[Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)” on page 2337, or “[Configuring VLANs for EX Series Switches \(J-Web Procedure\)](#)” on page 2334.

To configure virtual routing instances:

1. Create a routing instance:

```
[edit routing-instances]user@switch# set routing-instance-name instance-type virtual-router
```



**NOTE:** EX Series switches only support the virtual-router instance type.

2. Bind each routing instance to the corresponding physical interfaces:

```
[edit routing-instances]user@switch# set routing-instance-name interface
interface-name.logical-unit-number
```

3. Create the logical interfaces that are bound to the routing instance.

- To create a logical interface with an IPv4 address:

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

- To create a logical interface with an IPv6 address:

```
[edit interfaces]user@switch# set interface-name unit logical-unit-number family inet6
address ipv6-address
```



**NOTE:** Do not create a logical interface using the family ethernet-switching option in this step. Binding an interface using the family ethernet-switching option to a routing instance can cause the interface to shutdown.

4. Enable VLAN tagging on each physical interface that was bound to the routing instance:

```
[edit interfaces]user@switch# set interface-name vlan-tagging
```

### Related Documentation

- [Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 2306](#)
- [Verifying That Virtual Routing Instances Are Working on page 2435](#)
- [Understanding Virtual Routing Instances on EX Series Switches on page 2263](#)

## Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)

---



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs Junos OS that does not support ELS, see *Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

---

Multiple VLAN Registration Protocol (MVRP) is used to manage dynamic VLAN registration in a LAN. You can use MVRP on EX Series switches.

MVRP is disabled by default on EX Series switches.

To enable MVRP or set MVRP options, follow these instructions:

- [Enabling MVRP on page 2348](#)
- [Disabling MVRP on page 2348](#)
- [Disabling Dynamic VLANs on page 2348](#)
- [Configuring Timer Values on page 2349](#)
- [Configuring MVRP Registration Mode on page 2350](#)
- [Using MVRP in a Mixed-Release Network on page 2350](#)

---

### Enabling MVRP

MVRP can only be enabled on trunk interfaces.

To enable MVRP on a trunk interface:

```
[edit protocols mvrp]  
user@switch# set interface interface-name
```

---

### Disabling MVRP

MVRP is disabled by default. You only need to perform this procedure if you have previously enabled MVRP.

You can disable MVRP globally only. . To disable MVRP on all trunk interfaces on a switch, use one of the following commands:

```
user@switch# deactivate protocols mvrp  
user@switch# delete protocols mvrp
```

---

### Disabling Dynamic VLANs

By default, dynamic VLANs can be created on interfaces participating in MVRP. Dynamic VLANs are VLANs created on one switch that are propagated to other switches dynamically; in this case, using MVRP.

Dynamic VLAN creation through MVRP cannot be disabled per switch interface. To disable dynamic VLAN creation for interfaces participating in MVRP, you must disable it for all interfaces on the switch.

To disable dynamic VLAN creation:

```
[edit protocols mvrp]
user@switch# set no-dynamic-vlan
```

### Configuring Timer Values

The timers in MVRP define the amount of time all interfaces on a switch or a specific interface wait to join or leave MVRP, or to send or process the MVRP information for the switch after receiving an MVRP PDU. The join timer controls the amount of time the switch waits to accept a registration request, the leave timer controls the period of time that the switch waits in the Leave state before changing to the unregistered state, and the leaveall timer controls the frequency with which the LeaveAll messages are communicated.

The default MVRP timer values are 200 ms for the join timer, 1000 ms for the leave timer, and 10 seconds for the leaveall timer.



**BEST PRACTICE:** Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

On an EX Series switch that uses Junos OS with support for ELS, if the timer value set on an interface level is different from the value set on a switch level, then the value on the interface level takes precedence.

To set the join timer for all interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set join-timer milliseconds
```

To set the join timer for a specific interface:

```
[edit protocols mvrp]
user@switch# set interface interface-name join-timer milliseconds
```

To set the leave timer for all interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set leave-timer milliseconds
```

To set the leave timer for a specific interface:

```
[edit protocols mvrp]
user@switch# set interface interface-name leave-timer milliseconds
```

To set the leaveall timer for all interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set leaveall-timer seconds
```

To set the leaveall timer for a specific interface:

```
[edit protocols mvrp]
user@switch# set interface interface-name leaveall-timer seconds
```

### Configuring MVRP Registration Mode

---

The default MVRP registration mode for any interface participating in MVRP is normal. An interface in normal registration mode participates in MVRP when MVRP is enabled on the switch.

You can change the registration mode of a specific interface to **forbidden**. An interface in forbidden registration mode does not participate in MVRP even if MVRP is enabled on the switch.

To set an interface to forbidden registration mode:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 registration forbidden
```

To set an interface to normal registration mode:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 registration normal
```

### Using MVRP in a Mixed-Release Network

---

Except in Junos OS Releases 11.2 and earlier, MVRP has conformed with IEEE standard 802.1ak and IEEE Draft 802.1Q regarding the inclusion of an extra byte in the protocol data units (PDUs) sent and received by MVRP. As a result of changes in the standards with regard to the extra byte, MVRP in Junos OS Releases 13.2 and later for EX Series switches with support for the Enhanced Layer 2 Software (ELS) includes the extra byte, while MVRP in Junos OS Releases 11.3 and later for EX Series switches that do not support ELS does not include the extra byte. A compatibility issue arises, wherein the ELS version of MVRP does not recognize PDUs without the extra byte sent by the non-ELS version of MVRP.

For more information about this issue, see [“Understanding Multiple VLAN Registration Protocol \(MVRP\) on EX Series Switches” on page 2264](#).

You can recognize an MVRP version compatibility issue by observing the switch running the ELS version of MVRP. Because a switch running the ELS version of MVRP cannot interpret an unmodified PDU from a switch running the non-ELS version of MVRP, the switch will not add VLANs from the non-ELS version of MVRP. When you use the **show mvrp statistics** command in the ELS version of MVRP, the values for **Received Join Empty** and **Received Join In** will incorrectly display zero, even though the value for the **Received MVRP PDUs without error** has been increased. Another indication that MVRP is having a version compatibility issue is that unexpected VLAN activity, such as multiple VLAN creation, takes place on the switch running the ELS version of MVRP.

If your network includes a mix of EX Series switches running ELS and non-ELS versions of MVRP, you can eliminate the compatibility issue by entering the following command on the switches running the ELS version of MVRP:

```
[edit protocols mvrp]
user@switch# set no-attribute-length-in-pdu
```

The `no-attribute-length-in-pdu` statement prevents the ELS version of MVRP from sending PDUs with the extra byte, thereby eliminating the compatibility issue with the non-ELS version of MVRP.

#### Related Documentation

- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 2310](#)
- [Verifying That MVRP Is Working Correctly on page 2438](#)
- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on EX Series Switches on page 2264](#)

## Configuring Q-in-Q Tunneling (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Q-in-Q Tunneling (CLI Procedure)*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

Q-in-Q tunneling enables service providers on Ethernet access networks to segregate or bundle customer traffic into different VLANs by adding another layer of 802.1Q tags. You can configure Q-in-Q tunneling on EX Series switches.



**NOTE:** You cannot configure 802.1X user authentication on interfaces that have been enabled for Q-in-Q tunneling.

When Q-in-Q tunneling is configured on EX Series switches, trunk interfaces are assumed to be part of the service-provider network and access interfaces are assumed to be part of the customer network. Therefore, this topic also refers to trunk interfaces as service-provider VLAN (S-VLAN) interfaces (network-to-network interfaces [NNI]), and to access interfaces as customer VLAN (C-VLAN) interfaces (user-network interfaces [UNI]).

Before you begin configuring Q-in-Q tunneling, make sure you set up your VLANs. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 2337](#) or [“Configuring VLANs for EX Series Switches \(J-Web Procedure\)” on page 2334](#).

Configure Q-in-Q tunneling by using one of the following methods to map C-VLANs to S-VLANs:

- [Configuring All-in-One Bundling on page 2352](#)
- [Configuring Many-to-Many Bundling on page 2353](#)
- [Configuring a Specific Interface Mapping with VLAN Rewrite Option on page 2355](#)

## Configuring All-in-One Bundling

You can configure Q-in-Q tunneling by using the all-in-one bundling method, which maps packets from all C-VLAN interfaces on a switch to an S-VLAN.

To configure the all-in-one bundling method on a C-VLAN interface:

1. Enable the transmission of packets with no or a single 802.1Q VLAN tag:

```
[edit interfaces interface-name]  
user@switch# set flexible-vlan-tagging
```

2. Enable extended VLAN bridge encapsulation:

```
[edit interfaces interface-name]  
user@switch# set encapsulation extended-vlan-bridge
```

3. Map packets from all C-VLANs to a logical interface:

```
[edit interfaces interface-name unit logical-unit-number]  
user@switch# set vlan-id-list vlan-id-numbers
```

4. Enable a C-VLAN interface to send and receive untagged packets:

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

When specifying a native VLAN ID on a C-VLAN physical interface, the value must be included in the VLAN ID list specified on the C-VLAN logical interface in step 3.

5. Specify that packets traveling from a C-VLAN interface to an S-VLAN interface are tagged with the VLAN ID of the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]  
user@switch# set input-vlan-map push
```

6. Specify that the 802.1Q S-VLAN tag is removed as packets exit an S-VLAN interface.

```
[edit interfaces interface-name unit logical-unit-number]  
user@switch# set output-vlan-map pop
```

7. Configure a name for the S-VLAN, and associate the logical interface configured in step 3 with the S-VLAN:

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

The following configuration on the C-VLAN interface ge-0/0/1 enables Q-in-Q tunneling and maps packets from C-VLANs 100 through 200 to logical interface 10, which is in turn associated with S-VLAN v10. In this sample configuration, a packet originated in C-VLAN 100 includes a tag with the VLAN ID 100. When this packet travels from the interface ge-0/0/1 to the S-VLAN interface, a tag with VLAN ID 10 is added to it. As the packet exits the S-VLAN interface, the tag with the VLAN ID 10 is removed. .

```
set interfaces ge-0/0/1 flexible-vlan-tagging  
set interfaces ge-0/0/1 encapsulation extended-vlan-bridge  
set interfaces ge-0/0/1 unit 10 vlan-id-list 100-200  
set interfaces ge-0/0/1 native-vlan-id 150  
set interfaces ge-0/0/1 unit 10 input-vlan-map push  
set interfaces ge-0/0/1 unit 10 output-vlan-map pop  
set vlans v10 interface ge-0/0/1.10
```



To configure the all-in-one bundling method on an S-VLAN interface:

1. Enable the transmission of packets with no, one, or two 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# set flexible-vlan-tagging
```

2. Enable extended VLAN bridge encapsulation:

```
[edit interfaces interface-name]
user@switch# set encapsulation extended-vlan-bridge
```

3. Map packets from the logical interface specified in the C-VLAN interface configuration to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# set vlan-id number
```

4. Enable the S-VLAN interface to send and receive untagged packets:

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

When specifying a native VLAN ID on an S-VLAN physical interface, the value must match the VLAN ID specified on the S-VLAN logical interface in step 3.

5. Associate the S-VLAN interface with the S-VLAN that was configured in the C-VLAN interface procedure:

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

For example, the following configuration on the S-VLAN interface ge-1/1/1 enables Q-in-Q tunneling and maps packets with a VLAN ID tag of 10 to logical interface 10, which is in turn associated with S-VLAN v10. .

```
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation extended-vlan-bridge
set interfaces ge-1/1/1 unit 10 vlan-id 10
set interfaces ge-1/1/1 native-vlan-id 10
set vlans v10 interface ge-1/1/1.10
```

### Configuring Many-to-Many Bundling

You can configure Q-in-Q tunneling by using the many-to-many bundling method, which maps packets from multiple C-VLANs to multiple S-VLANs.

To configure the many-to-many bundling method on a C-VLAN interface:

1. Enable the transmission of packets with no or a single 802.1Q VLAN tag:

```
[edit interfaces interface-name]
user@switch# set flexible-vlan-tagging
```

2. Enable extended VLAN bridge encapsulation:

```
[edit interfaces interface-name]
user@switch# set encapsulation extended-vlan-bridge
```

3. Map packets from specified C-VLANs to a logical interface:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# set vlan-id-list vlan-id-numbers
```

4. Enable a C-VLAN interface to send and receive untagged packets:

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

When specifying a native VLAN ID on a C-VLAN physical interface, the value must be included in the VLAN ID list specified on the C-VLAN logical interface in step 3.

5. Specify that packets traveling from a C-VLAN interface to an S-VLAN interface are tagged with the VLAN ID of the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# set input-vlan-map push
```

6. Specify that the 802.1Q S-VLAN tag is removed as packets exit an S-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# set output-vlan-map pop
```

7. Configure a name for an S-VLAN, and associate the logical interface configured in step 3 with the S-VLAN:

```
[edit vlans vlan-name]
```

```
user@switch# set interface interface-name.logical-unit-number
```

The following configuration on the C-VLAN interface ge-0/0/1 for customer 1 enables Q-in-Q tunneling and maps packets from C-VLANs 100 through 120 to logical interface 10, which is in turn associated with S-VLAN v10.

The configuration on the C-VLAN interface ge-0/0/2 for customer 2 enables Q-in-Q tunneling and maps packets from C-VLANs 30 through 40, 50 through 60, and 70 through 80 to logical interface 30, which is in turn associated with S-VLAN v30.

In this sample configuration, a packet originated in C-VLAN 100 includes a tag with the VLAN ID 100. When this packet travels from the interface ge-0/0/1 to the S-VLAN interface, a tag with a VLAN ID of 10 is added to it. As the packet exits the S-VLAN interface, the tag with the VLAN ID of 10 is removed.

Customer 1

```
set interfaces ge-0/0/1 flexible-vlan-tagging
set interfaces ge-0/0/1 encapsulation extended-vlan-bridge
set interfaces ge-0/0/1 unit 10 vlan-id-list 100-120
set interfaces ge-0/0/1 native-vlan-id 100
set interfaces ge-0/0/1 unit 10 input-vlan-map push
set interfaces ge-0/0/1 unit 10 output-vlan-map pop
set vlans v10 interface ge-0/0/1.10
```

Customer 2

```
set interfaces ge-0/0/2 flexible-vlan-tagging
set interfaces ge-0/0/2 encapsulation extended-vlan-bridge
set interfaces ge-0/0/2 unit 30 vlan-id-list 30-40
set interfaces ge-0/0/2 unit 30 vlan-id-list 50-60
set interfaces ge-0/0/2 unit 30 vlan-id-list 70-80
set interfaces ge-0/0/2 native-vlan-id 30
set interfaces ge-0/0/2 unit 30 input-vlan-map push
set interfaces ge-0/0/2 unit 30 output-vlan-map pop
set vlans v30 interface ge-0/0/2.30
```

To configure the many-to-many bundling method on an S-VLAN interface:

1. Enable the transmission of packets with no, one, or two 802.1Q VLAN tags:

```
[edit interfaces interface-name]
```

```
user@switch# set flexible-vlan-tagging
```

2. Enable extended VLAN bridge encapsulation:

```
[edit interfaces interface-name]
```

```
user@switch# set encapsulation extended-vlan-bridge
```

3. Map packets from each logical interface specified in the C-VLAN interface configuration to an S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# set vlan-id number
```

4. Enable an S-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
```

```
user@switch# set native-vlan-id vlan-id
```

When specifying a native VLAN ID on an S-VLAN physical interface, the value must match an S-VLAN ID specified on the S-VLAN logical interface in step 3.

5. Associate the S-VLAN interface with the S-VLANs that were configured in the C-VLAN interface procedure:

```
[edit vlans vlan-name]
```

```
user@switch# set interface interface-name.logical-unit-number
```

For example, the following configuration on the S-VLAN interface ge-1/1/1 enables Q-in-Q tunneling and maps incoming C-VLAN packets to logical interfaces 10 and 30, which are in turn associated with S-VLANs v10 and v30, respectively.

```
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation extended-vlan-bridge
set interfaces ge-1/1/1 unit 10 vlan-id 10
set interfaces ge-1/1/1 unit 30 vlan-id 30
set interfaces ge-1/1/1 native-vlan-id 10
set vlans v10 interface ge-1/1/1.10
set vlans v30 interface ge-1/1/1.30
```

### Configuring a Specific Interface Mapping with VLAN Rewrite Option

You can configure Q-in-Q tunneling by mapping packets from a specified C-VLAN to a specified S-VLAN. In addition, while the packets are transmitted to and from the S-VLAN, you can specify that the 802.1Q C-VLAN tag be removed and replaced with the S-VLAN tag or vice versa.

To configure a specific interface mapping with VLAN rewriting on the C-VLAN interface:

1. Enable the transmission of packets with no or one 802.1Q VLAN tag:

```
[edit interfaces interface-name]
```

```
user@switch# set flexible-vlan-tagging
```

2. Enable extended VLAN bridge encapsulation:

```
[edit interfaces interface-name]
```

```
user@switch# set encapsulation extended-vlan-bridge
```

3. Map packets from a specified C-VLAN to a logical interface:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# set vlan-id number
```

4. Enable the C-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
```

```
user@switch# set native-vlan-id vlan-id
```

When specifying a native VLAN ID on a C-VLAN physical interface, the value must match the VLAN ID specified on the C-VLAN logical interface in step 3.

5. Specify that the existing 802.1Q C-VLAN tag is removed from packets traveling from a C-VLAN interface to an S-VLAN interface and replaced with the 802.1Q S-VLAN tag:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# set input-vlan-map swap
```

6. Specify that the existing 802.1Q S-VLAN tag is removed from packets traveling from an S-VLAN interface to a C-VLAN interface and replaced with the 802.1Q C-VLAN tag:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# set output-vlan-map swap
```

7. Configure a name for the S-VLAN, and associate the logical interface configured in step 3 with the S-VLAN:

```
[edit vlans vlan-name]
```

```
user@switch# set interface interface-name.logical-unit-number
```

For example, the following configuration on the C-VLAN interface ge-0/0/1 enables Q-in-Q tunneling and maps incoming packets from C-VLAN 150 to logical interface 200, which is in turn associated with VLAN v200. Also, as packets travel from the C-VLAN interface ge-0/0/1 to an S-VLAN interface, the C-VLAN tag 150 is removed and replaced with the S-VLAN tag 200. As packets travel from an S-VLAN interface to C-VLAN interface ge-0/0/1, the S-VLAN tag 200 is removed and replaced with the C-VLAN tag of 150.

```
set interfaces ge-0/0/1 flexible-vlan-tagging
set interfaces ge-0/0/1 encapsulation extended-vlan-bridge
set interfaces ge-0/0/1 unit 200 vlan-id 150
set interfaces ge-0/0/1 native-vlan-id 150
set interfaces ge-0/0/1 unit 200 input-vlan-map swap
set interfaces ge-0/0/1 unit 200 output-vlan-map swap
set vlans v200 interface ge-0/0/1.200
```

To configure a specific interface mapping with VLAN rewriting on the S-VLAN interface:

1. Enable the transmission of packets with no, one, or two 802.1Q VLAN tags:

```
[edit interfaces interface-name]
```

```
user@switch# set flexible-vlan-tagging
```

2. Enable extended VLAN bridge encapsulation:

```
[edit interfaces interface-name]
```

```
user@switch# set encapsulation extended-vlan-bridge
```

3. Map packets from the logical interface specified in the C-VLAN interface configuration to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# set vlan-id number
```

4. Enable the S-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
```

```
user@switch# set native-vlan-id vlan-id
```

When specifying a native VLAN ID on an S-VLAN physical interface, the value must match the VLAN ID specified on the S-VLAN logical interface in step 3.

5. Associate the S-VLAN interface with the S-VLAN that was configured in the C-VLAN interface procedure: :

```
[edit vlans vlan-name]
```

```
user@switch# set interface interface-name.logical-unit-number
```

For example, the following configuration on the S-VLAN interface ge-1/1/1 enables Q-in-Q tunneling and maps packets with VLAN ID 200 to logical interface 200, which is in turn associated with S-VLAN v200.

```
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation extended-vlan-bridge
set interfaces ge-1/1/1 unit 200 vlan-id 200
set interfaces ge-1/1/1 native-vlan-id 200
set vlans v200 interface ge-1/1/1.200
```

**Related  
Documentation**

- [Understanding Q-in-Q Tunneling on EX Series Switches on page 2269](#)

## Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)



**NOTE:** This topic applies to Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

Layer 2 protocol tunneling (L2PT) enables you to send Layer 2 protocol data units (PDUs) across a service provider network and deliver them to EX Series switches at a remote location. This feature is useful when you have a network that includes remote sites that are connected across a service provider network and you want to run Layer 2 protocols on switches connected across the service provider network.



**NOTE:** The implementation of L2PT on EX4300 switches does not support the drop-threshold or the shutdown-threshold configuration statements.

To configure L2PT on an EX Series switch, you must first configure a Q-in-Q interface or group of interfaces. See “[Configuring Q-in-Q Tunneling \(CLI Procedure\)](#)” on page 2351 for information about how to configure Q-in-Q tunneling.

- To configure L2PT on a specified Q-in-Q interface, enable MAC address rewriting for Layer 2 protocol tunneling and select the Layer 2 protocol to be tunneled from the list of options:



**NOTE:** You can select only one layer 2 protocol at a time. If you want an interface to support more than one layer 2 protocol, you must enter the `mac-rewrite` statement multiple times to select the desired protocols.

[edit protocols]

```
user@switch# set layer2-control mac-rewrite interface interface-name protocol (802.3AH | CDP | LACP | LLDP | MVRP | STP | VTP | GVRP | VSTP)
```

#### Related Documentation

- [Understanding Layer 2 Protocol Tunneling on EX Series Switches](#)

## Configuring Redundant Trunk Groups (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

A redundant trunk link provides a simple solution for network recovery when a trunk interface goes down. Traffic is routed to another trunk interface, keeping network convergence time to a minimum. You can configure redundant trunk groups (RTGs) with a primary link and a secondary link on trunk interfaces, or configure dynamic selection of the active interface. If the primary link fails, the secondary link automatically takes over without waiting for normal Spanning Tree Protocol (STP) convergence. An RTG can be created only if the following conditions are satisfied:

- A minimum of two trunk interfaces that are not part of any RTG are available.
- All the selected trunk interfaces to be added to the RTG have the same VLAN configuration.
- The selected trunk interfaces are not part of a spanning-tree configuration.

To configure an RTG by using the J-Web interface:

1. Select **Configure > Switching > RTG**.

The RTG Configuration page displays a list of existing RTGs. If you select a specific RTG, the details of the selected RTG are displayed in the Details of group section.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—Creates an RTG.
- **Edit**—Modifies an RTG.
- **Delete**—Deletes an RTG.

When you are adding or editing an RTG, enter information as described in [Table 241 on page 2359](#).

3. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.

**Table 241: RTG Configuration Fields**

| Field                                  | Function                                                                                                                                         | Your Action                                                                                                              |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Group Name                             | Specifies a unique name for the RTG.                                                                                                             | Enter a name.<br><br><b>NOTE:</b> Only on EX4300 switches, you can select the name from a list.                          |
| Member Interface 1                     | Specifies a logical interface containing multiple trunk interfaces.                                                                              | Select a trunk interface from the list.                                                                                  |
| Member Interface 2                     | Specifies a trunk interface containing multiple VLANs.                                                                                           | Select a trunk interface from the list.                                                                                  |
| Select Primary Interface               | Enables you to specify one of the interfaces in the RTG as the primary link. The interface without this option is the secondary link in the RTG. | <ol style="list-style-type: none"> <li>1. Select the option button.</li> <li>2. Select the primary interface.</li> </ol> |
| Dynamically select my active interface | Specifies that the system dynamically select the active interface.                                                                               | Select the option button.                                                                                                |

**Related Documentation**

- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 2326](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery](#)
- [Understanding Redundant Trunk Links on page 2276](#)

## Configuring Proxy ARP (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches and QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Proxy ARP (CLI Procedure)* or *Configuring Proxy ARP*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

You can configure proxy Address Resolution Protocol (ARP) on your switch to enable the switch to respond to ARP queries for network addresses by offering its own media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

To configure proxy ARP on a single interface:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number proxy-arp (restricted |
unrestricted)
```



**BEST PRACTICE:** We recommend that you configure proxy ARP in restricted mode. In restricted mode, the switch does not act as a proxy if the source and target IP addresses are on the same subnet. If you decide to use unrestricted mode, disable gratuitous ARP requests on the interface to avoid a situation wherein the switch's response to a gratuitous ARP request appears to the host to be an indication of an IP conflict.

To configure proxy ARP on an integrated routing and bridging (IRB) interface:

```
[edit interfaces]
user@switch# set irb.logical-unit-number proxy-arp restricted
```

### Related Documentation

- [Example: Configuring Proxy ARP on an EX Series Switch on page 2331](#)
- [Verifying That Proxy ARP Is Working Correctly on page 2439](#)
- [Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\) on page 2340](#)



## Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches and Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

The Ethernet switching table, also known as the forwarding table, specifies the known locations of VLAN nodes and the addresses of devices within those nodes. There are two ways to populate the Ethernet switching table on a switch. The easiest method is to let the switch update the table with MAC addresses.

The second way to populate the Ethernet switching table is to manually insert addresses into the table. You can do this to reduce flooding and speed up the switch's automatic learning process.

Before configuring a static MAC address, be sure that you have:

- Set up the VLAN. See “[Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)” on page 2337.

To configure an interface to have a static MAC address:

```
[edit vlans vlan-name switch-options interface interface-name]  
user@switch# set static-mac mac-address
```

### Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches](#) on page 2245

## Configuring MAC Limiting (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches and QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring MAC Limiting (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

This topic describes various ways of configuring a limitation on MAC addresses in packets that are received and forwarded by the switch.



**NOTE:** On a QFX Series Virtual Chassis, if you include the shutdown option at the [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level and issue the commit operation, the system generates a commit error. The system does not generate an error if you include the shutdown option at the [edit switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level.

The different ways of setting a MAC limit are described in the following sections:

- [Limiting the Number of MAC Addresses Learned by an Interface on page 2362](#)
- [Limiting the Number of MAC Addresses Learned by a VLAN on page 2362](#)

### Limiting the Number of MAC Addresses Learned by an Interface

To secure a port, you can set the maximum number of MAC addresses that can be learned by an interface:

- Set the MAC limit on an interface, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit switch-options]
user@switch# set interface interface-name interface-mac-limit limit packet-action
action
```

After you set a new MAC limit for the interface, the system clears existing entries in the MAC address forwarding table associated with the interface.

### Limiting the Number of MAC Addresses Learned by a VLAN

To limit the number of MAC addresses learned by a VLAN, perform both of the following steps:

1. Set the maximum number of MAC addresses that can be learned by a VLAN, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit vlans]
user@switch# set vlan-name switch-options mac-table-size limit packet-action
action
```

2. Set the maximum number of MAC addresses that can be learned by one or all interfaces in the VLAN, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit vlans]
user@switch# set vlan-name switch-options interface interface-name
interface-mac-limit limit packet-action action
[edit vlans]
user@switch# set vlan-name switch-options interface-mac-limit limit packet-action
action
```



**NOTE:** If you specify a MAC limit and packet action for all interfaces in the VLAN *and* a specific interface in the VLAN, the MAC limit and packet action specified at the specific interface level takes precedence. Also, at the VLAN interface level, only the drop and drop-and-log options are supported.

After you set new MAC limits for a VLAN by using the **mac-table-size** statement or for interfaces associated with a VLAN by using the **interface-mac-limit** statement, the system clears the corresponding existing entries in the MAC address forwarding table.

**Related Documentation**

- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 4569](#)

## Configuring MAC Notification (CLI Procedure)



**NOTE:** This task uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring MAC Notification (CLI Procedure)* or *Configuring MAC Notification*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

When a switch learns or unlearns a MAC address, SNMP notifications can be sent to the network management system at regular intervals to record the addition or removal of the MAC address. This process is known as MAC notification.

The MAC notification interval defines how often Simple Network Management Protocol (SNMP) notifications logging the addition or removal of MAC addresses on the switch are sent to the network management system.

MAC notification is disabled by default. When MAC notification is enabled, the default MAC notification interval is 30 seconds.

To enable or disable MAC notification, or to set the MAC notification interval, perform these tasks:

- [Enabling MAC Notification on page 2363](#)
- [Disabling MAC Notification on page 2364](#)
- [Setting the MAC Notification Interval on page 2364](#)

### Enabling MAC Notification

MAC notification is disabled by default. You need to perform this procedure to enable MAC notification.

To enable MAC notification on the switch with the default MAC notification interval of 30 seconds:

[edit switch-options]

```
user@switch# set mac-notification
```

To enable MAC notification on the switch with any other MAC notification interval (here, the MAC notification interval is set to 60 seconds):

```
[edit switch-options]
user@switch# set mac-notification notification-interval 60
```

---

### Disabling MAC Notification

MAC notification is disabled by default. Perform this procedure only if MAC notification was previously enabled on your switch.

To disable MAC notification on the switch:

```
[edit switch-options]
user@switch# delete mac-notification
```

To disable MAC notification on a specific interface (here, the interface is ge-0/0/3):

```
[edit switch-options]
user@switch# set interface ge-0/0/3 no-mac-notification
```

---

### Setting the MAC Notification Interval

The default MAC notification interval is 30 seconds. The procedure to change the MAC notification interval to a different interval is identical to the procedure to enable MAC notification on the switch with a nondefault value for the MAC notification interval.

To set the MAC notification interval on the switch (here, the MAC notification interval is set to 5 seconds):

```
[edit switch-options]
user@switch# set mac-notification notification-interval 5
```

#### Related Documentation

- [Verifying That MAC Notification Is Working Properly on page 2439](#)

---

## Configuration Statements

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)
- [\[edit switch-options\] Configuration Statement Hierarchy on EX Series Switches on page 2367](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX4300 Switches on page 2368](#)
- [\[edit vlans\] Configuration Statement Hierarchy on EX Series Switches on page 2370](#)
- [address on page 2373](#)
- [aggregated-ether-options on page 2374](#)
- [community-vlan on page 2375](#)
- [description \(Interfaces\) on page 2376](#)
- [description \(VLANs\) on page 2377](#)
- [encapsulation \(Physical Interface\) on page 2378](#)

- [ether-options](#) on page 2383
- [ethernet-switch-profile](#) on page 2384
- [filter \(VLANs\)](#) on page 2386
- [flexible-vlan-tagging](#) on page 2387
- [global-mac-table-aging-time](#) on page 2388
- [global-no-mac-learning](#) on page 2388
- [input-vlan-map](#) on page 2389
- [inter-switch-link](#) on page 2390
- [interface \(MVRP\)](#) on page 2391
- [interface \(VLANs\)](#) on page 2392
- [interface-mac-limit](#) on page 2393
- [interface-mode](#) on page 2395
- [isolated-vlan](#) on page 2396
- [join-timer \(MVRP\)](#) on page 2397
- [l3-interface \(VLANs\)](#) on page 2398
- [layer2-control](#) on page 2399
- [leaveall-timer \(MVRP\)](#) on page 2401
- [leave-timer \(MVRP\)](#) on page 2402
- [mac \(Static MAC-Based VLANs\)](#) on page 2403
- [mac-notification](#) on page 2403
- [mac-rewrite](#) on page 2404
- [mac-table-size](#) on page 2405
- [members](#) on page 2407
- [mvrp](#) on page 2409
- [native-vlan-id](#) on page 2410
- [no-attribute-length-in-pdu](#) on page 2411
- [no-dynamic-vlan](#) on page 2412
- [no-gratuitous-arp-request](#) on page 2413
- [no-mac-learning](#) on page 2414
- [output-vlan-map](#) on page 2416
- [notification-interval](#) on page 2417
- [packet-action](#) on page 2418
- [pop](#) on page 2420
- [private-vlan](#) on page 2421
- [proxy-arp](#) on page 2422
- [push](#) on page 2423
- [redundant-trunk-group](#) on page 2424

- [registration on page 2425](#)
- [swap on page 2426](#)
- [tag-protocol-id \(TPIDs Expected to Be Sent or Received\) on page 2427](#)
- [vlan \(802.1Q Tagging\) on page 2428](#)
- [vlan-id \(802.1Q Tagging\) on page 2429](#)
- [vlan-id \(VLAN Tagging and Layer 3 Subinterfaces\) on page 2430](#)
- [vlan-id-list on page 2431](#)
- [vlans on page 2433](#)

## **[edit interfaces] Configuration Statement Hierarchy on EX Series Switches**

Each of the following topics lists the statements at a subhierarchy of the **[edit interfaces]** hierarchy:

- [\[edit interfaces ae\] Configuration Statement Hierarchy on EX Series Switches on page 2700](#)
- [\[edit interfaces et\] Configuration Statement Hierarchy on EX Series Switches on page 2705](#)
- [\[edit interfaces ge\] Configuration Statement Hierarchy on EX Series Switches on page 2711](#)
- [\[edit interfaces interface-range\] Configuration Statement Hierarchy on EX Series Switches on page 2716](#)
- [\[edit interfaces irb\] Configuration Statement Hierarchy on EX Series Switches on page 2725](#)
- [\[edit interfaces lo\] Configuration Statement Hierarchy on EX Series Switches on page 360](#)
- [\[edit interfaces me\] Configuration Statement Hierarchy on EX Series Switches on page 363](#)
- [\[edit interfaces vme\] Configuration Statement Hierarchy on EX Series Switches on page 370](#)
- [\[edit interfaces xe\] Configuration Statement Hierarchy on EX Series Switches on page 2738](#)

### **Related Documentation**

- [EX Series Switches Interfaces Overview on page 2577](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 2689](#)
- [Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\) on page 2340](#)
- [Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis \(CLI Procedure\)](#)
- [Junos OS Interfaces Fundamentals Configuration Guide](#)

- *Junos OS Ethernet Interfaces Configuration Guide*

## [edit switch-options] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit switch-options]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit switch-options\] Hierarchy Level on page 2367](#)
- [Unsupported Statements in the \[edit switch-options\] Hierarchy Level on page 2368](#)

### Supported Statements in the [edit switch-options] Hierarchy Level

The following hierarchy shows the **[edit switch-options]** configuration statements supported on EX Series switches:

```
switch-options {
  authentication-whitelist mac-address {
    interface interface-name;
    vlan-assignment (vlan-id | vlan-name);
  }
  interface interface-name {
    interface-mac-limit number {
      packet-action action;
    }
    no-mac-learning;
    persistent-learning
  }
  no-mac-learning;
  redundant-trunk-group {
    group name {
      description text;
      interface interface-name {
        primary;
      }
      preempt-cutover-timer seconds
    }
  }
  unknown-unicast-forwarding {
    vlan (all | vlan-name | vlan-tag) {
      interface interface-name;
    }
  }
}
```

```

voip {
  interface (all | [interface-name | access-ports]) {
    forwarding-class (assured-forwarding | best-effort | expedited-forwarding | mcast-af
      | mcast-be | mcast-ef | mcast-nc | network-control);
    vlan vlan-name;
  }
}

```

### Unsupported Statements in the [edit switch-options] Hierarchy Level

All statements in the **[edit switch-options]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 242: Unsupported [edit switch-options] Configuration Statements on EX Series Switches**

| Statement                                                                                  | Hierarchy Level                          |
|--------------------------------------------------------------------------------------------|------------------------------------------|
| NOTE: Variables, such as <i>filename</i> , are not shown in the statements or hierarchies. |                                          |
| port-error-disable                                                                         | [edit switch-options]                    |
| disable-timeout                                                                            | [edit switch-options port-error-disable] |

### [edit protocols] Configuration Statement Hierarchy on EX4300 Switches

Each of the following topics lists the statements at a subhierarchy of the **[edit protocols]** hierarchy:

- [\[edit protocols bfd\] Configuration Statement Hierarchy on EX Series Switches on page 391](#)
- [\[edit protocols bgp\] Configuration Statement Hierarchy on EX Series Switches on page 392](#)
- [\[edit protocols dot1x\] Configuration Statement Hierarchy on EX Series Switches on page 404](#)
- [\[edit protocols igmp\] Configuration Statement Hierarchy on EX Series Switches on page 406](#)
- [\[edit protocols igmp-snooping\] Configuration Statement Hierarchy on page 3743](#)
- [\[edit protocols isis\] Configuration Statement Hierarchy on EX Series Switches on page 408](#)
- [\[edit protocols lacp\] Configuration Statement Hierarchy on EX Series Switches on page 411](#)
- [\[edit protocols l2-learning\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols layer2-control\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols lldp\] Configuration Statement Hierarchy on EX Series Switches on page 413](#)



- [\[edit protocols lldp-med\] Configuration Statement Hierarchy on EX Series Switches on page 415](#)
- [\[edit protocols msdp\] Configuration Statement Hierarchy on EX Series Switches on page 429](#)
- [\[edit protocols mstp\] Configuration Statement Hierarchy on EX Series Switches on page 431](#)
- [\[edit protocols mvrp\] Configuration Statement Hierarchy on EX Series Switches on page 433](#)
- [\[edit protocols neighbor-discovery\] Configuration Statement Hierarchy on EX Series Switches on page 434](#)
- [\[edit protocols oam\] Configuration Statement Hierarchy on EX Series Switches on page 435](#)
- [\[edit protocols ospf\] Configuration Statement Hierarchy on EX Series Switches on page 438](#)
- [\[edit protocols ospf3\] Configuration Statement Hierarchy on EX Series Switches on page 441](#)
- [\[edit protocols pim\] Configuration Statement Hierarchy on EX Series Switches on page 444](#)
- [\[edit protocols rip\] Configuration Statement Hierarchy on EX Series Switches on page 447](#)
- [\[edit protocols ripng\] Configuration Statement Hierarchy on EX Series Switches on page 450](#)
- [\[edit protocols router-advertisement\] Configuration Statement Hierarchy on EX Series Switches on page 451](#)
- [\[edit protocols router-discovery\] Configuration Statement Hierarchy on EX Series Switches on page 452](#)
- [\[edit protocols rstp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches on page 459](#)
- [\[edit protocols uplink-failure-detection\] Configuration Statement Hierarchy on EX Series Switches on page 461](#)
- [\[edit protocols vrrp\] Configuration Statement Hierarchy on EX Series Switches on page 462](#)
- [\[edit protocols vstp\] Configuration Statement Hierarchy on EX Series Switches](#)

**Related  
Documentation**

- [EX Series Switch Software Features Overview](#)

## [edit vlans] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit vlans]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit vlans\] Hierarchy Level on page 2370](#)
- [Unsupported Statements in the \[edit vlans\] Hierarchy Level on page 2372](#)

### Supported Statements in the [edit vlans] Hierarchy Level

---

The following hierarchy shows the **[edit vlans]** configuration statements supported on one or more of the EX Series switches:

```
vlans {
  vlan-name {
    description text-description;
    domain-type bridge;
    forwarding-options {
      dhcp-security {
        arp-inspection;
        group group-name {
          interface interface-name {
            static-ip ip-address {
              mac mac-address;
            }
          }
        }
        overrides {
          no-option82;
          trusted;
        }
      }
    }
    ip-source-guard;
    no-dhcp-snooping;
    option-82 {
      circuit-id {
        prefix {
          host-name;
          logical-system-name;
          routing-instance-name;
        }
      }
      use-interface-description (device | logical);
      use-vlan-id;
    }
  }
}
```

```

        remote-id {
            host-name;
            use-interface-description (device | logical);
            use-string string;
        }
        vendor-id {
            use-string string;
        }
    }
}
filter {
    input filter-name;
    output filter-name;
}
flood {
    input filter-name;
}
}
l3-interface irb.logical-unit-number;
multicast-snooping-options {
    flood-groups [group-names];
    forwarding-cache {
        threshold {
            reuse threshold;
            suppress threshold;
        }
    }
}
graceful-restart {
    disable;
    restart-duration duration;
}
host-outbound-traffic {
    dot1p bits;
    forwarding-class forwarding-class;
}
multichassis-lag-replicate-state;
nexthop-hold-time time;
options {
    syslog {
        level level;
        mark interval;
        upto level;
    }
}
}
traceoptions {
    file filename {
        files number;
        no-world-readable;
        size file-size;
        world-readable;
    }
    flag flag {
        disable;
    }
}
}
}

```

```

switch-options {
  interface interface-name {
    interface-mac-limit limit {
      packet-action action;
    }
    static-mac mac-address;
  }
  interface-mac-limit limit {
    packet-action action;
  }
  mac-move-limit limit {
    packet-action action;
  }
  mac-table-size limit {
    packet-action drop;
  }
  no-mac-learning;
}
vlan-id number;
vlan-id-list [vlan-id | vlan-id-vlan-id];
}

```

#### Unsupported Statements in the [edit vlans] Hierarchy Level

All statements in the **[edit vlans]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 243: Unsupported [edit vlans] Configuration Statements on EX Series Switches**

| Statement                                                                                         | Hierarchy Level |
|---------------------------------------------------------------------------------------------------|-----------------|
| <b>NOTE:</b> Variables, such as <i>filename</i> , are not shown in the statements or hierarchies. |                 |
| mcae-mac-synchronize                                                                              | [edit vlans]    |
| no-irb-layer-2-copy                                                                               | [edit vlans]    |

**Related Documentation**

- [Example: Connecting Access Switches to a Distribution Switch on page 2291](#)

## address

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>address (<i>ip-address</i>   <i>ipv6-address</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>],</p> <p>[edit dynamic-profiles <i>profile-name</i> interfaces demux0 unit <i>logical-unit-number</i> family <i>family</i>],</p> <p>[edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit “\$junos-interface-unit” family <i>family</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2.</p> <p>Support at the [edit dynamic-profiles <i>profile-name</i> interfaces pp0 unit “\$junos-interface-unit” family <i>family</i>] hierarchy level introduced in Junos OS Release 10.1.</p> <p>Support at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i>] hierarchy level introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p>                                                                                                                                                                                     |
| <b>Description</b>              | Configure the interface address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><i>ip-address</i>—IPv4 address of the interface.</p> <p><i>ipv6-address</i>—IPv6 address of the interface. When configuring an IPv6 address on a dynamically created interface, use the <i>\$junos-ipv6-address</i> dynamic variable.</p>                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring the Protocol Family</i></li> <li><i>Format for Specifying IP Addresses, Network Masks, and Prefixes in Junos OS Configuration Statements</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                         |

## aggregated-ether-options

```
Syntax  aggregated-ether-options {
        ethernet-switch-profile {
            tag-protocol-id;
        }
        (flow-control | no-flow-control);
        lacp {
            (active | passive);
            admin-key key;
            periodic interval;
            system-id mac-address;
        }
        (link-protection | no-link-protection);
        link-speed speed;
        local-bias;
        logical-interface-fpc-redundancy;
        (loopback | no-loopback);
        mc-ae {
            chassis-id chassis-id;
            events {
                iccp-peer-down {
                    force-icl-down;
                    prefer-status-control-active;
                }
            }
            init-delay-time seconds;
            mc-ae-id mc-ae-id;
            mode (active-active | active-standby);
            redundancy-group group-id;
            revert-time revert-time;
            status-control (active | standby);
            switchover-mode (non-revertive | revertive);
        }
        minimum-links number;
        system-priority
    }
```

**Hierarchy Level** [edit interfaces aex]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 12.3R2.

**Description** Configure the aggregated Ethernet properties of a specific aggregated Ethernet interface.  
  
The remaining statements are explained separately.


**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

- [Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 2671](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\) on page 2672](#)
- [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 2351](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

## community-vlan

|                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                        |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                | <code>community-vlan vlan <i>community-vlan-name</i>;</code>                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                       | [edit vlans <i>primary-vlan-name</i> vlan-id <i>primary-vlan-vlan-id</i> ]                                                                                                                                                                                             |
| <b>Release Information</b>                                                                                                                                                                                                                                                                   | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                          |
| <b>Description</b>                                                                                                                                                                                                                                                                           | Configure the specified community VLAN to be a secondary VLAN of the specified primary VLAN. A <i>community</i> VLAN is used to transport frames among members of a community (a subset of users within the VLAN), and to forward frames upstream to the primary VLAN. |
| <div>  <b>NOTE:</b> Before you specify this configuration statement, you must have already configured the specified community VLAN and assigned a VLAN ID to it. See <a href="#">private-vlan</a>. </div> |                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                              | system—To view this statement in the configuration.<br>system—control—To add this statement to the configuration.                                                                                                                                                      |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li>• <a href="#">Creating a Private VLAN on a Single Switch (CLI Procedure) on page 2343</a></li> <li>• <a href="#">Creating a Private VLAN Spanning Multiple Switches (CLI Procedure) on page 2345</a></li> </ul>                 |

## description (Interfaces)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description text;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <code>[edit interfaces <i>interface-name</i>],</code><br><code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the <b>show interfaces</b> commands, and is also exposed in the <b>ifAlias</b> Management Information Base (MIB) object. It has no effect on the operation of the interface on the router or switch.</p> <p>The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.</p>                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>text</b> —Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <b>interface</b> —To view this statement in the configuration.<br><b>interface-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Interface Description</i></li><li>• <a href="#">Adding a Logical Unit Description to the Configuration on page 2629</a></li><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li><li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li><li>• <a href="#">Using DHCP Relay Agent Option 82 Information on page 1464</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li><li>• <a href="#">Example: Connecting Access Switches to a Distribution Switch on page 2291</a></li></ul> |



## description (VLANs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description <i>text-description</i>;</code>                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> <i>vlan-name</i> ]                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Option <b>text-description</b> enhanced from supporting up to 128 characters to supporting up to 256 characters in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                     |
| <b>Description</b>              | Provide a textual description of the VLAN. The text has no effect on the operation of the VLAN or switch.                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b><i>text-description</i></b> —Text to describe the interface. It can contain letters, numbers, and hyphens (-) and can contain 256 characters. If the text includes spaces, enclose the entire text in quotation marks.                                                                                                                                                                    |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show vlans</a></li> <li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch</a></li> <li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 2281</a></li> <li>• <a href="#">Understanding Bridging and VLANs on EX Series Switches on page 2245</a></li> </ul> |

## encapsulation (Physical Interface)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>encapsulation (atm-ccc-cell-relay   atm-pvc   cisco-hdlc   cisco-hdlc-ccc   cisco-hdlc-tcc   ethernet-bridge   ethernet-ccc   ethernet-over-atm   ethernet-tcc   ethernet-vpls   ethernet-vpls-fr   ether-vpls-over-atm-llc   ethernet-vpls-ppp   extended-frame-relay-ccc   extended-frame-relay-ether-type-tcc   extended-frame-relay-tcc   extended-vlan-bridge   extended-vlan-ccc   extended-vlan-tcc   extended-vlan-vpls   flexible-ethernet-services   flexible-frame-relay   frame-relay   frame-relay-ccc   frame-relay-ether-type   frame-relay-ether-type-tcc   frame-relay-port-ccc   frame-relay-tcc   generic-services   multilink-frame-relay-uni-nni   ppp   ppp-ccc   ppp-tcc   vlan-ccc   vlan-vci-ccc   vlan-vpls);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | <code>[edit interfaces <i>interface-name</i>],</code><br><code>[edit interfaces rlsq <i>number:number</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers ( <b>flexible-ethernet-services</b> , <b>ethernet-ccc</b> , and <b>ethernet-tcc</b> options only).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | Specify the physical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>             | <b>ppp</b> —Use serial PPP encapsulation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>             | <b>atm-ccc-cell-relay</b> —Use ATM cell-relay encapsulation.<br><br><b>atm-pvc</b> —Use ATM PVC encapsulation.<br><br><b>cisco-hdlc</b> —Use Cisco-compatible High-Level Data Link Control (HDLC) framing.<br><br><b>cisco-hdlc-ccc</b> —Use Cisco-compatible HDLC framing on CCC circuits.<br><br><b>cisco-hdlc-tcc</b> —Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.<br><br><b>ethernet-bridge</b> —Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.<br><br><b>ethernet-ccc</b> —Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, CCC is not supported.<br><br><b>ethernet-over-atm</b> —For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i> , this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination |

IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.

**ethernet-tcc**—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

**ethernet-vpls**—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

**ethernet-vpls-fr**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

**ethernet-vpls-ppp**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.

**ether-vpls-over-atm-llc**—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

**extended-frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC.

**extended-frame-relay-ether-type-tcc**—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.

**extended-frame-relay-tcc**—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

**extended-vlan-bridge**—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

**extended-vlan-ccc**—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.

**extended-vlan-tcc**—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

**extended-vlan-vpls**—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



**NOTE:** The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

---

**flexible-ethernet-services**—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

**flexible-frame-relay**—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.

**frame-relay**—Use Frame Relay encapsulation.

**frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits.

**frame-relay-ether-type**—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay.

**frame-relay-ether-type-tcc**—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media.

**frame-relay-port-ccc**—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. When you use this encapsulation type, you can configure the **ccc** family only.

**frame-relay-tcc**—Use Frame Relay encapsulation on TCC circuits to connect different media.

**generic-services**—Use generic services encapsulation for services with a hierarchical scheduler.

**multilink-frame-relay-uni-nni**—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

**ppp**—Use serial PPP encapsulation.

**ppp-ccc**—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

**ppp-tcc**—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

**vlan-ccc**—Use Ethernet VLAN encapsulation on CCC circuits.

**vlan-vci-ccc**—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.

**vlan-vpls**—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



#### NOTE:

- Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.
- Starting with Junos OS release 13.3, a commit error occurs when you configure **vlan-vpls** encapsulation on a physical interface and configure **family inet** on one of the logical units. Previously, it was possible to commit this invalid configuration.

|                                 |                                                               |
|---------------------------------|---------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.        |
|                                 | interface-control—To add this statement to the configuration. |

**Related  
Documentation**

- *Configuring Interface Encapsulation on Physical Interfaces*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM Interface Encapsulation*
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring VLAN Encapsulation*
- *Configuring Extended VLAN Encapsulation*
- *Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces*
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- *Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)*
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Configuring TCC*
- *Configuring VPLS Interface Encapsulation*
- *Configuring Interfaces for VPLS Routing*
- *Defining the Encapsulation for Switching Cross-Connects*
- [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 2351](#)

## ether-options

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ether-options {   802.3ad {     aex;     (backup   primary);     lacp {       force-up;       port-priority     }   }   (auto-negotiation   no-auto-negotiation);   ethernet-switch-profile {     tag-protocol-id;   }   (flow-control   no-flow-control);   ieee-802-3az-eee;   link-mode <i>mode</i>;   (loopback   no-loopback);   speed (<i>speed</i>   auto-negotiation); }</pre>                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <pre>[edit interfaces <i>interface-name</i>], [edit interfaces interface-range <i>range</i>]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3R2.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Configure Ethernet properties for a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 2619</a></li> <li>• <a href="#">Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 2672</a></li> <li>• <a href="#">Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</a></li> <li>• <i>Junos OS Ethernet Interfaces Configuration Guide</i></li> </ul> |

## ethernet-switch-profile

**Syntax**

```

ethernet-switch-profile {
  ethernet-policer-profile {
    input-priority-map {
      ieee802.1p premium [values];
    }
    output-priority-map {
      classifier {
        premium {
          forwarding-class class-name {
            loss-priority (high | low);
          }
        }
      }
    }
  }
  policer cos-policer-name {
    aggregate {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    premium {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
  }
  tag-protocol-id tpid;
}
mac-learn-enable;

```

**Hierarchy Level** [edit interfaces *interface-name* gigether-options],  
 [edit interfaces *interface-name* aggregated-ether-options],  
 [edit interfaces *interface-name* **aggregated-ether-options**],  
 [edit interfaces *interface-name* ether-options]

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.  
 Statement introduced in Junos OS Release 13.2 for the QFX Series.  
 Statement introduced in Junos OS Release 13.2X50-D15 for the EX Series switches.

**Description**



**NOTE:** On QFX Series standalone switches, the **ethernet-policer-profile** CLI hierarchy and the **mac-learn-enable** statement are supported only on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, aggregated Ethernet with Gigabit Ethernet IQ interfaces, the built-in Gigabit Ethernet port on the M7i router); 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit



Ethernet, and aggregated Ethernet interfaces on EX Series switches, configure VLAN tag and MAC address accounting and filtering properties.

The remaining statements are explained separately.



**NOTE:** When you gather interfaces into a bridge domain, the `no-mac-learn-enable` statement at the `[edit interfaces interface-name together-options ethernet-switch-profile]` hierarchy level is not supported. You must use the `no-mac-learning` statement at the `[edit bridge-domains bridge-domain-name bridge-options interface interface-name]` hierarchy level to disable MAC learning on an interface in a bridge domain. For information on disabling MAC learning for a bridge domain, see the *MX Series Layer 2 Configuration Guide*.

|                                 |                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | If the <b>ethernet-switch-profile</b> statement is not configured, Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) behave like Gigabit Ethernet interfaces.                                                 |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Gigabit Ethernet Policers</i></li> <li>• <i>Configuring MAC Address Filtering</i></li> <li>• <i>Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview</i></li> <li>• <a href="#">Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</a></li> </ul> |

## filter (VLANs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>filter (input   output) <i>filter-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <b>vlan</b> <i>vlan-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Apply a firewall filter to traffic coming into or exiting from the VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>                  | All incoming traffic is accepted unmodified to the VLAN, and all outgoing traffic is sent unmodified from the VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b><i>filter-name</i></b> —Name of a firewall filter defined in a <b>filter</b> statement.</p> <ul style="list-style-type: none"><li>• <b>input</b>—Apply a firewall filter to VLAN ingress traffic.</li><li>• <b>output</b>—Apply a firewall filter to VLAN egress traffic.</li></ul>                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li><li>• <a href="#">Configuring Firewall Filters (CLI Procedure) on page 4804</a></li><li>• <a href="#">Configuring Firewall Filters (J-Web Procedure) on page 4813</a></li><li>• <a href="#">Firewall Filters for EX Series Switches Overview on page 4696</a></li><li>• <a href="#">Configuring VLANs for EX Series Switches (CLI Procedure) on page 2337</a></li></ul> |

## flexible-vlan-tagging

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | flexible-vlan-tagging;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit interfaces aex],<br>[edit interfaces ge- <i>fpc/pic/port</i> ],<br>[edit interfaces et- <i>fpc/pic/port</i> ],<br>[edit interfaces ps0],<br>[edit interfaces xe- <i>fpc/pic/port</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Support for aggregated Ethernet added in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.<br>Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.</p> <p>This statement is supported on M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2, IQ2-E, and IQ PICs or in MX Series DPCs, or on Ethernet interfaces for PTX Series Packet Transport Routers or 100-Gigabit Ethernet Type 5 PIC with CFP. This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches.</p> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Mixed Tagging</i></li> <li>• <i>Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## global-mac-table-aging-time

---

|                                 |                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | global-mac-table-aging-time <i>seconds</i> ;                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols l2-learning]                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                            |
| <b>Description</b>              | Configure the timeout interval for entries in the MAC table.                                                                                                                                                           |
| <b>Default</b>                  | 300 seconds                                                                                                                                                                                                            |
| <b>Options</b>                  | <b>seconds</b> —Time elapsed before MAC table entries are timed out and entries are deleted from the table.<br><b>Range:</b> For MX Series routers: 10 through 1 million; for EX Series switches: 60 through 1 million |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the MAC Table Timeout Interval</i></li><li>• <a href="#">Configuring MAC Table Aging (CLI Procedure) on page 2341</a></li></ul>                                 |

## global-no-mac-learning

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | global-no-mac-learning;                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit protocols l2-learning],<br>[edit protocols l2-learning]                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                    |
| <b>Description</b>              | Disable MAC learning on the entire device.                                                                                                                                                     |
| <b>Default</b>                  | MAC learning is enabled.                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Disabling Layer 2 Learning and Forwarding</i></li><li>• <a href="#">Understanding Q-in-Q Tunneling on EX Series Switches on page 2269</a></li></ul> |

## input-vlan-map

|                                 |                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>input-vlan-map {   (pop   pop-pop   pop-swap   push   push-push   swap   swap-push   swap-swap);   inner-tag-protocol-id <i>tpid</i>;   inner-vlan-id <i>number</i>;   tag-protocol-id <i>tpid</i>;   vlan-id <i>number</i>; }</pre>                                                                                                                                    |
| <b>Hierarchy Level</b>          | <pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</pre>                                                                                                                                                                       |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>pop-pop</b>, <b>pop-swap</b>, <b>push-push</b>, <b>swap-push</b>, and <b>swap-swap</b> statements introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.</p> |
| <b>Description</b>              | <p>For Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP interfaces, 100-Gigabit Ethernet Type 5 PIC with CFP only as well as Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, define the rewrite profile to be applied to incoming frames on this logical interface.</p> <p>The statements are explained separately.</p>         |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Stacking a VLAN Tag</i></li> <li>• <a href="#">output-vlan-map on page 2416</a></li> <li>• <a href="#">Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</a></li> </ul>                                                                                                                                                  |

## inter-switch-link

---

|                                 |                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | inter-switch-link vlan members <i>primary-vlan-name</i> ;                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching <b>interface-mode</b> trunk]                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Use this configuration statement when a private VLAN (PVLAN) spans multiple switches. The Inter-Switch Link protocol (ISL) must be configured on a trunk port of the primary VLAN in order to connect the switches composing the PVLAN to each other. You do not need to configure an ISL when a PVLAN is configured on a single switch. |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system—control—To add this statement to the configuration.                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Creating a Private VLAN Spanning Multiple Switches (CLI Procedure)</a> on page 2345</li></ul>                                                                                                                                                                                        |

## interface (MVRP)

|                            |                                                                                                                                                                                                                   |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>interface (all   <i>interface-name</i>) {   disable;   join-timer <i>milliseconds</i>;   leave-timer <i>milliseconds</i>;   leaveall-timer <i>milliseconds</i>;   registration (forbidden   normal); }</pre> |
| <b>Hierarchy Level</b>     | [edit protocols mvrp]                                                                                                                                                                                             |
| <b>Release Information</b> | Statement introduced in Junos OS Release 10.0 for EX Series switches.<br>Statement introduced in Junos OS Release 13.1 for the QFX Series.                                                                        |
| <b>Description</b>         | Specify interfaces on which to configure Multiple VLAN Registration Protocol (MVRP).                                                                                                                              |



**NOTE:** On QFX Series switches, you must configure specific interfaces—you cannot specify interface all. You can enable MVRP on an interface range.

|                                 |                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | By default, MVRP is disabled.                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>all</b>—All interfaces on the switch.</p> <p><b><i>interface-name</i></b>—Names of interface to be configured for MVRP.</p> <p>The remaining statements are explained separately.</p>                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches</i></li> <li>• <i>Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)</i></li> </ul> |

## interface (VLANs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> {<br/>    egress;<br/>    ingress;<br/>    mapping (native (push   swap)   policy   tag (push   swap));<br/>    pvlan-trunk;<br/>}</pre>                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <pre>[edit <a href="#">vlangs</a> <i>vlan-name</i>],<br/>[edit <a href="#">vlangs on page 2370</a> <i>vlan-name</i>],<br/>[edit <a href="#">vlangs</a> <i>vlan-name</i> <i>vlan-id number</i>],<br/>[edit <a href="#">vlangs on page 2370</a> <i>vlan-name</i> <i>vlan-id number</i>],<br/>[edit <a href="#">vlangs on page 2370</a> <i>vlan-name</i> <i>vlan-id-list number</i>]</pre> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | For a specific VLAN, configure an interface.                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><i>interface-name</i>—Name of a Gigabit Ethernet interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch</i></li><li>• <i>Configuring VLANs for EX Series Switches (CLI Procedure)</i></li><li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li><li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</i></li></ul>                                       |



## interface-mac-limit

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>interface-mac-limit <i>limit</i> {     <b>packet-action</b> drop; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>     | <p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>         [edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>         [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> switch-options],<br/>         [edit logical-systems <i>logical-system-name</i> switch-options interface <i>interface-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>         [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> switch-options],<br/>         [edit routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>],<br/>         [edit switch-options],<br/>         [edit switch-options interface <i>interface-name</i>],<br/>         [edit switch-options interface <i>interface-name</i>],<br/>         [edit vlans <i>vlan-name</i> switch-options],<br/>         [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>         | <p>(MX Series routers or EX Series switches only) Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |



**NOTE:** For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the `interface-mac-limit` statement or changing the `interface-mac-limit` configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the `interface-mac-limit` statement or use the `commit at` configuration statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers by using the `clear bridge mac-table` command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

---

**Default** For an access port, the default MAC limit is 1024 MAC addresses. For a trunk port, the default MAC limit is 8192 MAC addresses.

**Options** *limit*—Maximum number of MAC addresses learned from an interface.

**Range:** 1 through 131,071 MAC addresses per interface



The remaining statement is explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Layer 2 Learning and Forwarding for Bridge Domains Overview*
- *Layer 2 Learning and Forwarding for VLANs Overview*
- *Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

## interface-mode

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | interface-mode (access   trunk);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Hierarchy Level          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Release Information      | Statement introduced in Junos OS Release 9.2.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Description              | <p> <b>NOTE:</b> This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see <i>port-mode</i>. For ELS details, see <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a>.</p> <p>(QFX Series 3500 and 3600 standalone switches)—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the <b>trunk</b> option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the <b>vlan-id</b> or <b>vlan-id-list</b> statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the <b>access</b> option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the <b>vlan-id</b> statement.</p> <p> <b>NOTE:</b> On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure <b>interface-mode</b> and <b>irb</b> for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see <i>Configuring a Trunk Interface on a Bridge Network</i>.</p> |
| Options                  | <p><b>access</b>—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the <b>vlan-id</b> statement.</p> <p><b>trunk</b>—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the <b>vlan-id</b> or <b>vlan-id-list</b> statement.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Related Documentation    | <ul style="list-style-type: none"> <li>• <i>Configuring a Logical Interface for Access Mode</i></li> <li>• <i>Configuring a Logical Interface for Trunk Mode</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

- [Example: Connecting Access Switches to a Distribution Switch on page 2291](#)

## isolated-vlan

---

|                            |                                                                                                                                                                                                            |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>isolated-vlan vlan-name <i>isolated-vlan-name</i> vlan-id <i>isolated-vlan-id</i>;</code>                                                                                                            |
| <b>Hierarchy Level</b>     | <code>[edit vlans <i>primary-vlan-name</i> vlan-id <i>primary-vlan-vlan-id</i>]</code>                                                                                                                     |
| <b>Release Information</b> | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                              |
| <b>Description</b>         | Configure the specified isolated VLAN to be a secondary VLAN of the specified primary VLAN. An isolated VLAN receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN. |



**NOTE:** Before you specify this configuration statement, you must have already configured an isolated VLAN and assigned a VLAN ID to it. See [private-vlan](#).

|                                 |                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system—control—To add this statement to the configuration.                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Creating a Private VLAN on a Single Switch (CLI Procedure) on page 2343</a></li><li>• <a href="#">Creating a Private VLAN Spanning Multiple Switches (CLI Procedure) on page 2345</a></li></ul> |

## join-timer (MVRP)



|                                 |                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>join-timer <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols mvrp <b>interface</b> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.<br>Statement introduced in Junos OS Release 13.1 for the QFX Series.                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Configure the maximum number of milliseconds interfaces must wait before sending Multiple VLAN Registration Protocol (MVRP) protocol data units (PDUs).</p> <p>Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.</p> |
| <b>Default</b>                  | 200 milliseconds                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <i>milliseconds</i> —Number of milliseconds that the interface must wait before sending MVRP PDUs.                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">leave-timer on page 2402</a></li> <li>• <a href="#">leaveall-timer on page 2401</a></li> <li>• <i>Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches</i></li> <li>• <i>Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)</i></li> </ul>            |

## l3-interface (VLANs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>l3-interface <i>l3-interface-name.logical-interface-number</i> {<br/>l3-interface-ingress-counting;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> <i>vlan-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed.                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                  | No Layer 3 (routing) interface is associated with the VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b><i>l3-interface-name.logical-interface-number</i></b>—Name of the Layer 3 interface and number of the logical interface defined by using the <b>set interfaces <a href="#">vlan</a> unit</b> command. The name of the Layer 3 interface is <b>irb</b> for an integrated routing and bridging (IRB) interface, and <b>vlan</b> for a routed VLAN interface (RVI). The number of the logical interface is the same number that you configure in the <b>unit</b> statement.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>show ethernet-switching interfaces</i></li><li>• <a href="#">show ethernet-switching interface on page 2442</a></li><li>• <i>show vlans</i></li><li>• <a href="#">show vlans on page 2474</a></li><li>• <i>Configuring Routed VLAN Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Integrated Routing and Bridging Interfaces (CLI Procedure) on page 2340</a></li></ul>                                                                                                       |

## layer2-control

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <pre> layer2-control {   bpd-block {     disable-timeout <i>seconds</i>;     interface <i>interface-name</i>;   }   mac-rewrite {     interface <i>interface-name</i> {       enable-all-ifl       protocol ( 802.1X   802.3AH   CDP   LACP   LLDP   MVRP   STP   VTP   GVRP   VSTP );     }   }   nonstop-bridging;   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>maximum-file-size</i>&gt; &lt;world-readable         no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;;   } } </pre> |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | [edit protocols]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <p>Statement introduced in Junos OS Release 8.4.</p> <p><b>bpd-block</b> statement added in Junos OS Release 9.4.</p> <p><b>enable-all-if</b> statement added in Junos OS Release 13.3.</p> <p>Support for PVSTP protocol introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.</p>                                                                                                                                                                                |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Configure Layer 2 control protocols to enable features such as Layer 2 protocol tunneling or nonstop bridging.                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <div>  <p><b>NOTE:</b> On EX4300 switches, the <b>mac-rewrite</b> statement does not support the 802.1X authentication protocol.</p> </div> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <div>  <p><b>NOTE:</b> For a detailed description of configuring the <b>nonstop-bridging</b> statement, see the <i>Junos OS High Availability Library for Routing Devices</i>. When this statement is configured on routing platforms with two Routing Engines, a master Routing Engine switches over gracefully to a backup Routing Engine and preserves Layer 2 Control Protocol (L2CP) information.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                    |

**Related  
Documentation**

- *Layer 2 Protocol Tunneling Through a Network Overview*
- *Layer 2 Protocol Tunnel Interface*
- *Layer 2 Protocol to be Tunneled*
- *Understanding Layer 2 Protocol Tunneling on EX Series Switches*
- *Configuring Layer 2 Protocol Tunneling*
- [Configuring Layer 2 Protocol Tunneling on EX Series Switches \(CLI Procedure\) on page 2357](#)
- *instance-type*



## leaveall-timer (MVRP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>leaveall-timer <i>interval</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>For platforms with ELS:           <ul style="list-style-type: none"> <li><code>[edit protocols <a href="#">mvrp</a>],</code></li> <li><code>[edit protocols <a href="#">mvrp interface</a> <i>interface-name</i>]</code></li> </ul> </li> <li>For platforms without ELS:           <ul style="list-style-type: none"> <li><code>[edit protocols mvrp <a href="#">interface</a> (all   <i>interface-name</i>)]</code></li> </ul> </li> </ul>                                                                                                                                                 |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Hierarchy level <code>[edit protocols <a href="#">mvrp</a>]</code> introduced in Junos OS Release 13.2X50-D10 (ELS). (See “<a href="#">Getting Started with Enhanced Layer 2 Software</a>” on page 3 for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.1 for the QFX Series.</p>                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>For Multiple VLAN Registration Protocol (MVRP), configure the interval at which the LeaveAll state operates on the interface.</p> <p>Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP. However, if you choose to change the default values, keep in mind that on an EX Series switch that uses Junos OS with support for ELS, if the timer value set on an interface level is different from the value set on a switch level, then the value on the interface level takes precedence.</p> |
| <b>Options</b>                  | <p><b><i>interval</i></b>—Number of seconds or milliseconds between the sending of Leave All messages.</p> <p><b>Default:</b> 10 seconds, or 10,000 milliseconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">join-timer on page 2397</a></li> <li>• <a href="#">leave-timer on page 2402</a></li> <li>• <i>Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches</i></li> <li>• <i>Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 2310</i></li> <li>• <i>Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)</i></li> <li>• <i>Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure ) on page 2348</i></li> </ul>                                                           |

## leave-timer (MVRP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>leave-timer <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols mvrp <b>interface</b> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.<br>Statement introduced in Junos OS Release 13.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>For Multiple VLAN Registration Protocol (MVRP), configure the number of milliseconds the switch retains a VLAN in the Leave state before the VLAN is unregistered. If the interface receives a join message before this timer expires, the VLAN remains registered.</p> <p>Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.</p> |
| <b>Default</b>                  | 1000 milliseconds                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <i>milliseconds</i> —Number of milliseconds that the switch retains a VLAN in the Leave state before the VLAN is unregistered. At a minimum, set the leave-timer interval at twice the join-timer interval.                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">join-timer on page 2397</a></li><li>• <a href="#">leaveall-timer on page 2401</a></li><li>• <i>Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches</i></li><li>• <i>Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)</i></li></ul>                                                                                                                              |

## mac (Static MAC-Based VLANs)

|                                 |                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mac <i>mac-address</i> {<br/>    next-hop <i>interface-name</i>;<br/>}</code>                                                          |
| <b>Hierarchy Level</b>          | <code>[edit ethernet-switching-options static vlan <i>vlan-name</i>]</code>                                                                  |
| <b>Description</b>              | Specify the MAC address to add to the Ethernet switching table.<br><br>The remaining statement is explained separately.                      |
| <b>Options</b>                  | <i>mac-address</i> —MAC address                                                                                                              |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)</i></li> </ul> |

## mac-notification

|                                 |                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mac-notification {<br/>    notification-interval <i>seconds</i>;<br/>}</code>                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit ethernet-switching-options]<br/>[edit switch-options]</code>                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Hierarchy level <code>[edit switch-options]</code> added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.                                                               |
| <b>Description</b>              | Enable MAC notification for a switch. If you configure this statement without setting a notification interval, MAC notification is enabled with the default MAC notification interval of 30 seconds.<br><br>The remaining statement is explained separately. |
| <b>Default</b>                  | MAC notification is disabled by default.                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring MAC Notification (CLI Procedure)</i></li> <li>• <a href="#">Configuring MAC Notification (CLI Procedure) on page 2363</a></li> </ul>                                                                 |

## mac-rewrite

---

**Syntax**    `mac-rewrite {  
              interface interface-name {  
                  enable-all-ifl  
                  protocol (802.1X | 802.3AH | CDP | LACP | LLDP | MVRP | STP | VTP | GVRP | VSTP);  
                  }  
              }  
              }`

**Hierarchy Level**    [edit protocols [layer2-control](#)]

**Release Information**    Statement introduced in Junos OS Release 9.1.  
                              Statement introduced in Junos OS Release 13.2 for QFX series.  
                              Support for PVSTP introduced in Junos OS Release 13.3.  
                              **enable-all-if** statement added in Junos OS Release 13.3.  
                              Support for PVSTP protocol introduced in Junos OS Release 13.3.  
                              Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

**Description**    Enable rewriting of the MAC address for Layer 2 protocol tunneling.



**NOTE:** On EX4300 switches, the `mac-rewrite` statement does not support the 802.1X authentication protocol.

---

The remaining statements are explained separately.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**

- *Layer 2 Protocol Tunneling Through a Network Overview*
- *Understanding Layer 2 Protocol Tunneling on EX Series Switches*
- [Configuring Layer 2 Protocol Tunneling on EX Series Switches \(CLI Procedure\) on page 2357](#)

## mac-table-size

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>mac-table-size <i>limit</i> {<br/>    <code>packet-action</code> drop;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>     | [edit bridge-domains <i>bridge-domain-name</i> bridge-options],<br>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],<br>[edit logical-systems <i>logical-system-name</i> switch-options],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br>[edit routing-instances <i>routing-instance-name</i> switch-options],<br>[edit switch-options],<br>[edit vlans <i>vlan-name</i> switch-options]                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.4.<br>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.<br>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.<br>Support for logical systems added in Junos OS Release 9.6.<br><b>[edit switch-options]</b> and <b>[edit vlans <i>vlan-name</i> switch-options]</b> hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Support at the <b>[edit vlans <i>vlan-name</i> switch-options]</b> hierarchy level introduced in Junos OS Release 13.2 for the QFX Series. |
| <b>Description</b>         | Modify the size of the MAC address table for the bridge domain or VLAN, a set of bridge domains or VLANs associated with a trunk port, or a virtual switch. The default is 5120 MAC addresses.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |



**NOTE:** For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the **mac-table-size** statement or changing the **mac-table-size** configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the **mac-table-size** statement or use the **commit at** configuration statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers by using the **clear bridge mac-table** command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

**Options** *limit*—Specify the maximum number of addresses in the MAC address table.

**Range:** 16 through 1,048,575 MAC addresses

**Default:** 5120 MAC addresses There is no default MAC address limit for the **mac-table-size** statement at the **[edit switch-options]** hierarchy level. The number of MAC addresses that can be learned is only limited by the platform, 65,535 MAC addresses for EX Series switches and 1,048,575 MAC addresses for other devices.

The remaining statement is explained separately.

**Required Privilege** routing—To view this statement in the configuration.

**Level** routing-control—To add this statement to the configuration.

**Related  
Documentation**

- *Layer 2 Learning and Forwarding for Bridge Domains Overview*
- *Layer 2 Learning and Forwarding for VLANs Overview*
- *Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

## members

|                            |                                                                                                                                                                                                  |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>members [(all   <i>names</i>   <i>vlan-ids</i>)];</code>                                                                                                                                   |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>family</b> ethernet-switching <b>vlan</b> ]                                                                     |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated with enhanced ? (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches. |
| <b>Description</b>         | For trunk interfaces, configure the VLANs that can carry traffic.                                                                                                                                |



**TIP:** To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after `vlan` or `vlands` in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.



**NOTE:** The number of VLANs supported per switch varies for each model. Use the configuration-mode command `set vlans id vlan-id ?` to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum.

On an EX Series switch that runs Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style, the maximum number of VLAN members allowed on the switch is 8 times the maximum number of VLANs the switch supports (`vmember limit = vlan max * 8`). If the switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (`eswd`) due to memory allocation failure.

On an EX Series switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs the switch supports (`vmember limit = vlan max * 24`). If the configuration of one of these switches exceeds the recommended VLAN member maximum, a warning message appears in the system log (`syslog`).

|                |                                                                                                                                                                                                                                |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b> | <b>all</b> —Specifies that this trunk interface is a member of all the VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN. |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|



**NOTE:** Since VLAN members are limited, specifying all could cause the number of VLAN members to exceed the limit at some point.

**names**—Name of one or more VLANs. VLAN IDs are applied automatically in this case.



**NOTE:** all cannot be a VLAN name.

**vlan-ids**—Numeric identifier of one or more VLANs. For a series of tagged VLANs, specify a range; for example, 10–20 or 10–20 23 27–30.



**NOTE:** Each configured VLAN must have a specified VLAN ID to successfully commit the configuration; otherwise, the configuration commit fails.

**Required Privilege  
Level**

interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related  
Documentation**

- *show ethernet-switching interfaces*
- [show ethernet-switching interface on page 2442](#)
- *show vlans*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 2281](#)
- *Example: Connecting an Access Switch to a Distribution Switch*
- [Example: Connecting Access Switches to a Distribution Switch on page 2291](#)
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 2619](#)
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- [Configuring VLANs for EX Series Switches \(CLI Procedure\) on page 2337](#)
- *Creating a Series of Tagged VLANs (CLI Procedure)*
- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)



## mvrp

**Syntax**

```
mvrp {
  interface interface-name {
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer seconds;
    registration (forbidden | normal);
  }
  join-timer milliseconds;
  leave-timer milliseconds;
  leaveall-timer seconds;
  no-attribute-length-in-pdu
  no-dynamic-vlan;
  traceoptions (Spanning Trees) {
    file filename <files number > <size size > <world-readable | no-world-readable>;
    flag <flag> <disable>;
  }
}
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

**Description**



**NOTE:** If your switch command-line interface (CLI) displays different options for the mvrp statement from the options shown in this document, see *mvrp*.

Configure Multiple VLAN Registration Protocol (MVRP) on a trunk interface to ensure that the VLAN membership information on the trunk interface is updated as the switch's access interfaces become active or inactive in the configured VLANs.



**NOTE:** In Junos OS Release 11.3, MVRP was updated to conform to the IEEE standard 802.1ak. This update might result in compatibility issues in mixed release networks. For details, see “[Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\)](#)” on page 2348.

The remaining statements are explained separately.

**Default** MVRP is disabled by default.

**Required Privilege Level**

|                                                             |
|-------------------------------------------------------------|
| routing—To view this statement in the configuration.        |
| routing-control—To add this statement to the configuration. |

**Related Documentation**

- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches on page 2310](#)

## native-vlan-id

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>native-vlan-id <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit interfaces <i>ge-fpc/pic/port</i> ],<br>[edit interfaces <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure mixed tagging support for untagged packets on a port for the following:</p> <ul style="list-style-type: none"><li>• M Series routers with Gigabit Ethernet IQ PICs with SFP and Gigabit Ethernet IQ2 PICs with SFP configured for 802.1Q flexible VLAN tagging</li><li>• MX Series routers with Gigabit Ethernet DPCs and MICs, Tri-Rate Ethernet DPCs and MICs, and 10-Gigabit Ethernet DPCs and MICs and MPCs configured for 802.1Q flexible VLAN tagging</li><li>• T4000 routers with 100-Gigabit Ethernet Type 5 PIC with CFP</li><li>• EX Series switches with Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces</li></ul> <p>When the <b>native-vlan-id</b> statement is included with the <a href="#">flexible-vlan-tagging</a> statement, untagged packets are accepted on the same mixed VLAN-tagged port.</p> <p>The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface. To configure the logical interface, include the <b>vlan-id</b> statement (matching the <b>native-vlan-id</b> statement on the physical interface) at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p> <p>When the <b>native-vlan-id</b> statement is included with the <a href="#">interface-mode</a> statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.</p> |
| <b>Options</b>                  | <p><b><i>number</i></b>—VLAN ID number.</p> <p><b>Range:</b> (ACX Series routers and EX Series switches) 0 through 4094.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Mixed Tagging Support for Untagged Packets</a></li><li>• <a href="#">Configuring a Logical Interface for Access Mode</a></li><li>• <a href="#">Configuring the Native VLAN Identifier (CLI Procedure) on page 2342</a></li><li>• <a href="#">Understanding Bridging and VLANs on EX Series Switches on page 2245</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

- [flexible-vlan-tagging on page 2387](#)
- [Understanding Q-in-Q Tunneling on EX Series Switches on page 2269](#)

## no-attribute-length-in-pdu

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-attribute-length-in-pdu;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mvrp</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Include an extra byte in protocol data units (PDUs) sent by the Multiple VLAN Registration Protocol (MVRP). You can disable the extra byte to address a compatibility issue between MVRP in Junos OS Releases 13.2 and later for EX Series switches with support for the Enhanced Layer 2 Software (ELS), which includes the extra byte, and MVRP in Junos OS Releases 11.3 and later for EX Series switches that do not support ELS, which does not include the extra byte. If this compatibility issue arises, the ELS version of MVRP does not recognize PDUs without the extra byte sent by the non-ELS version of MVRP.</p> <p>You can recognize an MVRP version compatibility issue by observing the switch running the ELS version of MVRP. Because a switch running the ELS version of MVRP cannot interpret an unmodified PDU from a switch running the non-ELS version of MVRP, the switch will not add VLANs from the non-ELS version of MVRP. When you execute the command <b>show mvrp statistics</b> in the ELS version of MVRP, the values for <b>Received Join Empty</b> and <b>Received Join In</b> will incorrectly display zero, even though the value for the <b>Received MVRP PDUs without error</b> has been increased. Another indication that MVRP is having a version compatibility issue is that unexpected VLAN activity, such as multiple VLAN creation, takes place on the switch running the ELS version of MVRP.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure ) on page 2348</a></li> <li>• <a href="#">Understanding Multiple VLAN Registration Protocol (MVRP) on EX Series Switches on page 2264</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## no-dynamic-vlan

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-dynamic-vlan;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols mvrp]<br>[edit protocols <a href="#">mvrp</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Disable the dynamic creation of VLANs using Multiple VLAN Registration Protocol (MVRP) for interfaces participating in MVRP.</p> <p>Dynamic VLAN configuration can be enabled on an interface independent of MVRP. The MVRP dynamic VLAN configuration setting does not override the interface configuration dynamic VLAN configuration setting. If dynamic VLAN creation is disabled on the interface in the interface configuration, no dynamic VLANs are created on the interface, including dynamic VLANs created using MVRP.</p> <p>This option can be applied globally; it cannot be applied per interface.</p> |
| <b>Default</b>                  | If MVRP is enabled, the dynamic creation of VLANs as a result of MVRP protocol exchange messages is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)</a></li><li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)</a> on page 2348</li></ul>                                                                                                                                                                                                                                                                                                                                                                   |

---

## no-gratuitous-arp-request

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-gratuitous-arp-request;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on Layer 2 Ethernet switching interfaces, and integrated routing and bridging (IRB) interfaces or routed VLAN interfaces (RVIs). (On EX Series switches that use Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style, the feature is known as an IRB interface. On EX Series switches that use Junos OS that does not support ELS, the feature is known as an RVI.) |
| <b>Default</b>                  | Gratuitous ARP responses are enabled on all Ethernet switching interfaces, and IRB interfaces or RVIs.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Proxy ARP on an EX Series Switch on page 2331</a></li><li>• <a href="#">Configuring Proxy ARP (CLI Procedure)</a></li><li>• <a href="#">Configuring Proxy ARP (CLI Procedure) on page 2360</a></li></ul>                                                                                                                                                                                                                                      |

## no-mac-learning

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | no-mac-learning;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | <p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],<br/>[edit logical-systems <i>logical-system-name</i> switch-options],<br/>[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols evpn],<br/>[edit routing-instances <i>routing-instance-name</i> protocols evpn interface <i>interface-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> switch-options],<br/>[edit switch-options],<br/>[edit <a href="#">switch-options on page 2212</a>],<br/>[edit <a href="#">switch-options on page 2212</a> interface <i>interface-name</i>],<br/>[set vlans <i>vlan-name</i> switch-options]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3 R2 for EX Series switches.</p> <p>Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.</p> <p>Hierarchy levels [edit switch-options interface <i>interface-name</i>] and [edit vlans <i>vlan-name</i> switch-options] introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | <p>For MX Series routers and EX Series switches, disable MAC learning for a virtual switch, for a bridge domain or VLAN, for a specific logical interface in a bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port. On platforms that support EVPNs, you can disable MAC learning on an EVPN.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |



**NOTE:** When MAC learning is disabled for a VPLS routing instance, traffic is not load-balanced and only one of the equal-cost next hops is used.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | MAC learning is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring EVPN Routing Instances</i></li> <li>• <i>Layer 2 Learning and Forwarding for Bridge Domains Overview</i></li> <li>• <i>Layer 2 Learning and Forwarding for VLANs Overview</i></li> <li>• <i>Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i></li> <li>• <a href="#">Understanding Bridging and VLANs on EX Series Switches on page 2245</a></li> <li>• <a href="#">Understanding Q-in-Q Tunneling on EX Series Switches on page 2269</a></li> </ul> |

## output-vlan-map

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>output-vlan-map {<br/>    (pop   pop-pop   pop-swap   push   push-push   swap   swap-push   swap-swap);<br/>    inner-tag-protocol-id <i>tpid</i>;<br/>    inner-vlan-id <i>number</i>;<br/>    tag-protocol-id <i>tpid</i>;<br/>    vlan-id <i>number</i>;<br/>}</pre>                                                                                            |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ]                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br><b>pop-pop</b> , <b>pop-swap</b> , <b>push-push</b> , <b>swap-push</b> , and <b>swap-swap</b> statements added in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.                     |
| <b>Description</b>              | For Gigabit Ethernet IQ, 10-Port 10-Gigabit Ethernet SFPP interfaces, 100-Gigabit Ethernet Type 5 PIC with CFP only, Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, define the rewrite operation to be applied to outgoing frames on this logical interface.<br><br>The remaining statements are explained separately. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Stacking and Rewriting Gigabit Ethernet VLAN Tags</i></li><li>• <a href="#">input-vlan-map on page 2389</a></li><li>• <a href="#">Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</a></li></ul>                                                                                                                    |



## notification-interval

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | notification-interval <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit ethernet-switching-options <a href="#">mac-notification</a> ]<br>[edit switch-options <a href="#">mac-notification</a> ]                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Hierarchy level <a href="#">[edit switch-options]</a> added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the MAC notification interval for a switch.<br><br>The MAC notification interval is the amount of time the switch waits before sending learned or unlearned MAC address SNMP notifications to the network management server. For instance, if the MAC notification interval is set to 10, all of the MAC address addition and removal SNMP notifications will be sent to the network management system every 10 seconds. |
| <b>Options</b>                  | <b><i>seconds</i></b> —The MAC notification interval, in seconds.<br><b>Range:</b> 1 through 60<br><b>Default:</b> 30                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring MAC Notification (CLI Procedure)</a></li> <li><a href="#">Configuring MAC Notification (CLI Procedure) on page 2363</a></li> </ul>                                                                                                                                                                                                                                  |

## packet-action

**Syntax** `packet-action action;`

**Hierarchy Level** [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* switch-options **interface-mac-limit** *limit*],  
 [edit protocols l2-learning global-mac-limit *limit*],  
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* protocols evpn interface-mac-limit (vpls)],  
 [edit routing-instances *routing-instance-name* protocols evpn interface *interface-name* interface-mac-limit (vpls)],  
 [edit routing-instances *routing-instance-name* protocols evpn mac-table-size *limit*],  
 [edit routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],  
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit switch-options **interface-mac-limit** *limit*],  
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit switch-options **interface-mac-limit** *limit*],  
 [edit switch-options **mac-table-size** *limit*],  
 [edit [switch-options on page 2212](#) interface *interface-name* **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options **mac-table-size** *limit*]  
 [edit [vlans on page 2370](#) *vlan-name* switch-options **interface-mac-limit** *limit*],  
 [edit [vlans on page 2370](#) *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit [vlans on page 2370](#) *vlan-name* switch-options **mac-table-size** *limit*]

**Release Information** Statement introduced in Junos OS Release 8.4.  
 Support for the **switch-options** statement added in Junos OS Release 9.2.  
 Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy

supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

[edit switch-options interface *interface-name* interface-mac-limit *limit*], [edit switch-options interface-mac-limit *limit*], [edit switch-options mac-table-size *limit*], [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*], [edit vlans *vlan-name* switch-options interface-mac-limit *limit*], and [edit vlans *vlan-name* switch-options mac-table-size *limit*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Support for EVPNs introduced in Junos OS Release 13.2 on MX Series 3D Universal Edge Routers.

Support at the [edit switch-options interface *interface-name* interface-mac-limit *limit*] hierarchy level and hierarchy levels under [edit vlans *vlan-name*] introduced in Junos OS Release 13.2X50-D10 for EX Series switches and Junos OS Release 13.2 for the QFX Series.

**Description** Specify the action taken when packets with new source MAC addresses are received after the MAC address limit is reached. If this statement is not configured, packets with new source MAC addresses are forwarded by default.

**Default**



**NOTE:** On a QFX Series Virtual Chassis, if you include the shutdown option at the [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level and issue the commit operation, the system generates a commit error. The system does not generate an error if you include the shutdown option at the [edit switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level.

Disabled. The default is for packets for new source MAC addresses to be forwarded after the MAC address limit is reached.

**Options**

- drop**—Drop packets with new source MAC addresses, and do not learn the new source MAC addresses.
- drop-and-log**—(EX Series switches and QFX Series only) Drop packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.
- log**—(EX Series switches and QFX Series only) Hold packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.
- none**—(EX Series switches and QFX Series only) Forward packets with new source MAC addresses, and learn the new source MAC address.
- shutdown**—(EX Series switches and QFX Series only) Disable the specified interface, and generate an alarm, an SNMP trap, or a system log entry.

**Required Privilege Level**


- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring EVPN Routing Instances](#)
  - [Configuring MAC Limiting \(CLI Procedure\) on page 2361](#)
  - [Configuring Persistent MAC Learning \(CLI Procedure\) on page 4569](#)
  - [Layer 2 Learning and Forwarding for Bridge Domains Overview](#)
  - [Layer 2 Learning and Forwarding for VLANs Overview](#)
  - [Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports](#)
  - [Layer 2 Learning and Forwarding for VLANs Overview](#)
  - [Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port](#)

---

## pop

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | pop;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ]                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <div> <b>NOTE:</b> On EX4300 switches, pop is not supported at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a>] hierarchy level.</div> <p>For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2, and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.</p> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Removing a VLAN Tag</a></li><li>• <a href="#">Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## private-vlan

|                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                            | <code>private-vlan (isolated   community) <b>vlan-id</b> <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                   | <code>[edit vlans <i>vlan-name</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                               | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                       | Configure a secondary VLAN (either an isolated VLAN or a community VLAN) within a private VLAN (PVLAN) and specify a VLAN ID for that secondary VLAN. This statement essentially converts a VLAN into a PVLAN, by carving out discrete subdomains (secondary VLANs) within the primary VLAN. You must specify a VLAN ID for each secondary PVLAN.                                                                                                                                                                                                                                                                                                                                                                            |
| <div>  <b>NOTE:</b> After you have configured the secondary VLAN, you must also configure its association with a specific primary VLAN. See <a href="#">isolated-vlan</a> and <a href="#">community-vlan</a> for additional information.         </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                           | <ul style="list-style-type: none"> <li>• <b>isolated</b> — The VLAN specified by <i>vlan-name</i> is defined as an <i>isolated</i> VLAN and a VLAN-ID is assigned to it. An isolated VLAN receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN. The VLAN name is optional. The VLAN ID is required.</li> <li>• <b>community</b> — The VLAN specified by <i>vlan-name</i> is defined as community VLAN and a VLAN-ID is assigned to it. A <i>community</i> VLAN used to transport frames among members of a community, which is a subset of users within the VLAN, and to forward frames upstream to the primary VLAN. The VLAN name is optional. The VLAN ID is required.</li> </ul> |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                          | system—To view this statement in the configuration.<br>system—control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li>• <a href="#">Creating a Private VLAN on a Single Switch (CLI Procedure) on page 2343</a></li> <li>• <a href="#">Creating a Private VLAN Spanning Multiple Switches (CLI Procedure) on page 2345</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## proxy-arp

---

|                            |                                                                                                                                                                                                                                                                            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>proxy-arp (restricted   unrestricted);</code>                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ]                                                                           |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.6 for EX Series switches.<br><b>restricted</b> added in Junos OS Release 10.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series. |
| <b>Description</b>         | For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.                                                                                       |




**NOTE:** You must configure the IP address and the inet family for the interface when you enable proxy ARP.

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <ul style="list-style-type: none"><li>• <b>none</b>—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.</li><li>• <b>restricted</b>—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address.</li><li>• <b>unrestricted</b>—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.</li></ul> |
|                                 | <b>Default:</b> <b>unrestricted</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Restricted and Unrestricted Proxy ARP on page 2665</a></li><li>• <a href="#">Configuring Proxy ARP (CLI Procedure)</a></li><li>• <a href="#">Configuring Proxy ARP (CLI Procedure) on page 2360</a></li><li>• <a href="#">Example: Configuring Proxy ARP on an EX Series Switch on page 2331</a></li><li>• <a href="#">Configuring Gratuitous ARP on page 2663</a></li></ul>                                                                                                                                                                                                                                                                         |

## push

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>push;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a>],</code><br><code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code><br><code>  <a href="#">input-vlan-map</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code><br><code>  output-vlan-map]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p> <b>NOTE:</b> On EX4300 switches, <code>push</code> is not supported at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</code> hierarchy level.</p> <p>Specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.</p> <p>You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.</p> <p>If you include the <b>push</b> statement in the configuration, you must also include the <a href="#">pop</a> statement at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</code> hierarchy level.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Stacking a VLAN Tag</a></li> <li>• <a href="#">Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## redundant-trunk-group

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>redundant-trunk-group {<br/>  group <i>name</i> {<br/>    interface <i>interface-name</i> &lt;primary&gt;;<br/>    interface <i>interface-name</i>;<br/>    preempt-cutover-timer <i>seconds</i>;<br/>  }<br/>}</pre>                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>• For platforms with ELS:<br/>[edit switch-options]</li><li>• For platforms without ELS:<br/>[edit ethernet-switching-options]</li></ul>                                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level <b>[edit switch-options]</b> introduced in Junos OS Release 13.2X50-D10 (ELS). (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> |
| <b>Description</b>              | <p>Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over without waiting for normal spanning-tree protocol convergence.</p> <p>The remaining statements are explained separately.</p>                                                                                                            |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery</a></li><li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2326</a></li><li>• <a href="#">Understanding Redundant Trunk Links on page 2276</a></li></ul>                                                             |



## registration

---

|                                 |                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | registration (forbidden   normal);                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols mvrp <b>interface</b> (all   <i>interface-name</i> )],<br>[edit protocols <b>mvrp interface</b> <i>interface-name</i> ]                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                             |
| <b>Description</b>              | Specifies the Multiple VLAN Registration Protocol (MVRP) registration mode for the interface if MVRP is enabled.                                                                                                                                  |
| <b>Default</b>                  | normal                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <b>forbidden</b> —The interface or interfaces do not register and do not participate in MVRP.<br><br><b>normal</b> —The interface or interfaces accept MVRP messages and participate in MVRP.                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)</i></li> <li>• <a href="#">Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure ) on page 2348</a></li> </ul> |

## swap

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | swap;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a> ]                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.</p> <p>On MX Series routers, you can enter this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet using Gigabit Ethernet IQ interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP. On EX Series switches, you can enter this statement on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.</p> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Rewriting the VLAN Tag on Tagged Frames</a></li><li>• <a href="#">Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                    |

## tag-protocol-id (TPIDs Expected to Be Sent or Received)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>tag-protocol-id [tpids];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit interfaces <i>interface-name</i> <b>gigether-options</b> <a href="#">ethernet-switch-profile</a>],</p> <p>[edit interfaces <i>interface-name</i> <b>aggregated-ether-options</b> <a href="#">ethernet-switch-profile</a>],</p> <p>[edit interfaces <i>interface-name</i> <b>aggregated-ether-options</b> <a href="#">ethernet-switch-profile</a>],</p> <p>[edit interfaces <i>interface-name</i> <b>ether-options</b> <a href="#">ethernet-switch-profile</a>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, and the built-in Gigabit Ethernet port on the M7i router), define the TPIDs expected to be sent or received on a particular VLAN. For each Gigabit Ethernet port, you can configure up to eight TPIDs using the <b>tag-protocol-id</b> statement; but only the first four TPIDs are supported on IQ2 and IQ2-E interfaces.</p> <p>For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers only the default TPID value (<b>0x8100</b>) is supported.</p> <p>For Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches, define the TPIDs expected to be sent or received on a particular VLAN. The default TPID value is <b>0x8100</b>. Other supported values are <b>0x88a8</b>, <b>0x9100</b>, and <b>0x9200</b>.</p> |
| <b>Options</b>                  | <b>tpids</b> —TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p><b>interface</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames</a></li> <li><a href="#">Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## vlan (802.1Q Tagging)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vlan {<br/>    <b>members</b> [(all   <i>names</i>   <i>vlan-ids</i>)];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>family</b> ethernet-switching]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Bind an 802.1Q VLAN tag ID to a logical interface.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>show ethernet-switching interfaces</i></li><li>• <a href="#">show ethernet-switching interface on page 2442</a></li><li>• <i>Example: Setting Up Bridging with Multiple VLANs for EX Series Switches</i></li><li>• <i>Configuring Routed VLAN Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Integrated Routing and Bridging Interfaces (CLI Procedure) on page 2340</a></li><li>• <a href="#">Understanding Bridging and VLANs on EX Series Switches on page 2245</a></li><li>• <i>Junos OS Ethernet Interfaces Configuration Guide</i></li></ul> |

## vlan-id (802.1Q Tagging)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vlan-id <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <b>vlan</b> <i>vlan-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Configure an 802.1Q tag to apply to all traffic that originates on the VLAN.</p> <p>The number zero is reserved for priority tagging and the number 4095 is also reserved.</p>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | If you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><i>number</i> —VLAN tag identifier</p> <p><b>Range:</b></p> <ul style="list-style-type: none"> <li>• 1 through 4094 (all switches except EX8200 Virtual Chassis)</li> <li>• 1 through 4092 (EX8200 Virtual Chassis only)</li> </ul> <p><b>Default:</b> 1</p>                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Setting Up Bridging with Multiple VLANs for EX Series Switches</i></li> <li>• <a href="#">Example: Connecting Access Switches to a Distribution Switch on page 2291</a></li> <li>• <i>Example: Configuring a Private VLAN on a Single EX Series Switch</i></li> <li>• <i>Example: Configuring a Private VLAN Spanning Multiple EX Series Switches</i></li> <li>• <i>Creating a Private VLAN on a Single EX Series Switch (CLI Procedure)</i></li> <li>• <i>Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)</i></li> </ul> |

## vlan-id (VLAN Tagging and Layer 3 Subinterfaces)

---


|                            |                                                                                 |
|----------------------------|---------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>vlan-id <i>vlan-id-number</i>;</code>                                     |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.2 for EX Series switches.            |
| <b>Description</b>         | Bind an 802.1Q VLAN tag ID to a logical interface.                              |



**NOTE:** The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <b><i>vlan-id-number</i></b> —A valid VLAN identifier.<br><b>Range:</b> 1 through 4094                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">vlan-tagging on page 2834</a></li><li>• <i>Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch</i></li><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 2619</a></li><li>• <a href="#">Configuring a Layer 3 Subinterface (CLI Procedure) on page 2689</a></li><li>• <a href="#">Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</a></li><li>• <i>Junos OS Ethernet Interfaces Configuration Guide</i></li></ul> |

## vlan-id-list

|                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                               | <code>vlan-id-list [ <i>vlan-id-numbers</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>                                                                                                                                                                      | <p>[edit bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>]</p> <p>[edit interfaces <i>interface-name</i> unit 0],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</p> <p>[edit vlans <i>vlan-name</i>]</p>                                                                                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                  | <p>Statement introduced in Junos OS Release 9.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                                                                                                                                                                          | <p>Specify a VLAN identifier list to use for a bridge domain or VLAN in trunk mode.</p> <p>Specify the <b>trunk</b> option in the <b>interface-mode</b> statement to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the <b>vlan-id-list</b> statement to forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the <b>access</b> option to accept packets with no VLAN ID to forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the <b>vlan-id</b> statement.</p> <p>This statement also enables you to bind a logical interface to a list of VLAN IDs, thereby configuring the logical interface to receive and forward a frame with a tag that matches the specified VLAN ID list.</p> |
| <b>Options</b>                                                                                                                                                                              | <p><b><i>vlan-id-numbers</i></b>—Valid VLAN identifiers. You can combine individual numbers with range lists including a hyphen.</p> <p><b>Range:</b> 0 through 4095</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <div>  <p><b>NOTE:</b> On EX Series switches and the QFX Series, the range is 0 through 4094.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                                                                                                                                                             | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>                                                                                                                                                                | <ul style="list-style-type: none"> <li>• <i>Configuring a Bridge Domain</i></li> <li>• <i>Configuring a VLAN</i></li> <li>• <a href="#">Configuring VLANs for EX Series Switches (CLI Procedure) on page 2337</a></li> <li>• <i>Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

- *Configuring VLAN Identifiers for VLANs and VPLS Routing Instances*
- [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 2351](#)



## vlan

```
Syntax  vlan {
        vlan-name {
            description text-description;
            dot1q-tunneling {
                customer-vlans (id | range)
                layer2-protocol-tunneling all | protocol-name {
                    drop-threshold number;
                    shutdown-threshold number;
                }
            }
            filter input filter-name;
            filter output filter-name;
            interface interface-name {
                egress;
                ingress;
                mapping (native (push | swap) | policy | tag (push | swap));
                pvlan-trunk;
            }
            isolation-id id-number;
            l3-interface l3-interface-name.logical-interface-number;
            l3-interface-ingress-counting layer-3-interface-name;
            mac-limit limit action action;
            mac-table-aging-time seconds;
            no-local-switching;
            no-mac-learning;
            primary-vlan vlan-name;
            vlan-id number;
            vlan-prune;
            vlan-range vlan-id-low-vlan-id-high;
        }
    }
```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure VLAN properties on EX Series switches. The following configuration guidelines apply:

- Only private VLAN (PVLAN) firewall filters can be used when the VLAN is enabled for Q-in-Q tunneling.
- An S-VLAN tag is added to the packet if the VLAN is Q-in-Q-tunneled and the packet is arriving from an access interface.
- You cannot use a firewall filter to assign an integrated routing and bridging (IRB) interface or a routed VLAN interface (RVI) to a VLAN.
- VLAN assignments performed using a firewall filter override all other VLAN assignments.

**Options** *vlan-name*—Name of the VLAN. The name can include letters, numbers, hyphens (-), and periods (.) and can contain up to 255 characters long.

The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system—control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring VLANs for EX Series Switches (CLI Procedure)</i></li><li>• <a href="#">Configuring VLANs for EX Series Switches (CLI Procedure) on page 2337</a></li><li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li><li>• <i>Creating a Series of Tagged VLANs (CLI Procedure)</i></li><li>• <i>Configuring Routed VLAN Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Integrated Routing and Bridging Interfaces (CLI Procedure) on page 2340</a></li><li>• <a href="#">Understanding Bridging and VLANs on EX Series Switches on page 2245</a></li></ul> |

# Administration

- [Routine Monitoring on page 2435](#)
- [Operational Commands on page 2440](#)

## Routine Monitoring

---

- [Verifying That Virtual Routing Instances Are Working on page 2435](#)
- [Verifying Integrated Routing and Bridging Interface Status and Statistics on page 2436](#)
- [Verifying That MVRP Is Working Correctly on page 2438](#)
- [Verifying That MAC Notification Is Working Properly on page 2439](#)
- [Verifying That Proxy ARP Is Working Correctly on page 2439](#)

## Verifying That Virtual Routing Instances Are Working

**Purpose** After creating a virtual routing instance, make sure it is set up properly.

**Action** 1. Use the **show route instance** command to list all of the routing instances and their properties:

```
user@switch> show route instance
Instance          Type
Primary RIB
master            forwarding
inet.0            3/0/0
__juniper_private1__ forwarding
__juniper_private1__.inet.0 1/0/3
__juniper_private2__ forwarding
instance1         forwarding
r1                virtual-router
r1.inet.0         1/0/0
r2                virtual-router
r2.inet.0         1/0/0
```

2. Use the **show route forwarding-table** command to view the forwarding table information for each routing instance:

```
user@switch> show route forwarding-table
```

Routing table: r1.inet

Internet:

| Destination        | Type | RtRef | Next hop    | Type | Index | NhRef | Netif      |
|--------------------|------|-------|-------------|------|-------|-------|------------|
| default            | perm | 0     |             | rjct | 539   | 2     |            |
| 0.0.0.0/32         | perm | 0     |             | dscd | 537   | 1     |            |
| 103.1.1.0/24       | ifdn | 0     |             | rslv | 579   | 1     | ge-0/0/3.0 |
| 103.1.1.0/32       | iddn | 0     | 103.1.1.0   | recv | 577   | 1     | ge-0/0/3.0 |
| 103.1.1.1/32       | user | 0     |             | rjct | 539   | 2     |            |
| 103.1.1.1/32       | intf | 0     | 103.1.1.1   | loc1 | 578   | 2     |            |
| 103.1.1.1/32       | iddn | 0     | 103.1.1.1   | loc1 | 578   | 2     |            |
| 103.1.1.255/32     | iddn | 0     | 103.1.1.255 | bcst | 576   | 1     | ge-0/0/3.0 |
| 224.0.0.0/4        | perm | 0     |             | mdsc | 538   | 1     |            |
| 224.0.0.1/32       | perm | 0     | 224.0.0.1   | mcst | 534   | 1     |            |
| 255.255.255.255/32 | perm | 0     |             | bcst | 535   | 1     |            |

**Meaning** The output confirms that the virtual routing instances are created and the links are up and displays the routing table information.

**Related Documentation**

- [Configuring Virtual Routing Instances \(CLI Procedure\) on page 2347](#)
- [Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 2306](#)

## Verifying Integrated Routing and Bridging Interface Status and Statistics

**Purpose** Determine status information and traffic statistics for integrated routing and bridging (IRB) interfaces.

**Action** Display IRB interfaces and their current states:

```
user@switch> show interfaces irb terse
```

| Interface | Admin | Link | Proto | Local            | Remote |
|-----------|-------|------|-------|------------------|--------|
| irb       | up    | up   |       |                  |        |
| irb.111   | up    | up   | inet  | 111.111.111.1/24 |        |
| ...       |       |      |       |                  |        |

Display Layer 2 VLANs, including any tags assigned to the VLANs and the interfaces associated with the VLANs:

```
user@switch> show vlans
```

| Routing instance | VLAN name | Tag | Interfaces  |
|------------------|-----------|-----|-------------|
| default-switch   | irb       | 101 |             |
| default-switch   | support   | 111 |             |
|                  |           |     | ge-0/0/18.0 |
| ...              |           |     |             |

Display Ethernet switching table entries for the VLAN that is attached to the IRB interface:

```
user@switch> show ethernet-switching table
```

MAC flags (S -static MAC, D -dynamic MAC, L -locally learned  
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

```
Routing instance : default-switch
```

| Vlan Name | MAC address | MAC flags | Age | Logical interface |
|-----------|-------------|-----------|-----|-------------------|
|-----------|-------------|-----------|-----|-------------------|

```
support          00:01:02:03:04:05  S          - ge-0/0/18.0
...
```

Display the ingress-counting statistics of an IRB interface with either the **show interfaces irb detail** command or the **show interfaces irb extensive** command. Ingress counting is displayed as **Input bytes** and **Input packets** and egress counting is displayed as **Output bytes** and **Output packets** under **Transit Statistics**.

```
user@switch> show interfaces irb.111 detail
```

```
Logical interface irb.111 (Index 65) (SNMP ifIndex 503) (HW Token 100) (Generation
131)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Bandwidth: 1000mbps
Routing Instance: default-switch Bridging Domain: irb+111
Traffic statistics:
  Input bytes:    17516756
  Output bytes:   411764
  Input packets: 271745
  Output packets: 8256
Local statistics:
  Input bytes:    3240
  Output bytes:   411764
  Input packets:  54
  Output packets: 8256
Transit statistics:
  Input bytes:    17513516  0 bps
  Output bytes:   0        0 bps
  Input packets: 271745    0 pps
  Output packets: 0        0 pps
Protocol inet, MTU: 1514, Generation: 148, Route table: 0
Flags: None
Addresses, Flags: iS-Preferred Is-Primary
  Destination: 50.1.1/24, Local: 50.1.1.1, Broadcast: 50.1.1.255, Generation: 136
```

- Meaning**
- **show interfaces irb terse** displays a list of interfaces, including IRB interfaces, and their current states (up, down).
  - **show vlans** displays a list of VLANs, including any tags assigned to the VLANs and the interfaces associated with the VLANs.
  - **show ethernet-switching table** displays the Ethernet switching table entries, including VLANs attached to the IRB interface.
  - **show interfaces irb detail** displays IRB interface ingress counting as **Input Bytes** and **Input Packets** under **Transit Statistics**.

- Related Documentation**
- [Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\) on page 2340](#)

## Verifying That MVRP Is Working Correctly

### Purpose



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Verifying That MVRP Is Working Correctly*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

After configuring your EX Series switch to participate in MVRP, verify that the configuration is properly set and that MVRP messages are being sent and received on your switch.

### Action

1. Confirm that MVRP is enabled on your switch.

```
user@switch> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface      Join   Leave  LeaveAll
  xe-0/1/1       200   1000   10000
```

2. Confirm that MVRP messages are being sent and received on your switch.

```
user@switch> show mvrp statistics
MVRP statistics for routing instance 'default-switch'

Interface name           : xe-0/1/1
VLAN IDs registered      : 117
Sent MVRP PDUs           : 118824
Received MVRP PDUs without error: 118848
Received MVRP PDUs with error : 0
Transmitted Join Empty   : 5229
Transmitted Leave All    : 2
Received Join In         : 11884924
Transmitted Join In      : 1835
Transmitted Empty        : 93606408
Transmitted Leave        : 888
Transmitted In           : 13780024
Transmitted New          : 2692
Received Leave All       : 118761
Received Leave           : 97
Received In              : 3869
Received Empty           : 828
Received Join Empty      : 2020152
Received New             : 224
...
```

### Meaning

The output of **show mvrp** shows that interface xe-0/1/1 is enabled for MVRP participation.

The output for **show mvrp statistics** confirms that MVRP messages are being transmitted and received on interface xe-0/1/1.



**NOTE:** You can identify an MVRP compatibility issue by observing the output from this command. If **Received Join Empty** and **Received Join In** incorrectly display zero, even though the value for **Received MVRP PDUs without error** has been increased, you are probably running different versions of Junos OS on the switches in this network. Another indication that MVRP is having a version problem is that unexpected VLAN activity, such as multiple VLAN creation, takes place on the switch running the earlier release version. To remedy these problems, see [“Configuring Multiple VLAN Registration Protocol \(MVRP\) \(CLI Procedure\)”](#) on page 2348.

- Related Documentation**
- [Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches](#) on page 2310
  - [Understanding Multiple VLAN Registration Protocol \(MVRP\) on EX Series Switches](#) on page 2264

## Verifying That MAC Notification Is Working Properly

**Purpose** Verify that MAC notification is enabled or disabled, and that the MAC notification interval is set to the specified value.

**Action** Verify that MAC notification is enabled while also verifying the MAC notification interval setting.

```
user@switch> show ethernet-switching mac-notification
Notification Status: Enabled
Notification Interval: 30
```

**Meaning** The output in the **Notification Status** field shows that MAC notification is enabled. The output in the **Notification Status** field would display **Disabled** if MAC notification was disabled.

The **Notification Interval** field output shows that the MAC notification interval is set to 30 seconds.

- Related Documentation**
- [Configuring MAC Notification \(CLI Procedure\)](#)
  - [Configuring MAC Notification \(CLI Procedure\)](#) on page 2363

## Verifying That Proxy ARP Is Working Correctly

**Purpose** Verify that the switch is sending proxy ARP messages.

**Action** List the system statistics for ARP:

```
user@switch> show system statistics arp
arp:
    90060 datagrams received
```

```
34 ARP requests received
610 ARP replies received
2 resolution request received
0 unrestricted proxy requests
0 restricted proxy requests
0 received proxy requests
0 unrestricted proxy requests not proxied
0 restricted proxy requests not proxied
0 datagrams with bogus interface
0 datagrams with incorrect length
0 datagrams for non-IP protocol
0 datagrams with unsupported op code
0 datagrams with bad protocol address length
0 datagrams with bad hardware address length
0 datagrams with multicast source address
0 datagrams with multicast target address
0 datagrams with my own hardware address
0 datagrams for an address not on the interface
0 datagrams with a broadcast source address
294 datagrams with source address duplicate to mine
89113 datagrams which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
309 ARP requests sent
35 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor
```

**Meaning** The statistics show that two proxy ARP requests were received. The **unrestricted proxy requests not proxied** and **restricted proxy requests not proxied** fields indicate that all the unproxied ARP requests received have been proxied by the switch.

**Related Documentation**

- [Configuring Proxy ARP](#)
- [Configuring Proxy ARP \(CLI Procedure\) on page 2360](#)

---

## Operational Commands

- [show ethernet-switching interface](#)
- [show ethernet-switching mac-notification](#)
- [show ethernet-switching table](#)
- [show interfaces irb](#)
- [show mac-rewrite interface](#)
- [show mvrp](#)
- [show mvrp dynamic-vlan-memberships](#)



- `show mvrp statistics`
- `show redundant-trunk-group`
- `show route forwarding-table`
- `show system statistics arp`
- `show vlans`

## show ethernet-switching interface

|                                 |                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show ethernet-switching interface</b><br><b>&lt;brief   detail   extensive&gt;</b><br><b>&lt;interface-name&gt;</b>                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.3R2.<br>Command introduced in Junos OS Release 12.3R2 for EX Series switches.                                                                                                                                                          |
| <b>Description</b>              | (MX Series routers and EX Series switches only) Display Layer 2 learning information for all the interfaces.                                                                                                                                                                     |
| <b>Options</b>                  | <b>none</b> —Display Ethernet-switching information for all interfaces.<br><br><b>brief   detail   extensive</b> —(Optional) Display the specified level of output.<br><br><b>interface-name</b> —(Optional) Display Ethernet-switching information for the specified interface. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>    | <a href="#">show ethernet-switching interface ae10.0 on page 2443</a><br><a href="#">show ethernet-switching interface detail on page 2443</a>                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 244 on page 2442</a> describes the output fields for the <b>show ethernet-switching interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                             |

**Table 244: show ethernet-switching interface Output Fields**

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                        |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Logical interface</b>       | Name of the logical interface.                                                                                                                                                                                                                                                                                                           |
| <b>VLAN members</b>            | VLANs associated with this interface.                                                                                                                                                                                                                                                                                                    |
| <b>Tag</b>                     | VLAN ID.                                                                                                                                                                                                                                                                                                                                 |
| <b>MAC limit</b>               | Number of MAC addresses that can be associated with the interface.                                                                                                                                                                                                                                                                       |
| <b>STP state</b>               | Spanning tree protocol (STP) state.                                                                                                                                                                                                                                                                                                      |
| <b>Logical interface flags</b> | Status of Layer 2 learning properties for each interface: <ul style="list-style-type: none"> <li><b>DL</b>—MAC learning is disabled.</li> <li><b>LH</b>—MAC interface limit has been reached..</li> <li><b>AD</b>—Packets are dropped after the MAC interface limit is reached.</li> <li><b>DN</b>—The MAC interface is down.</li> </ul> |
| <b>Tagging</b>                 | Tagging state of the VLAN.                                                                                                                                                                                                                                                                                                               |

## Sample Output

### show ethernet-switching interface ae10.0

```

user@host> show ethernet-switching interface ae10.0
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members limit state   interface flags
ae10.0
      VLAN70.. 701   1024   Forwarding
      VLAN70.. 702   1024   Forwarding
      VLAN70.. 703   1024   Forwarding
      VLAN70.. 704   1024   Forwarding
      VLAN70.. 705   1024   Forwarding
      VLAN70.. 706   1024   Forwarding
      VLAN70.. 707   1024   Forwarding
      VLAN70.. 708   1024   Forwarding
      VLAN70.. 709   1024   Forwarding
      VLAN71.. 710   1024   Forwarding
      VLAN71.. 711   1024   Forwarding
      VLAN71.. 712   1024   Forwarding
      VLAN71.. 713   1024   Forwarding
      VLAN71.. 714   1024   Forwarding
      VLAN71.. 715   1024   Forwarding
[...output truncated...]

```

### show ethernet-switching interface detail

```

user@host> show ethernet-switching interface detail
Information for interface family:
Name: ge-1/0/3.0
  Type: IFF
  Index: 331
  IFD index: 141
  IFL index: 331
  Sequence number: 0
  MAC limit: 65535
  Static MACs learned: 0
Name: ge-1/0/3.0
  Type: IFBD (static)
  Index:
  Trunk id: 0
  IFD index:
  IFL index:
  Handle: 0x8bba280
  Generation: 159
  Flags: UP,
  Routing/Vlan index: 4
  Address family: 50
  MAC sequence number: 0
  MACs learned: 0
  Non configured static MACs learned: 0
  Handle: 0x8bb6e00
  Generation: 129
  Flags: UP,
  Routing/Vlan index: 2
  Address family:

```

Sequence number: 1  
MAC limit: 65535  
Static MACs learned: 0  
VSTP index: 11  
Name: ge-1/0/3.0  
Type: IFBD (static)  
Index:  
Trunk id: 0  
IFD index:  
IFL index:  
Sequence number: 1  
MAC limit: 65535  
Static MACs learned: 0  
VSTP index: 11

MAC sequence number: 1  
MACs learned: 0  
Non configured static MACs learned: 0  
Rewrite op:  
  
Handle: 0x8bb6f00  
Generation: 130  
Flags: UP,  
Routing/Vlan index: 3  
Address family:  
MAC sequence number: 1  
MACs learned: 0  
Non configured static MACs learned: 0  
Rewrite op:

## show ethernet-switching mac-notification

|                                 |                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show ethernet-switching mac-notification                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.6 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                 |
| <b>Description</b>              | Display information about MAC notification.                                                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Verifying That MAC Notification Is Working Properly on page 2439</a></li> </ul>                                                                  |
| <b>List of Sample Output</b>    | <a href="#">show ethernet-switching mac-notification (MAC Notification Enabled) on page 2445</a><br><a href="#">show ethernet-switching mac-notification (MAC Notification Disabled) on page 2445</a> |
| <b>Output Fields</b>            | <a href="#">Table 245 on page 2445</a> lists the output fields for the <b>show ethernet-switching mac-notification</b> command. Output fields are listed in the order in which they appear.           |

**Table 245: show ethernet-switching mac-notification Output Fields**

| Field Name                         | Field Description                                                                                                                                                                 |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Notification Status</b>         | MAC notification status: <ul style="list-style-type: none"> <li>• <b>Enabled</b>—MAC notification is enabled.</li> <li>• <b>Disabled</b>—MAC notification is disabled.</li> </ul> |
| <b>Notification Interval</b>       | MAC notification interval in seconds.                                                                                                                                             |
| <b>Notifications Sent</b>          | Number of notifications sent to SNMP when MACs are learned or when MACs age out.                                                                                                  |
| <b>Notifications Table Maxsize</b> | Maximum size of the notification table, which is populated when notifications are sent to the SNMP server.                                                                        |

### Sample Output

#### show ethernet-switching mac-notification (MAC Notification Enabled)

```

user@switch> show ethernet-switching mac-notification
Notification Status           : Enabled
Notification Interval         : 30
Notifications Sent            : 0
Notifications Table Maxsize   : 256

```

### Sample Output

#### show ethernet-switching mac-notification (MAC Notification Disabled)

```

user@switch> show ethernet-switching mac-notification
Notification Status           : Disabled
Notification Interval         : 0

```

Notifications Sent : 0  
Notifications Table Maxsize : 256

## show ethernet-switching table

**Syntax** `show ethernet-switching table`  
`<brief | count | detail | extensive | summary>`  
`<address>`  
`<instance instance-name>`  
`<interface interface-name>`  
`isid isid`  
`<logical-system logical-system-name>`  
`<persistent-learning (interface interface-name | mac mac-address)>`  
`<address>`  
`<vlan-id (all-vlan | vlan-id)>`  
`<vlan-name (all | vlan-name)>`

**Release Information** Command introduced in Junos OS Release 12.3R2.  
 Command introduced in Junos OS Release 12.3R2 for EX Series switches.  
 Options **logical-system**, **persistent-learning**, and **summary** introduced in Junos OS Release 13.2X50-D10 (ELS).

### Description



**NOTE:** If your EX Series switch CLI displays different options for the `show ethernet-switching table` command from the options shown in this document, then see [show ethernet-switching table](#).

(MX Series routers and EX Series switches only) Display Layer 2 MAC address information.

**Options** **none**—Display all learned Layer 2 MAC address information.

**brief | count | detail | extensive | summary**—(Optional) Display the specified level of output.

**address**—(Optional) Display the specified learned Layer 2 MAC address information.

**instance *instance-name***—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.

**interface *interface-name***—(Optional) Display learned Layer 2 MAC addresses for the specified interface.

**isid *isid***—(Optional) Display learned Layer 2 MAC addresses for the specified ISID.

**logical-system *logical-system-name***—(Optional) Display Ethernet-switching statistics information for the specified logical system.

**persistent-learning (interface *interface-name* | mac *mac-address*)**—(Optional) Display dynamically learned MAC addresses that are retained despite device restarts and interface failures for a specified interface, or information about a specified MAC address.

**vlan-id (all-vlan | *vlan-id*)**—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

**vlan-name (all | *vlan-name*)**—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

**Additional Information** When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.

**Required Privilege Level** view

**List of Sample Output** [show ethernet-switching table on page 2449](#)  
[show ethernet-switching table brief on page 2450](#)  
[show ethernet-switching table count on page 2451](#)  
[show ethernet-switching table extensive on page 2452](#)

**Output Fields** [Table 228 on page 2239](#) describes the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

**Table 246: show ethernet-switching table Output fields**

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Routing instance</b>   | Name of the routing instance.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>VLAN name</b>          | Name of the VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>MAC address</b>        | MAC address or addresses learned on a logical interface.                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>MAC flags</b>          | Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> <li>• <b>S</b>—Static MAC address is configured.</li> <li>• <b>D</b>—Dynamic MAC address is configured.</li> <li>• <b>L</b>—Locally learned MAC address is configured.</li> <li>• <b>SE</b>—MAC accounting is enabled.</li> <li>• <b>NM</b>—Non-configured MAC.</li> <li>• <b>R</b>—Locally learned MAC address is configured.</li> </ul> |
| <b>Age</b>                | This field is not supported.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Logical interface</b>  | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>MAC count</b>          | Number of MAC addresses learned on the specific routing instance or interface.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Learning interface</b> | Name of the logical interface on which the MAC address was learned.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Learning VLAN</b>      | VLAN ID of the routing instance or VLAN in which the MAC address was learned.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Layer 2 flags</b>      | Debugging flags signifying that the MAC address is present in various lists.                                                                                                                                                                                                                                                                                                                                                               |



Table 246: show ethernet-switching table Output fields (*continued*)

| Field Name      | Field Description                                                                                     |
|-----------------|-------------------------------------------------------------------------------------------------------|
| Epoch           | Spanning-tree-protocol epoch number identifying when the MAC address was learned. Used for debugging. |
| Sequence number | Sequence number assigned to this MAC address. Used for debugging.                                     |
| Learning mask   | Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.         |
| IPC generation  | Creation time of the logical interface when this MAC address was learned. Used for debugging.         |

## Sample Output

### show ethernet-switching table

```

user@host> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags    interface
  VLAN101   88:e0:f3:bb:07:f0  D        -        ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags    interface
  VLAN102   88:e0:f3:bb:07:f0  D        -        ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags    interface
  VLAN103   88:e0:f3:bb:07:f0  D        -        ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags    interface
  VLAN104   88:e0:f3:bb:07:f0  D        -        ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

```

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN1101     | 00:1f:12:32:f5:c1 | D            | -   | ae0.0                |

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN1102     | 00:1f:12:32:f5:c1 | D            | -   | ae0.0                |

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN1103     | 00:1f:12:32:f5:c1 | D            | -   | ae0.0                |

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN1104     | 00:1f:12:32:f5:c1 | D            | -   | ae0.0                |

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN1105     | 00:1f:12:32:f5:c1 | D            | -   | ae0.0                |

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN1106     | 00:1f:12:32:f5:c1 | D            | -   | ae0.0                |

[...output truncated...]

### show ethernet-switching table brief

user@host> show ethernet-switching table brief

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

| Vlan<br>name | MAC<br>address | MAC<br>flags | Age | Logical<br>interface |
|--------------|----------------|--------------|-----|----------------------|
|--------------|----------------|--------------|-----|----------------------|

```
VLAN101          88:e0:f3:bb:07:f0  D          -   ae20.0
```

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN102      | 88:e0:f3:bb:07:f0 | D            | -   | ae20.0               |

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN103      | 88:e0:f3:bb:07:f0 | D            | -   | ae20.0               |

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN104      | 88:e0:f3:bb:07:f0 | D            | -   | ae20.0               |

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned  
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN1101     | 00:1f:12:32:f5:c1 | D            | -   | ae0.0                |

[...output truncated...]

### show ethernet-switching table count

```
user@host> show ethernet-switching table count
```

```
0 MAC address learned in routing instance default-switch VLAN VLAN1000
ae26.0:1000
```

```
1 MAC address learned in routing instance default-switch VLAN VLAN101
ae20.0:101
```

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count | Static MAC count |
|---------------|-----------|------------------|
| 101           | 1         | 0                |

```
1 MAC address learned in routing instance default-switch VLAN VLAN102
ae20.0:102
```

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count | Static MAC count |
|---------------|-----------|------------------|
| 102           | 1         | 0                |

```
1 MAC address learned in routing instance default-switch VLAN VLAN103
ae20.0:103
```

```
MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count      Static MAC count
        103             1             0

1 MAC address learned in routing instance default-switch VLAN VLAN104
ae20.0:104

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count      Static MAC count
        104             1             0

0 MAC address learned in routing instance default-switch VLAN VLAN105
ae20.0:105

0 MAC address learned in routing instance default-switch VLAN VLAN106
ae20.0:106

0 MAC address learned in routing instance default-switch VLAN VLAN107
ae20.0:107

0 MAC address learned in routing instance default-switch VLAN VLAN108
ae20.0:108

0 MAC address learned in routing instance default-switch VLAN VLAN109
ae20.0:109

0 MAC address learned in routing instance default-switch VLAN VLAN110
ae20.0:110

1 MAC address learned in routing instance default-switch VLAN VLAN1101
ae0.0:1101

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count      Static MAC count
        1101             1             0

1 MAC address learned in routing instance default-switch VLAN VLAN1102
ae0.0:1102

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count      Static MAC count
        1102             1             0
[...output truncated...]
```

#### show ethernet-switching table extensive

```
user@host> show ethernet-switching table extensive

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 101
Learning interface: ae20.0
Layer 2 flags: in_hash,in_ifid,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                      Sequence number: 2
Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 102
Learning interface: ae20.0
```

```
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 103
Learning interface: ae20.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 104
Learning interface: ae20.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1101
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1102
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1103
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1104
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008
```

## show interfaces irb

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces irb &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance <i>instance-name</i>&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 12.3R2.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for the QFX Series</p>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Display integrated routing and bridging interfaces information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance <i>instance-name</i></b>—(Optional) Display information for the interface with the specified SNMP index.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the interface with the specified SNMP index.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Additional Information</b>   | Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another VLAN that has a Layer 3 protocol configured.                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces irb extensive on page 2458</a></p> <p><a href="#">show interfaces irb snmp-index on page 2459</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | <a href="#">Table 247 on page 2454</a> lists the output fields for the <b>show interfaces irb</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Table 247: show interfaces irb Output Fields**

| Field Name                | Field Description                                                                                                                             | Level of Output |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Physical Interface</b> |                                                                                                                                               |                 |
| <b>Physical interface</b> | Name of the physical interface.                                                                                                               | All levels      |
| <b>Enabled</b>            | State of the physical interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels      |

Table 247: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                          | Level of Output                    |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| <b>Proto</b>                   | Protocol configured on the interface.                                                                                                                                                                                                                      | <b>terse</b>                       |
| <b>Interface index</b>         | Physical interface index number, which reflects its initialization sequence.                                                                                                                                                                               | <b>detail extensive none</b>       |
| <b>SNMP ifIndex</b>            | SNMP index number for the physical interface.                                                                                                                                                                                                              | <b>detail extensive none</b>       |
| <b>Type</b>                    | Physical interface type.                                                                                                                                                                                                                                   | <b>detail extensive none</b>       |
| <b>Link-level type</b>         | Encapsulation being used on the physical interface.                                                                                                                                                                                                        | <b>detail extensive brief none</b> |
| <b>MTU</b>                     | MTU size on the physical interface.                                                                                                                                                                                                                        | <b>detail extensive brief none</b> |
| <b>Clocking</b>                | Reference clock source: <b>Internal</b> or <b>External</b> . Always unspecified on IRB interfaces.                                                                                                                                                         | <b>detail extensive brief</b>      |
| <b>Speed</b>                   | Speed at which the interface is running. Always unspecified on IRB interfaces.                                                                                                                                                                             | <b>detail extensive brief</b>      |
| <b>Device flags</b>            | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .                                                                                                         | <b>detail extensive brief none</b> |
| <b>Interface flags</b>         | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .                                                                                                            | <b>detail extensive brief none</b> |
| <b>Link type</b>               | Physical interface link type: <b>full duplex</b> or <b>half duplex</b> .                                                                                                                                                                                   | <b>detail extensive none</b>       |
| <b>Link flags</b>              | Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .                                                                                                                     | <b>detail extensive none</b>       |
| <b>Physical Info</b>           | Physical interface information.                                                                                                                                                                                                                            | All levels                         |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                        | <b>detail extensive</b>            |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                    | <b>detail extensive none</b>       |
| <b>Hardware address</b>        | MAC address of the hardware.                                                                                                                                                                                                                               | <b>detail extensive none</b>       |
| <b>Alternate link address</b>  | Backup address of the link.                                                                                                                                                                                                                                | <b>detail extensive</b>            |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> . | <b>detail extensive none</b>       |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                          | <b>detail extensive</b>            |

Table 247: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output         |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> |
| <b>IPv6 transit statistics</b> | <p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Input errors</b>            | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>           | <b>detail extensive</b> |
| <b>Output errors</b>           | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the DPC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>detail extensive</b> |

#### Logical Interface



Table 247: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output                 |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Logical interface</b>       | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels                      |
| <b>Index</b>                   | Index number of the logical interface (which reflects its initialization sequence).                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>            | SNMP interface index number of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b><br>none |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive</b>         |
| <b>Flags</b>                   | Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b>         |
| <b>Encapsulation</b>           | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>         |
| <b>Bandwidth</b>               | Speed at which the interface is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b>         |
| <b>Routing Instance</b>        | Routing instance IRB is configured under.                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>         |
| <b>Bridging Domain</b>         | Bridging domain IRB is participating in.                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b>         |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the logical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                             | <b>detail extensive</b>         |
| <b>IPv6 transit statistics</b> | <p>Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b>         |
| <b>Local statistics</b>        | Statistics for traffic received from and transmitted to the Routing Engine.                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>         |
| <b>Transit statistics</b>      | Statistics for traffic transiting the router.                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b>         |
| <b>Protocol</b>                | Protocol family configured on the local interface. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>         |
| <b>MTU</b>                     | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b>         |
| <b>Maximum labels</b>          | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b><br>none |

Table 247: show interfaces irb Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                               | Level of Output         |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                                                               | <b>detail extensive</b> |
| <b>Route table</b>      | Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.                                             | <b>detail extensive</b> |
| <b>Addresses, Flags</b> | Information about address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .                 | <b>detail extensive</b> |
| <b>Policer</b>          | The policer that is to be evaluated when packets are received or transmitted on the interface.                                                                  | <b>detail extensive</b> |
| <b>Flags</b>            | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | <b>detail extensive</b> |

## Sample Output

### show interfaces irb extensive

```

user@host> show interfaces irb extensive
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0
    Output packets: 0
  IPv6 transit statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0
    Output packets: 0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Giants: 0, Policed discards:
0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0

```

```

Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Protocol inet, MTU: 1500, Generation: 154, Route table: 0
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255,
    Generation: 155
Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
  Flags: Is-Primary
  Policer: Input: __default_arp_policer

```

### show interfaces irb snmp-index

```

user@host> show interfaces irb snmp-index 25
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 25
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags : Present Running
  Interface flags: SNMP-Traps
  Link type : Full-Duplex
  Link flags : None
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Last flapped : Never
    Input packets : 0
    Output packets: 0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1500
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
  Protocol multiservice, MTU: 1500
    Flags: Is-Primary

```

## show mac-rewrite interface

|                                 |                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show mac-rewrite interface</b><br><b>&lt;brief   detail&gt;</b><br><b>&lt;interface-name&gt;</b>                                                                                                                                                               |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.1.<br>Command introduced in Junos OS Release 14.1X53-D10 for EX Series switches.                                                                                                                                         |
| <b>Description</b>              | Display Layer 2 protocol tunneling information.                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>interface interface-name</b> —(Optional) Display Layer 2 protocol tunneling information for the specified interface.                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">layer2-control on page 2399</a></li> <li>• <a href="#">mac-rewrite on page 2404</a></li> <li>• <a href="#">Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches on page 2322</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show mac-rewrite interface on page 2460</a><br><a href="#">show mac-rewrite interface (EX Series Switch) on page 2461</a>                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 248 on page 2460</a> lists the output fields for the <b>show mac-rewrite interface</b> command.<br>Output fields are listed in the approximate order in which they appear.                                                                      |

**Table 248: show mac-rewrite interface Output Fields**

| Field Name       | Field Description                                                                                                                                                                                                                                                                                                     | Level of Output     |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| <b>Interface</b> | Name of the interface that has Layer 2 protocol tunneling configured on it.                                                                                                                                                                                                                                           | <b>brief detail</b> |
| <b>Protocols</b> | Layer 2 protocols being tunneled on this interface: Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), Per-VLAN Spanning Tree Plus (PVSTP+), or VLAN Trunk Protocol (VTP)<br><br>On EX Series switches, the following Layer 2 protocols are supported: 802.3A, CDP, LACP, LLDP, MVRP, STP, VTP, GVRP, VSTP | <b>brief detail</b> |

## Sample Output

### show mac-rewrite interface

```

user@host> show mac-rewrite interface
Interface          Protocols
-----
ge-1/0/5           STP VTP CDP PVSTP+

```

**show mac-rewrite interface (EX Series Switch)**

```
user@switch> show mac-rewrite interface
Interface      Protocols
-----
ge-0/0/1      STP
```

## show mvrp

|                                 |                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show mvrp                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.1.<br>Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.                                                                                                                                                                  |
| <b>Description</b>              | Display Multiple VLAN Registration Protocol (MVRP) configuration information.                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>show mvrp applicant-state</i></li> <li>• <a href="#">show mvrp dynamic-vlan-memberships on page 2464</a></li> <li>• <i>show mvrp interface</i></li> <li>• <i>show mvrp registration-state</i></li> <li>• <i>show mvrp statistics</i></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show mvrp on page 2462</a>                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | <a href="#">Table 249 on page 2462</a> lists the output fields for the <b>show mvrp</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                    |

**Table 249: show mvrp Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MVRP dynamic VLAN creation | Displays whether global MVRP dynamic Virtual LAN (VLAN) creation is <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| MVRP BPDU MAC address      | Displays the multicast media access control (MAC) address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the customer MVRP multicast MAC address is used.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| MVRP timers (ms)           | Displays MVRP timer information: <ul style="list-style-type: none"> <li>• <b>Interface</b>—The interface on which MVRP is configured.</li> <li>• <b>Join</b>—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements.</li> <li>• <b>Leave</b>—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message.</li> <li>• <b>LeaveAll</b>—The interval at which LeaveAll messages are sent on interfaces. LeaveAll messages maintain current MVRP VLAN membership information in the network.</li> </ul> |

## Sample Output

### show mvrp

```
user@host> show mvrp
```

```
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface      Join   Leave  LeaveAll
  ge-11/2/8      200    800    10000
  ge-11/0/9      200    800    10000
  ge-11/3/0      200    800    10000
```

## show mvrp dynamic-vlan-memberships

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show mvrp dynamic-vlan-memberships                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.1.<br>Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.                                                                                                                                         |
| <b>Description</b>              | Display all Virtual LANs (VLANs) that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the router or switch.                                                                                                                      |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show mvrp on page 2462</a></li> <li>• <i>show mvrp applicant-state</i></li> <li>• <i>show mvrp interface</i></li> <li>• <i>show mvrp registration-state</i></li> <li>• <i>show mvrp statistics</i></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show mvrp dynamic-vlan-memberships on page 2464</a>                                                                                                                                                                                                    |
| <b>Output Fields</b>            | <a href="#">Table 250 on page 2464</a> lists the output fields for the <b>show mvrp dynamic-vlan-memberships</b> command. Output fields are listed in the approximate order in which they appear.                                                                  |

**Table 250: show mvrp dynamic-vlan-memberships Output Fields**

| Field Name | Field Description                                                           |
|------------|-----------------------------------------------------------------------------|
| VLAN Id    | The VLAN ID of the dynamically created VLAN.                                |
| Interfaces | The interface or interfaces that are bound to the dynamically created VLAN. |

### Sample Output

#### show mvrp dynamic-vlan-memberships

```

user@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
(s) static vlan, (f) fixed registration

VLAN Id      Interfaces
  100 (s)    ge-11/3/0
  200 (s)    ge-11/3/0
  300 (s)

```



## show mvrp statistics

|                                 |                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show mvrp statistics<br><interface <i>interface-name</i> ><br><routing-instance <i>routing-instance-name</i> >                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X50-D10 (ELS).                                                                                                                           |
| <b>Description</b>              | Display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages.                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show mvrp on page 2462</a></li> <li>• <a href="#">Verifying That MVRP Is Working Correctly on page 2438</a></li> </ul>         |
| <b>List of Sample Output</b>    | <a href="#">show mvrp statistics on page 2466</a>                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 251 on page 2465</a> lists the output fields for the <b>show mvrp statistics</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 251: show mvrp statistics Output Fields**

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface name                   | Interface for which MVRP statistics are displayed.                                                                                                                                                                                                                                                                                                                                                                                                   |
| VLAN IDs registered              | Number of Virtual LAN (VLAN) IDs registered.                                                                                                                                                                                                                                                                                                                                                                                                         |
| Sent MVRP PDUs                   | Number of MRPDU messages transmitted from the switch.                                                                                                                                                                                                                                                                                                                                                                                                |
| Received MVRP PDUs without error | Number of MRPDU messages received on the switch.                                                                                                                                                                                                                                                                                                                                                                                                     |
| Received MVRP PDUs with error    | Number of invalid MRPDU messages received on the switch.                                                                                                                                                                                                                                                                                                                                                                                             |
| Transmitted Join Empty           | Number of JoinEmpty messages sent from the switch.                                                                                                                                                                                                                                                                                                                                                                                                   |
| Transmitted Leave All            | Number of MRP LeaveAll messages sent from the switch.                                                                                                                                                                                                                                                                                                                                                                                                |
| Received Join In                 | Number of MRP JoinIn messages received on the switch. Either this value or the value for <b>Received Join Empty</b> should increase when the value for <b>Received MVRP PDUs without error</b> increases. If this value is not incrementing when it should, you might have a Junos OS release compatibility issue. To resolve the issue, see <a href="#">“Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)” on page 2348</a> . |
| Transmitted Join In              | Number of MRP JoinIn messages sent from the switch.                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 251: show mvrp statistics Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Transmitted Empty</b>   | Number of MRP Empty messages sent from the switch.                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Transmitted Leave</b>   | Number of MRP LeaveEmpty messages sent from the switch.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Transmitted In</b>      | Number of MRP In messages sent from the switch.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Transmitted New</b>     | Number of New messages transmitted from the switch.                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Received Leave All</b>  | Number of LeaveAll messages received on the switch.                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Received Leave</b>      | Number of MRP Leave messages received on the switch.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Received In</b>         | Number of MRP In messages received on the switch.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Received Empty</b>      | Number of MRP Empty messages received on the switch.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Received Join Empty</b> | Number of MRP JoinEmpty messages received on the switch. Either this value or the value for <b>Received Join In</b> should increase when the value for <b>Received MVRP PDUs without error</b> increases. If this value is not incrementing when it should, you might have a Junos OS release compatibility issue. To resolve the issue, see <a href="#">“Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)”</a> on page 2348. |
| <b>Received New</b>        | Number of New messages received on the switch.                                                                                                                                                                                                                                                                                                                                                                                                      |

## Sample Output

### show mvrp statistics

```

user@host> show mvrp statistics
MVRP statistics for routing instance 'default-switch'

Interface name           : xe-0/1/1
VLAN IDs registered      : 117
Sent MVRP PDUs           : 118824
Received MVRP PDUs without error: 118848
Received MVRP PDUs with error : 0
Transmitted Join Empty   : 5229
Transmitted Leave All    : 2
Received Join In         : 11884924
Transmitted Join In      : 1835
Transmitted Empty        : 93606408
Transmitted Leave        : 888
Transmitted In           : 13780024
Transmitted New          : 2692
Received Leave All       : 118761
Received Leave           : 97
Received In              : 3869
Received Empty           : 828
Received Join Empty      : 2020152
Received New             : 224
...

```

## show redundant-trunk-group

|                                 |                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show redundant-trunk-group &lt;group-name group-name&gt;</code>                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.                                                                                                                                                                          |
| <b>Description</b>              | Display information about redundant trunk groups.                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <code>group-name group-name</code> —Display information about the specified redundant trunk group.                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery</a></li> <li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2326</a></li> <li>• <a href="#">Understanding Redundant Trunk Links on page 2276</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show redundant-trunk-group group-name Group1 on page 2467</a>                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 252 on page 2467</a> lists the output fields for the <code>show redundant-trunk-group</code> command. Output fields are listed in the approximate order in which they appear.                                                                                                                       |

Table 252: show redundant-trunk-group Output Fields

| Field Name        | Field Description                                                                                                                                                                                                                                                                        |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Group name        | Name of the redundant trunk port group.                                                                                                                                                                                                                                                  |
| Interface         | Name of an interface belonging to the trunk port group.                                                                                                                                                                                                                                  |
| State             | Operating state of the interface. <ul style="list-style-type: none"> <li>• <b>Up</b> denotes the interface is up.</li> <li>• <b>Down</b> denotes the interface is down.</li> <li>• <b>Pri</b> denotes a primary interface.</li> <li>• <b>Act</b> denotes an active interface.</li> </ul> |
| Time of last flap | Date and time at which the advertised link became unavailable, and then, available again.                                                                                                                                                                                                |
| Flap count        | Total number of flaps since the last switch reboot.                                                                                                                                                                                                                                      |

## Sample Output

### show redundant-trunk-group group-name Group1

```
user@switch> show redundant-trunk-group group-name Group1
```

| Group name | Interface | State | Time of last flap | Flap Count |
|------------|-----------|-------|-------------------|------------|
|------------|-----------|-------|-------------------|------------|

|        |             |            |       |   |
|--------|-------------|------------|-------|---|
| Group1 | ge-0/0/45.0 | UP/Pri/Act | Never | 0 |
|        | ge-0/0/47.0 | UP         | Never | 0 |

## show route forwarding-table

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;ccc interface-name&gt; &lt;destination&gt; &lt;family family   matching matching&gt; &lt;label name&gt; &lt;multicast&gt; &lt;table (default   routing-table-name)&gt; &lt;vpn vpn&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in JUNOS Release 9.2 for EX-series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Display the forwarding table of the routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>none</b>—Display the routes in the forwarding table.</p> <p><b>detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>ccc interface-name</b>—(Optional) Display route entries for the specified circuit cross-connect interface.</p> <p><b>destination</b>—(Optional) Destination prefix.</p> <p><b>family family</b>—(Optional) Display routing table entries for the specified family: <b>inet</b>, <b>inet6</b>, <b>iso</b>, <b>mpls</b>, <b>tnp</b>, <b>unix</b>, or <b>vpls</b>.</p> <p><b>label name</b>—(Optional) Display route entries for the specified label.</p> <p><b>matching matching</b>—(Optional) Display routing table entries matching the specified prefix or prefix length.</p> <p><b>multicast</b>—(Optional) Display routing table entries for multicast routes.</p> <p><b>table (default   routing-table-name)</b>—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing table.</p> <p><b>vpn vpn</b>—(Optional) Display routing table entries for a specified VPN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches on page 2306</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>    | <a href="#">show route forwarding-table on page 2471</a><br><a href="#">show route forwarding-table detail on page 2471</a><br><a href="#">show route forwarding-table extensive on page 2471</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | <p><a href="#">Table 253 on page 2470</a> lists the output fields for the <b>show route forwarding-table</b> command. Output fields are listed in the approximate order in which they appear. Field names may</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 253: show route forwarding-table Output Fields

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output         |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Routing table</b>           | Name of the routing table (for example, <b>inet</b> , <b>inet6</b> , <b>mpls</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels              |
| <b>Destination</b>             | Destination of the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Route Type (Type)</b>       | How the route was placed into the forwarding table. When the <b>detail</b> keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> <li>• <b>cloned (clon)</b>—(TCP or multicast only) Cloned route.</li> <li>• <b>destination (dest)</b>—Remote addresses directly reachable through an interface.</li> <li>• <b>destination down (iddn)</b>—Destination route for which the interface is unreachable.</li> <li>• <b>interface cloned (ifcl)</b>—Cloned route for which the interface is unreachable.</li> <li>• <b>route down (ifdn)</b>—Interface route for which the interface is unreachable.</li> <li>• <b>ignore (ignr)</b>—Ignore this route.</li> <li>• <b>interface (intf)</b>—Installed as a result of configuring an interface.</li> <li>• <b>permanent (perm)</b>—Routes installed by the kernel when the routing table is initialized.</li> <li>• <b>user</b>—Routes installed by the routing protocol process or as a result of the configuration.</li> </ul>                                                                                                                           | All levels              |
| <b>Route Reference (RtRef)</b> | Number of routes to reference.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Next hop</b>                | IP address of the next hop to the destination.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Next hop Type (Type)</b>    | Next-hop type. When the <b>detail</b> keyword is used, the next-hop type might be abbreviated (as indicated in parentheses): <ul style="list-style-type: none"> <li>• <b>broadcast (bcst)</b> —Broadcast.</li> <li>• <b>deny</b>—Deny.</li> <li>• <b>hold</b>—Next hop is waiting to be resolved into a unicast or multicast type.</li> <li>• <b>indexed (idxd)</b>—Indexed next hop.</li> <li>• <b>indirect (indr)</b>—Indirect next hop.</li> <li>• <b>local (locl)</b>—Local address on an interface.</li> <li>• <b>routed multicast (mcr)</b>—Regular multicast next hop</li> <li>• <b>multicast (mcst)</b>—Wire multicast next hop (limited to the LAN).</li> <li>• <b>multicast discard (mdsc)</b>—Multicast discard.</li> <li>• <b>multicast group (mgrp)</b> —Multicast group member.</li> <li>• <b>receive (rcv)</b>—Receive.</li> <li>• <b>reject (rjct)</b> Discard. An ICMP unreachable message was sent.</li> <li>• <b>resolve (rslv)</b>—Resolving the next hop.</li> <li>• <b>unicast (ucst)</b>—Unicast.</li> <li>• <b>unilist (ulst)</b>—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.</li> </ul> | <b>detail extensive</b> |

Table 253: show route forwarding-table Output Fields (*continued*)

| Field Name                        | Field Description                                                                    | Level of Output              |
|-----------------------------------|--------------------------------------------------------------------------------------|------------------------------|
| <b>Index</b>                      | Software index of the next hop that is used to route the traffic for a given prefix. | <b>detail extensive none</b> |
| <b>Reference (NhRef)</b>          | Number of routes that refer to this next hop.                                        | <b>detail extensive none</b> |
| <b>Next-hop interface (Netif)</b> | Interface used to reach the next hop.                                                | <b>detail extensive none</b> |

## Sample Output

### show route forwarding-table

```

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          user   1 0:12:f2:21:cf:0 ucst  331   4 me0.0
default          perm   0                      rjct   36    1
0.0.0.0/32       perm   0                      dscd   34    1
10.93.54.0/24    intf   0                      rslv  323    1 me0.0
10.93.54.0/32    dest   0 10.93.54.0        recv  321    1 me0.0
10.93.54.1/32    intf   0 10.93.54.1        locl  322    2
10.93.54.1/32    dest   0 10.93.54.1        locl  322    2
10.93.54.254/32  dest   0 0:12:f2:21:cf:0   ucst  331    4 me0.0
10.93.54.255/32  dest   0 10.93.54.255      bcst  320    1 me0.0
172.23.5.193/32 clon    1 10.93.54.254      ucst  331    4 me0.0
224.0.0.0/4      perm   0                      mdsc   35    1
224.0.0.1/32     perm   0 224.0.0.1         mcst   31    1
255.255.255.255/32 perm   0                      bcst   32    1

```

### show route forwarding-table detail

```

user@host> show route forwarding-table detail
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          user   1 0:12:f2:21:cf:0 ucst  331   4 me0.0
default          perm   0                      rjct   36    1
0.0.0.0/32       perm   0                      dscd   34    1
10.93.54.0/24    intf   0                      rslv  323    1 me0.0
10.93.54.0/32    dest   0 10.93.54.0        recv  321    1 me0.0
10.93.54.1/32    intf   0 10.93.54.1        locl  322    2
10.93.54.1/32    dest   0 10.93.54.1        locl  322    2
10.93.54.254/32  dest   0 0:12:f2:21:cf:0   ucst  331    4 me0.0
10.93.54.255/32  dest   0 10.93.54.255      bcst  320    1 me0.0
172.23.5.193/32 clon    1 10.93.54.254      ucst  331    4 me0.0
224.0.0.0/4      perm   0                      mdsc   35    1
224.0.0.1/32     perm   0 224.0.0.1         mcst   31    1
255.255.255.255/32 perm   0                      bcst   32    1

```

### show route forwarding-table extensive

```

user@host> show route forwarding-table extensive
Routing table: default.inet [Index 0]
Internet:

```

```
Destination: default
Route type: user
Route reference: 1          Route interface-index: 0
Flags: sent to PFE, rt nh decoupled
Nexthop: 0:12:f2:21:cf:0
Next-hop type: unicast      Index: 331      Reference: 4
Next-hop interface: me0.0
```



## show system statistics arp

|                                 |                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show system statistics arp                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.6 for EX Series switches.                                                                                                                                                           |
| <b>Description</b>              | Display system-wide Address Resolution Protocol (ARP) statistics.                                                                                                                                                            |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Proxy ARP on an EX Series Switch on page 2331</a></li> <li>• <a href="#">Verifying That Proxy ARP Is Working Correctly on page 2439</a></li> </ul> |

## show system statistics arp

```

user@switch> show system statistics arp
arp:
    90060 datagrams received
    34 ARP requests received
    610 ARP replies received
    0 resolution request received
    0 unrestricted proxy requests
    0 restricted proxy requests
    0 received proxy requests
    0 unrestricted proxy requests not proxied
    0 restricted proxy requests not proxied
    0 datagrams with bogus interface
    0 datagrams with incorrect length
    0 datagrams for non-IP protocol
    0 datagrams with unsupported op code
    0 datagrams with bad protocol address length
    0 datagrams with bad hardware address length
    0 datagrams with multicast source address
    0 datagrams with multicast target address
    0 datagrams with my own hardware address
    0 datagrams for an address not on the interface
    0 datagrams with a broadcast source address
    294 datagrams with source address duplicate to mine
    89113 datagrams which were not for me
    0 packets discarded waiting for resolution
    0 packets sent after waiting for resolution
    309 ARP requests sent
    35 ARP replies sent
    0 requests for memory denied
    0 requests dropped on entry
    0 requests dropped during retry
    0 requests dropped due to interface deletion
    0 requests on unnumbered interfaces
    0 new requests on unnumbered interfaces
    0 replies for from unnumbered interfaces
    0 requests on unnumbered interface with non-subnetted donor
    0 replies from unnumbered interface with non-subnetted donor

```

## show vlans

**Syntax** `show vlans`  
`<brief | detail | extensive>`  
`<instance instance-name>`  
`<logical-system logical-system-name>`  
`<operational>`  
`<vlan-name>`  
`<interface interface-name>`

**Release Information** Command introduced in Junos OS Release 12.3R2.  
 Command introduced in Junos OS Release 12.3R2 for EX Series switches.  
 Option **interface** introduced in Junos OS Release 13.2X50-D10 (ELS).

### Description



**NOTE:** If your EX Series switch CLI displays different options for the `show vlans` command from the options shown in this document, see *show vlans*.

(MX Series routers and EX Series switches only) Display VLAN information.

### Options

**none**—Display information for all VLANs.

**brief | detail | extensive**—(Optional) Display the specified level of output.

**instance *instance-name***—(Optional) Display information for the specified routing instance.

**logical-system *logical-system-name***—(Optional) Display Ethernet-switching statistics information for the specified logical system.

**operational**—(Optional) Display information for the operational routing instances.

***vlan-name***— (Optional) Display information about the specified VLAN.

**interface *interface-name***—(Optional) Display information about the specified interface.

### Required Privilege Level

view

### List of Sample Output

[show vlans brief \(EX Series Switch\) on page 2474](#)  
[show vlans brief on page 2475](#)  
[show vlans detail \(EX Series Switch\) on page 2475](#)  
[show vlans detail on page 2477](#)  
[show vlans extensive \(EX Series Switch\) on page 2477](#)  
[show vlans extensive on page 2479](#)

## Sample Output

### show vlans brief (EX Series Switch)

```
user@switch> show vlans brief
Routing instance  VLAN name  Tag  Interfaces
default-switch  c1          20
```

| Routing instance | VLAN name | Tag | Interfaces                                |
|------------------|-----------|-----|-------------------------------------------|
| default-switch   | c2        | 30  | ge-0/0/0.0*<br>ge-1/0/0.0*<br>ge-2/0/0.0* |
| default-switch   | default   | 1   | ge-0/0/0.0*<br>ge-2/0/0.0*                |
| default-switch   | iso       | 10  | ge-0/0/1.0*                               |
| default-switch   | iso1      | 50  | ge-0/0/0.0*<br>ge-2/0/0.0*                |
| default-switch   | pri       | 100 | ge-0/0/0.0*<br>ge-1/0/0.0*<br>ge-2/0/0.0* |

### show vlans brief

```

user@host> show vlans brief
Routing instance  VLAN name  Tag  Interfaces
VPLS-1           __VPLS-1__  all  ae1.0
VPLS-2           __VPLS-2__  all  ae3.0
                                   ge-3/1/2.0
                                   vt-3/3/10.1048576
default-switch   VLAN1000    1000 ae26.0
default-switch   VLAN101     101  ae20.0
default-switch   VLAN102     102  ae20.0
default-switch   VLAN103     103  ae20.0
default-switch   VLAN104     104  ae20.0
default-switch   VLAN105     105  ae20.0
default-switch   VLAN106     106  ae20.0
default-switch   VLAN107     107  ae20.0
default-switch   VLAN108     108  ae20.0
[...output truncated...]

```

### show vlans detail (EX Series Switch)

```

user@switch> show vlans detail
Routing instance: default-switch
  VLAN Name: c1                               State: Active
  Tag: 20
  PVLAN type : Community
  Internal index: 16, Generation Index: 21, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk

```

```
ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3      , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: c2                                State: Active
  Tag: 30
  PVLAN type : Community
  Internal index: 17, Generation Index: 22, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
  Number of interfaces: Tagged 2      , Untagged 0
  Total MAC count: 0

Routing instance: default-switch
  VLAN Name: default                          State: Active
  Tag: 1
  Internal index: 5, Generation Index: 5, Origin: Static
  MAC aging time: 300 seconds
  Number of interfaces: Tagged 0      , Untagged 0
  Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso                              State: Active
  Tag: 10
  Internal index: 14, Generation Index: 19, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/1.0*,untagged,access
  Number of interfaces: Tagged 0      , Untagged 1
  Total MAC count: 0

Routing instance: default-switch
  VLAN Name: iso1                             State: Active
  Tag: 50
  PVLAN type : Isolated
  Internal index: 15, Generation Index: 20, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
  Number of interfaces: Tagged 2      , Untagged 0
  Total MAC count: 0

Routing instance: default-switch
  VLAN Name: pri                              State: Active
  Tag: 100
  PVLAN type : Primary
  Isolated VLAN :
  vlan-id : 50 vlan name : iso1
  Community VLAN :
  vlan-id : 20 vlan name : c1
  vlan-id : 30 vlan name : c2
  Internal index: 9, Generation Index: 14, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
```

```

Number of interfaces: Tagged 3      , Untagged 0
Total MAC count: 0

```

### show vlans detail

```

user@host> show vlans detail
Routing instance: VPLS-1
  VLAN Name: __VPLS-1__                State: Active
  Tag: all
  Internal index: 2, Generation Index:  , Origin: Dynamic
  Interfaces:
    ae1.0,tagged
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 0

Routing instance: VPLS-2
  VLAN Name: __VPLS-2__                State: Active
  Tag: all
  Internal index: 3, Generation Index:  , Origin: Dynamic
  Interfaces:
    ae3.0,tagged
    ge-3/1/2.0,tagged
    vt-3/3/10.1048576,tagged
  Number of interfaces: Tagged 3      , Untagged 0
  Total MAC count: 4

Routing instance: default-switch
  VLAN Name: VLAN1000                  State: Active
  Tag: 1000
  Internal index: 4, Generation Index: 1, Origin: Static
  Layer 3 interface: irb.1000
  Interfaces:
    ae26.0,tagged,trunk
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 0

Routing instance: default-switch
  VLAN Name: VLAN101                  State: Active
  Tag: 101
  Internal index: 5, Generation Index: 2, Origin: Static
  Layer 3 interface: irb.101
  Interfaces:
    ae20.0,tagged,trunk
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 1

Routing instance: default-switch
  VLAN Name: VLAN102                  State: Active
  Tag: 102
  Internal index: 6, Generation Index: 3, Origin: Static
  Layer 3 interface: irb.102
  Interfaces:
    ae20.0,tagged,trunk
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 1
[...output truncated...]

```

### show vlans extensive (EX Series Switch)

```

user@switch> show vlans extensive

```

Routing instance: default-switch  
VLAN Name: c1 State: Active  
Tag: 20  
PVLAN type : Community  
Internal index: 16, Generation Index: 21, Origin: Static  
MAC aging time: 300 seconds  
Interfaces:  
ge-0/0/0.0\*,tagged,trunk  
ge-1/0/0.0\*,tagged,trunk  
ge-2/0/0.0\*,tagged,trunk  
Number of interfaces: Tagged 3 , Untagged 0  
Total MAC count: 0

Routing instance: default-switch  
VLAN Name: c2 State: Active  
Tag: 30  
PVLAN type : Community  
Internal index: 17, Generation Index: 22, Origin: Static  
MAC aging time: 300 seconds  
Interfaces:  
ge-0/0/0.0\*,tagged,trunk  
ge-2/0/0.0\*,tagged,trunk  
Number of interfaces: Tagged 2 , Untagged 0  
Total MAC count: 0

Routing instance: default-switch  
VLAN Name: default State: Active  
Tag: 1  
Internal index: 5, Generation Index: 5, Origin: Static  
MAC aging time: 300 seconds  
Number of interfaces: Tagged 0 , Untagged 0  
Total MAC count: 0

Routing instance: default-switch  
VLAN Name: iso State: Active  
Tag: 10  
Internal index: 14, Generation Index: 19, Origin: Static  
MAC aging time: 300 seconds  
Interfaces:  
ge-0/0/1.0\*,untagged,access  
Number of interfaces: Tagged 0 , Untagged 1  
Total MAC count: 0

Routing instance: default-switch  
VLAN Name: iso1 State: Active  
Tag: 50  
PVLAN type : Isolated  
Internal index: 15, Generation Index: 20, Origin: Static  
MAC aging time: 300 seconds  
Interfaces:  
ge-0/0/0.0\*,tagged,trunk  
ge-2/0/0.0\*,tagged,trunk  
Number of interfaces: Tagged 2 , Untagged 0  
Total MAC count: 0

Routing instance: default-switch  
VLAN Name: pri State: Active  
Tag: 100  
PVLAN type : Primary  
Isolated VLAN :  
vlan-id : 50 vlan name : iso1

```

Community VLAN :
vlan-id : 20 vlan name : c1
vlan-id : 30 vlan name : c2
Internal index: 9, Generation Index: 14, Origin: Static
MAC aging time: 300 seconds
Interfaces:
    ge-0/0/0.0*,tagged,trunk
    ge-1/0/0.0*,tagged,trunk
    ge-2/0/0.0*,tagged,trunk
Number of interfaces: Tagged 3      , Untagged 0
Total MAC count: 0

```

### show vlans extensive

```

user@host> show vlans extensive
Routing instance: default-switch
  VLAN Name: VLAN_10                                State: Active
  Tag: 10
  Internal index: 2, Generation Index: 1, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-1/0/3.0*,tagged,trunk
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 0

Routing instance: default-switch
  VLAN Name: VLAN_20                                State: Active
  Tag: 20
  Internal index: 3, Generation Index: 2, Origin: Static
  MAC aging time: 300 seconds
  Interfaces:
    ge-1/0/3.0*,tagged,trunk
  Number of interfaces: Tagged 1      , Untagged 0
  Total MAC count: 0

```





## PART 13

# High Availability

- [Overview on page 2483](#)
- [Configuration on page 2505](#)
- [Administration on page 2545](#)
- [Troubleshooting Procedures on page 2573](#)



## CHAPTER 40

# Overview

- [High Availability Features Overview on page 2483](#)
- [Nonstop Active Routing Overview on page 2486](#)
- [Nonstop Bridging Overview on page 2487](#)
- [Nonstop Software Upgrade Overview on page 2488](#)
- [Power Management Overview on page 2495](#)
- [VRRP Overview on page 2501](#)

### High Availability Features Overview

---

- [High Availability Features for EX4300 Switches Overview on page 2483](#)

### High Availability Features for EX4300 Switches Overview

*High availability* refers to the hardware and software components that provide redundancy and reliability for network communications. This topic covers the following high availability features of Juniper Networks EX4300 Ethernet Switches:

- [Redundant Routing Engines on page 2483](#)
- [Virtual Chassis on page 2484](#)
- [VRRP on page 2484](#)
- [Graceful Routing Engine Switchover on page 2485](#)
- [Link Aggregation on page 2485](#)
- [Nonstop Active Routing and Nonstop Bridging on page 2485](#)

### Redundant Routing Engines

---

Redundant Routing Engines are two Routing Engines that are installed in a switch or a Virtual Chassis. When a switch has two Routing Engines, one functions as the master, while the other stands by as a backup in case the master Routing Engine fails. When a Virtual Chassis has two Routing Engines, the switch in the master role functions as the master Routing Engine and the switch in the backup role functions as the backup Routing Engine. Redundant Routing Engines are supported on Juniper Networks EX4300 Ethernet Switches configuring into a Virtual Chassis.

The master Routing Engine receives and transmits routing information, builds and maintains routing tables, communicates with interfaces and Packet Forwarding Engine components of the switch, and has full control over the control plane of the switch.

The backup Routing Engine stays in sync with the master Routing Engine in terms of protocol states, forwarding tables, and so forth. If the master becomes unavailable, the backup Routing Engine takes over the functions that the master Routing Engine performs.

Network reconvergence takes place more quickly on switches and on Virtual Chassis with redundant Routing Engines than on switches and on Virtual Chassis with a single Routing Engine.

### Virtual Chassis

---

A Virtual Chassis is multiple switches connected together that operate as a single network entity. The advantages of connecting multiple switches into a Virtual Chassis include better-managed bandwidth at a network layer, simplified configuration and maintenance because multiple devices can be managed as a single device, a simplified Layer 2 network topology that minimizes or eliminates the need for loop prevention protocols such as Spanning Tree Protocol (STP), and improved fault tolerance and high availability. A Virtual Chassis improves high availability for the following reasons:

- **Dual Routing Engine support.** A Virtual Chassis automatically has two Routing Engines—the switches in the master and backup **routing-engine** roles—and, therefore, provides more high availability options than standalone switches. Many high availability features are available for an EX Series Virtual Chassis that are not available on standalone EX Series switches.
- **Increased fault tolerance.** You increase your fault tolerance options when you configure your EX Series switches into a Virtual Chassis. You can, for instance, configure interfaces into a link aggregation group (LAG) with member interfaces on different member switches in the same Virtual Chassis to ensure network traffic is received by a Virtual Chassis even when a switch or physical interface in the Virtual Chassis fails.

You can configure up to ten EX4300 switches into an EX4300 Virtual Chassis. See [“Understanding EX4300 Virtual Chassis” on page 5065](#).

### VRRP

---

You can configure Virtual Router Redundancy Protocol (VRRP) for IP and IPv6 on most switch interfaces, including Gigabit Ethernet interfaces, high-speed uplink interfaces, and logical interfaces. When VRRP is configured, the switches act as virtual routing platforms. VRRP enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master routing platform fails, one of the backup routing platforms becomes the new master, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup switch can take over a failed default switch within a few seconds. This is done with minimum loss of VRRP traffic and without any interaction with the hosts.

---

### Graceful Routing Engine Switchover

---

You can configure graceful Routing Engine switchover (GRES) on a switch with redundant Routing Engines or on a Virtual Chassis, allowing control to switch from the master Routing Engine to the backup Routing Engine with minimal interruption to network communications. When you configure GRES, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and forwarding state. Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur. If the kernel on the master Routing Engine stops operating, the master Routing Engine experiences a hardware failure, or the administrator initiates a manual switchover, mastership switches to the backup Routing Engine.

When the backup Routing Engine assumes mastership in a redundant failover configuration (that is, when GRES is not enabled), the Packet Forwarding Engines initialize their state to the boot state before they connect to the new master Routing Engine. In contrast, in a GRES configuration, the Packet Forwarding Engines do not reinitialize their state, but resynchronize their state to that of the new master Routing Engine. The interruption to traffic is minimal.

---

### Link Aggregation

---

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a link aggregation group (LAG) or bundle. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation provides network redundancy by load-balancing traffic across all available links. If one of the links should fail, the system automatically load-balances traffic across all remaining links. In a Virtual Chassis, LAGs can be used to load-balance network traffic between member switches, which increases high availability by ensuring that network traffic is received by the Virtual Chassis even if a single interface fails for any reason.

The number of Ethernet interfaces you can include in a LAG and the number of LAGs you can configure on a switch depend on the switch model. For information about LAGs, see [“Understanding Aggregated Ethernet Interfaces and LACP” on page 2582](#).

---

### Nonstop Active Routing and Nonstop Bridging

---

Nonstop active routing (NSR) provides high availability in a switch or Virtual Chassis with redundant Routing Engines by enabling transparent switchover of the Routing Engines without requiring restart of supported Layer 3 routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbor routing devices, which do not detect that a change has occurred. The neighboring devices and other devices on the network do not, therefore, have to resynchronize their Layer 3 protocol states to respond to the downtime on the switch—a process that adds network overhead and risks disrupting network performance—when a Routing Engine switchover occurs when NSR is enabled.

Nonstop bridging (NSB) provides the same mechanism for Layer 2 protocols. NSB operates by synchronizing all protocol information for NSB-supported Layer 2 protocols between the master and backup Routing Engines. If the switch has a Routing Engine switchover, the NSB-supported Layer 2 protocol sessions remain active because they are already synchronized on the backup Routing Engine. The Routing Engine switchover

is transparent to neighbor devices, which do not detect any changes related to the Layer 2 protocol sessions. The neighboring devices and other devices on the network do not, therefore, have to resynchronize their Layer 2 protocol states to respond to the downtime on the switch—a process that adds network overhead and risks disrupting network performance—when a Routing Engine switchover occurs when NSB is enabled.

To use NSR or NSB, you must also configure GRES.

**Related Documentation**

- For more information about high availability features, see the [Junos OS High Availability Configuration Guide](#).
- [EX Series Virtual Chassis Overview](#)
- [Understanding EX4300 Virtual Chassis on page 5065](#)
- [Understanding VRRP on EX Series Switches on page 2501](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2582](#)

---

## Nonstop Active Routing Overview

- [Understanding Nonstop Active Routing on EX Series Switches on page 2486](#)

### Understanding Nonstop Active Routing on EX Series Switches

You can configure nonstop active routing (NSR) on an EX Series switch with redundant Routing Engines or on an EX Series Virtual Chassis to enable the transparent switchover of the Routing Engines in the event that one of the Routing Engines goes down.

Nonstop active routing provides high availability for Routing Engines by enabling transparent switchover of the Routing Engines without requiring restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbor routing devices, which do not detect that a change has occurred.

Enable nonstop active routing when neighbor routing devices are not configured to support graceful restart of protocols or when you want to ensure graceful restart of protocols for which graceful restart is not supported—such as PIM.

You do not need to start the two Routing Engines simultaneously to synchronize them for nonstop active routing. If both Routing Engines are not present or not up when you issue a **commit synchronize** statement, the candidate configuration is committed in the master Routing Engine and when the backup Routing Engine is inserted or comes online, its configuration is automatically synchronized with that of the master.

Nonstop active routing uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop active routing also saves routing protocol information by running the routing protocol process (**rpd**) on the backup Routing Engine. By saving this additional information, nonstop active routing does not rely on other routing devices to assist in restoring routing protocol information.



**NOTE:** After a graceful Routing Engine switchover, we recommend that you issue the `clear interface statistics (interface-name | all)` command to reset the cumulative values for local statistics on the new master Routing Engine.

If you suspect a problem with the synchronization of Routing Engines when nonstop active routing is enabled, you can gather troubleshooting information using trace options. For example, if certain protocols lose connectivity with neighbors after a graceful Routing Engine switchover with NSR enabled, you can use trace options to help isolate the problem. See [“Tracing Nonstop Active Routing Synchronization Events” on page 2573](#).



**NOTE:** Graceful restart and nonstop active routing are mutually exclusive. You will receive an error message upon commit if both are configured.



**NOTE:** Nonstop active routing provides a transparent switchover mechanism only for Layer 3 protocol sessions. Nonstop bridging (NSB) provides a similar mechanism for Layer 2 protocol sessions. See [“Understanding Nonstop Bridging on EX Series Switches” on page 2487](#).

#### Related Documentation

- [Configuring Nonstop Active Routing on Switches on page 2505](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 2514](#)

## Nonstop Bridging Overview

- [Understanding Nonstop Bridging on EX Series Switches on page 2487](#)

### Understanding Nonstop Bridging on EX Series Switches

You can configure nonstop bridging (NSB) to provide resilience for Layer 2 protocol sessions on a Juniper Networks EX Series Ethernet Switch or on an EX Series Virtual Chassis with redundant Routing Engines.

NSB operates by synchronizing all protocol information for NSB-supported Layer 2 protocols between the master and backup Routing Engines. If the switch has a Routing Engine switchover, the NSB-supported Layer 2 protocol sessions remain active because all session information is already synchronized to the backup Routing Engine. Traffic disruption for the NSB-supported Layer 2 protocol is minimal or nonexistent as a result of the switchover. The Routing Engine switchover is transparent to neighbor devices, which do not detect any changes related to the NSB-supported Layer 2 protocol sessions on the switch.

For a list of the EX Series switches and Layer 2 protocols that support NSB, see *EX Series Switch Software Features Overview* and *EX Series Virtual Chassis Software Features Overview*.



**NOTE:** Nonstop bridging provides a transparent switchover mechanism only for Layer 2 protocol sessions. Nonstop active routing (NSR) provides a similar mechanism for Layer 3 protocol sessions. See [“Understanding Nonstop Active Routing on EX Series Switches” on page 2486](#).

**Related Documentation**

- For information about configuring NSB on EX Series switches that do not support the Enhanced Layer 2 Software (ELS) CLI style, see [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\)](#)
- For information about configuring NSB on EX Series switches that support ELS, see [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 2507](#)

---

## Nonstop Software Upgrade Overview

- [Understanding Nonstop Software Upgrade on EX Series Switches on page 2488](#)

### Understanding Nonstop Software Upgrade on EX Series Switches

Nonstop software upgrade (NSSU) enables you to upgrade the software running on Juniper Networks EX Series Ethernet Switches with redundant Routing Engines and all member switches in EX Series Virtual Chassis by using a single command and with minimal network traffic disruption during the upgrade.

NSSU is supported on the following platforms:

- EX3300 Virtual Chassis
- EX4200 Virtual Chassis
- EX4300 Virtual Chassis
- EX4500 Virtual Chassis
- EX4550 Virtual Chassis
- All mixed Virtual Chassis composed of EX4200, EX4500, and EX4550 switches
- EX6200 switches
- EX8200 switches
- EX8200 Virtual Chassis

Performing an NSSU provides these benefits:

- No disruption to the control plane—An NSSU takes advantage of graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) to ensure no disruption to the control plane. During the upgrade process, interface, kernel, and routing protocol information is preserved.
- Minimal disruption to network traffic—An NSSU minimizes network traffic disruption by:



- Upgrading line cards one at a time in an EX6200 switch, EX8200 switch, or EX8200 Virtual Chassis, permitting traffic to continue to flow through the line cards that are not being upgraded.
- Upgrading member switches one at a time in an EX3300, EX4200, EX4300, EX4500, or mixed Virtual Chassis, permitting traffic to continue to flow through the members that are not being upgraded.

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



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

For EX6200 switches, EX8200 switches, and EX8200 Virtual Chassis, you can reduce the amount of time an upgrade takes by configuring line-card upgrade groups. The line cards in an upgrade group are upgraded simultaneously, reducing the amount of time it takes to complete an upgrade. See *Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure)*.

This topic covers:

- [Requirements for Performing an NSSU on page 2489](#)
- [How an NSSU Works on page 2490](#)
- [NSSU Limitations on page 2493](#)
- [NSSU and Junos OS Release Support on page 2493](#)
- [Overview of NSSU Configuration and Operation on page 2494](#)

### **Requirements for Performing an NSSU**

The following requirements apply to all switches and Virtual Chassis:

- All Virtual Chassis members and all Routing Engines must be running the same Junos OS release.
- Graceful Routing Engine switchover (GRES) must be enabled.
- Nonstop active routing (NSR) must be enabled.



**NOTE:** Although nonstop bridging (NSB) does not have to be enabled to perform an NSSU, we recommend enabling NSB before performing an NSSU. Enabling NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU. See *Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)*.

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

The following are requirements for EX3300, EX4200, EX4300, EX4500, and mixed Virtual Chassis:

- The Virtual Chassis members must be connected in a ring topology so that no member is isolated as a result of another member being rebooted. This topology prevents the Virtual Chassis from splitting during an NSSU.
- The Virtual Chassis master and backup must be adjacent to each other in the ring topology. Adjacency permits the master and backup to always be in sync, even when the switches in linecard roles are rebooting.
- The Virtual Chassis must be preprovisioned so that the linecard role has been explicitly assigned to member switches acting in a linecard role. During an NSSU, the Virtual Chassis members must maintain their roles—the master and backup must maintain their master and backup roles (although mastership will change), and the remaining switches must maintain their linecard roles.
- A two-member Virtual Chassis must have **no-split-detection** configured so that the Virtual Chassis does not split when an NSSU upgrades a member.



**NOTE:** For the EX4300 Virtual Chassis, you should enable the `vcp-no-hold-time` statement at the `[edit virtual-chassis]` hierarchy level before performing a software upgrade using NSSU. If you do not enable the `vcp-no-hold-time` statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see [“Understanding Split and Merge in a Virtual Chassis” on page 5087](#)

---

## How an NSSU Works

This section describes what happens when you request an NSSU on these switches and Virtual Chassis:

- [EX3300, EX4200, EX4300, EX4500, and Mixed Virtual Chassis on page 2491](#)
- [EX6200 and EX8200 Switches on page 2491](#)
- [EX8200 Virtual Chassis on page 2492](#)

***EX3300, EX4200, EX4300, EX4500, and Mixed Virtual Chassis***

When you request an NSSU on an EX3300, EX4200, EX4300, EX4500, or mixed Virtual Chassis:

1. The Virtual Chassis master verifies that:
  - The backup is online and running the same software version.
  - Graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) are enabled.
  - The Virtual Chassis has a preprovisioned configuration.
2. The master installs the new software image on the backup and reboots it.
3. The master resynchronizes the backup.
4. The master installs the new software image on member switches that are in the linecard role and reboots them, one at a time. The master waits for each member to become online and active before starting the software upgrade on the next member.
5. When all members that are in the linecard role have been upgraded, the master performs a graceful Routing Engine switchover, and the upgraded backup becomes the master.
6. The software on the original master is upgraded and the original master is automatically rebooted. After the original master has rejoined the Virtual Chassis, you can optionally return control to it by requesting a graceful Routing Engine switchover.

***EX6200 and EX8200 Switches***

When you request an NSSU on a standalone switch with redundant Routing Engines:

1. The switch verifies that:
  - Both Routing Engines are online and running the same software version.
  - Both Routing Engines have sufficient storage space for the new software image.
  - Graceful Routing Engine switchover and nonstop active routing are enabled.
2. The switch installs the new software image on the backup Routing Engine and reboots it.
3. The switch resynchronizes the backup Routing Engine to the master Routing Engine.
4. The line cards in the first upgrade group (or the line card in slot 0, if no upgrade groups are defined) download the new image and then restart. Traffic continues to flow through the line cards in the other upgrade groups during this process.
5. When line cards restarted in Step 4 are online again, the line cards in the next upgrade group download the new image and restart. This process continues until all online line cards have restarted with the new software.



**NOTE:** If you have taken a line card offline with the CLI before you start the NSSU, the line card is not restarted and remains offline.

6. The switch performs a graceful Routing Engine switchover, so that the upgraded backup Routing Engine becomes the master.
7. The switch installs the new software on the original master Routing Engine.

To complete the upgrade process, the original master Routing Engine must be rebooted. You can do so manually or have the switch perform an automatic reboot by including the **reboot** option when you request the NSSU. After the original master has been rebooted, you can optionally return control to it by requesting a graceful Routing Engine switchover.

8. (EX6200 switch only) The original master Routing Engine reboots to complete the software upgrade.



**NOTE:** To complete the upgrade process on an EX8200 switch, you must intervene to reboot the original master Routing Engine. You can reboot the original master Routing Engine manually or have the switch perform an automatic reboot by including the **reboot** option when you request the NSSU.

9. (Optional) After the original master has been rebooted, you can return control to it by requesting a graceful Routing Engine switchover.

The switch can maintain normal operations with either Routing Engine acting as the master Routing Engine after the software upgrade, so you only have to perform this switchover if you want to return Routing Engine control to the original master Routing Engine.

### ***EX8200 Virtual Chassis***

When you request an NSSU on an EX8200 Virtual Chassis:

1. The master external Routing Engine verifies that:
  - It has a backup external Routing Engine that is online.
  - All Virtual Chassis members have redundant Routing Engines and the Routing Engines are online.
  - All Routing Engines are running the same software version.
  - All Routing Engines have sufficient storage space for the new software image.
  - Graceful Routing Engine switchover and nonstop active routing (NSR) are enabled.
2. The master external Routing Engine installs the new software image on the backup external Routing Engine and reboots it.
3. The backup external Routing Engine resynchronizes with the master external Routing Engine.

4. The master external Routing Engine installs the new software on the backup Routing Engines in the member switches and reboots the backup Routing Engines.
5. When the reboot of the backup Routing Engines complete, the line cards in the first upgrade group download the new image and then restart. (If no upgrade groups are defined, the line card in slot 0 of member 0 downloads the new image and restarts.) Traffic continues to flow through the line cards in the other upgrade groups during this process.
6. When line cards restarted in Step 5 are online again, the line cards in the next upgrade group (or the next sequential line card) download the new image and restart. This process continues until all online line cards have restarted with the new software.



**NOTE:** If you have taken a line card offline with the CLI before you start the NSSU, the line card is not restarted and remains offline.

7. The new software image is installed on the master Routing Engines, both external and internal.
8. The member switches perform a graceful Routing Engine switchover, so that the upgraded backup Routing Engines become masters.
9. The master external Routing Engine performs a graceful Routing Engine switchover so that the backup external Routing Engine is now the master.

To complete the upgrade process, the original master Routing Engines, both external and internal, must be rebooted. You can do so manually by establishing a console connection to each Routing Engine or have the reboot performed automatically by including the **reboot** option when you request the NSSU. After the original master external Routing Engine has been rebooted, you can optionally return control to it by requesting a graceful Routing Engine switchover.

### NSSU Limitations

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

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

### NSSU and Junos OS Release Support

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

Table 254 on page 2494 lists the EX Series switches and Virtual Chassis that support NSSU and the Junos OS release at which they began supporting it.

Table 254: Platform and Release Support for NSSU

| Platform                                         | Junos OS Release     |
|--------------------------------------------------|----------------------|
| EX3300 Virtual Chassis                           | 12.2 or later        |
| EX4200 Virtual Chassis                           | 12.1 or later        |
| EX4300 Virtual Chassis                           | 13.2X51-D20 or later |
| EX4500 Virtual Chassis                           | 12.1 or later        |
| EX4550 Virtual Chassis                           | 12.2 or later        |
| Mixed EX4200 and EX4500 Virtual Chassis          | 12.1 or later        |
| Mixed EX4200 and EX4550 Virtual Chassis          | 12.2 or later        |
| Mixed EX4200, EX4500, and EX4550 Virtual Chassis | 12.2 or later        |
| Mixed EX4500 and EX4550 Virtual Chassis          | 12.2 or later        |
| EX6200 switch                                    | 12.2 or later        |
| EX8200 switch                                    | 10.4 or later        |
| EX8200 Virtual Chassis                           | 11.1 or later        |

### Overview of NSSU Configuration and Operation

You must ensure that the configuration of the switch or Virtual Chassis meets the requirements described in [“Requirements for Performing an NSSU” on page 2489](#). NSSU requires no additional configuration.

For EX6200 switches, EX8200 switches, and EX8200 Virtual Chassis, you can optionally configure line-card upgrade groups using the CLI. See *Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches*.

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

#### Related Documentation

- [Upgrading Software on an EX3300, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and Mixed Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 2508](#)
- [Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade \(CLI Procedure\)](#)
- [Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\)](#)

- [Configuring Nonstop Active Routing on Switches on page 2505](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 5116](#)
- [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches](#)

---

## Power Management Overview

---

- [Understanding Power Management on EX Series Switches on page 2495](#)

### Understanding Power Management on EX Series Switches

The power management feature for Juniper Networks Ethernet Switches helps ensure that normal operation of the system is not disrupted because of insufficient power to the switch. For example:

- Power management ensures that operating line cards continue to receive power if a user installs a new line card in an operating switch when power is insufficient for both the new and existing line cards.
- Power management reserves a certain amount of power to power supply redundancy, so that if a power supply fails, the switch can continue to operate normally. If power management must use some of this reserved power to provide power to switch components, it raises an alarm to indicate that power supply redundancy no longer exists and that normal operations might be disrupted if a power supply fails.
- If power supply failure requires power management to power down some components, it does so gracefully by powering down line cards and PoE ports in the order specified by the user.

Power management manages power to switch components by employing a power budget policy. In its power budget policy, power management:

- Budgets power for each installed switch component that requires power. With the exception of PoE power for line cards that support PoE, the amount that power management budgets for each component is the maximum power that component might consume under worst case operating conditions. For example, for the fan tray, power management budgets the amount of power required to run the fans at their maximum speed setting, even if the current fan speed is much lower.
- Reserves a set amount of power for power supply redundancy. In its default configuration, power management manages the switch for  $N+1$  power redundancy, which ensures uninterrupted system operation if one power supply fails. For example, if a switch has four online 3000 W power supplies, power management reserves 3000 W in its power budget policy for redundancy. It allocates the remaining 9000 W to normal operating power.
- Specifies the rules under which components receive power. These rules are designed to ensure the least disruption to switch operation under conditions of insufficient power.

For example, power management provides power to core system components, such as the Routing Engines, before it provides power to line cards.

You can configure certain aspects of power management's budget policy, specifically:

- The power priority of individual line cards. By assigning different power priorities to the line cards, you can determine which line cards are more likely to receive power in the event of insufficient power.
- The power redundancy configuration. The default power redundancy configuration is  $N+1$ ; you can optionally configure  $N+N$ . For example, if you have deployed two independent AC power feeds to the switch, configure  $N+N$  redundancy. When you configure power management for  $N+N$  redundancy, it reserves the appropriate amount of power in its power budget and reports insufficient power conditions accordingly.

These configurable items are discussed further in:

- [Power Priority of Line Cards on page 2496](#)
- [Power Supply Redundancy on page 2499](#)

---

### Power Priority of Line Cards

---

The power priority of line cards determines:

- The order in which line cards are allocated power
- The order in which line cards that support PoE are allocated power for PoE
- How power is reallocated in cases of changes in power availability or demand in an operating switch



**NOTE:** On EX6200 switches, the four 10-Gigabit Ethernet SFP+ uplink ports on a Switch Fabric and Routing Engine (SRE) module are treated like a line card in the power budget.

---

This section covers:

- [How a Line Card's Power Priority Is Determined on page 2496](#)
- [Line Card Priority and Line Card Power on page 2497](#)
- [Line Card Priority and PoE Power on page 2497](#)
- [Line Card Priority and Changes in the Power Budget on page 2498](#)

#### ***How a Line Card's Power Priority Is Determined***

Using the CLI, you can assign an explicit power priority to a line-card slot. If more than one slot has the same assigned priority, the power priority is determined by slot number, with the lowest-numbered slots receiving power first.

By default, all slots in an EX8200 switch are assigned the lowest priority. Thus if you do not explicitly assign priorities to slots, power priority is determined by slot number, with slot 0 having the highest priority.



In an EX6200 switch, all slots are assigned the lowest priority, except for the slots containing an SRE module. Slots containing an SRE module are automatically assigned the highest priority. This means that the line cards that represent the 10-Gigabit Ethernet SFP+ ports on SRE modules have the highest priority among the line cards.

### ***Line Card Priority and Line Card Power***

When an EX6200 or EX8200 switch is powered on, power management allocates power to components according to its power budget policy. After power management has allocated power to the base chassis components, it allocates the remaining available power to the line cards. It powers on the line cards in priority order until all line cards are powered on or the available power (including reserved power, if necessary) is exhausted. Thus if available power is exhausted before all line cards receive power, higher-priority cards are powered on while lower-priority cards remain powered off.

A lower-priority card might receive power while a higher-priority card does not if the remaining available power is sufficient to power on the lower-priority card but not the higher-priority card. For example, if a line card requiring 450 W is in a higher-priority slot than line card requiring 330 W, the line card requiring 330 W receives the power if there is less than 450 W but more than 330 W remaining in the power budget.

Line cards that have been administratively taken offline are not allocated power.



**NOTE:** Because power management does not allocate power to a line card that has been administratively taken offline, a line card that has been taken offline in an EX6200 or EX8200 switch is not automatically brought online when you commit a configuration. You must explicitly use the `request chassis fpc slot slot-number online` command to bring a line card online that was taken offline previously. This behavior differs from other platforms running Juniper Networks Junos operating system (Junos OS), which automatically bring an offline FPC online when you commit a configuration.

If power management cannot power on a line card because of insufficient power, it raises a major (red) alarm.

### ***Line Card Priority and PoE Power***

After all line cards have been powered on, power management allocates any remaining available power, including reserved power, to the PoE power budgets of line cards that have PoE ports. Power management allocates PoE power to line cards in the order of power priority. If enough power is available, a line card receives its full PoE power budget before power management allocates PoE power to the next highest-priority line card. If not enough power is available, a line card receives partial PoE power and lower-priority line cards receive no PoE power.

If power management is unable to allocate enough power to meet the PoE power budget for a line card, it logs a message to the system log.

The default PoE power budget for a line card is the amount of power needed to supply the maximum supported power to all PoE ports. In cases where powered devices do not require the maximum power or in which some PoE ports are not used for powered devices,

you can configure a smaller PoE power budget for a line card. By configuring a smaller PoE power budget, you make more power available for the PoE power budgets of lower-priority line cards.

You can also configure the power priority of the PoE ports on a line card. If power management is unable to allocate enough power to a line card to meet its PoE power budget, the line card PoE controller will turn off power to PoE ports in reverse priority order as required to meet the reduced power allocation.

See [“Configuring PoE on EX Series Switches \(CLI Procedure\)”](#) on page 4440 for more information on how to configure the PoE power budget for a line card and how to configure PoE port priorities.

### ***Line Card Priority and Changes in the Power Budget***

In an operating switch, power management dynamically reallocates power in response to changes in power availability or demand or changes in line card priority. Power management uses line card priority to determine how to reallocate power in response to the following events:

- A power supply fails, is removed, or is taken offline:
  - If power is insufficient to meet the PoE power allocations of all PoE line cards, power management deallocates PoE power from the line cards in reverse priority order until power is sufficient to meet the remaining PoE power allocations.
  - If power is insufficient to meet the base (non-PoE) power requirements of all the line cards, all PoE power is deallocated. If, after the deallocation of PoE power, power is still not sufficient, power management turns off line cards in reverse priority order until power is sufficient for the remaining line cards.
- A new line card is inserted or a line card is brought online:
  - If the line card supports PoE and there is insufficient power to meet its PoE power budget, PoE power is reallocated from lower-priority line cards. If not enough PoE power can be reallocated from lower-priority line cards, the new line card receives a partial PoE power allocation.
  - If there is insufficient power to power on the new line card, PoE power is removed from PoE line cards in reverse priority order until the new line card can be powered on.
  - If the removal of all PoE power is insufficient to free up enough power to power on the line card, the line card remains powered off and the PoE line cards continue to receive their PoE power allocations. To minimize disruption on an operating switch, lower-priority line cards are not turned off to provide power to the new line card. However, if you restart the switch, power management reruns the current power budget policy and powers line cards on or off based on their priority. As a result, line cards receive power strictly by priority order and previously operating line cards might no longer receive power.
- A new power supply is brought online:

- Any line cards that were powered off because of insufficient power are powered on in priority order.
- After all line cards are powered on, remaining power is allocated to the PoE power budgets of line cards in priority order.
- A line card is removed or taken offline, freeing up power:
  - Any line cards that were powered down because of insufficient power are powered on in priority order.
  - After all line cards are powered on, any remaining power is allocated to the PoE power budgets of line cards in priority order.
- A user changes the assigned power priority of one or more line cards when power is insufficient to meet the power budget:
  - PoE power to the line cards is reallocated based on the new power priorities.
  - Base power allocation to the line cards is not changed—in other words, power management does not power down line cards that had been receiving power because they are now a lower priority. However, if you restart the switch, power management reruns the current power budget policy and powers line cards on or off based on their priority. As a result, line cards receive power strictly by priority order and previously operating line cards might no longer receive power.

If, because of insufficient power, power management reduces or eliminates the PoE power budget for a line card, it logs a message to the system log. If power management must power down a line card because of insufficient power, it raises a major (red) alarm.

### Power Supply Redundancy

By default, power management in EX Series switches is configured to manage the power supplies for  $N+1$  redundancy, in which one power supply is held in reserve for backup if one of the other power supplies is removed or fails.

You can configure power management to manage the power supplies for  $N+N$  redundancy. In  $N+N$  redundancy, power management holds  $N$  power supplies in reserve for backup. For example, if your switch has six power supplies and you configure  $N+N$  redundancy, power management makes three power supplies available for normal operating power and reserves three power supplies for redundancy (3+3). If you have an odd number of power supplies, power management allocates one more power supply to normal operating power than to redundant power. For example, if you have five power supplies, the  $N+N$  configuration is 3+2.

Given the same number of power supplies, an  $N+N$  configuration usually provides less normal operating power than an  $N+1$  configuration because the  $N+N$  configuration holds more power in reserve for backup. [Table 255 on page 2500](#) shows the effect on normal operating power in  $N+1$  and  $N+N$  configurations.

**Table 255: Available Operating Power in N+1 and N+N Redundancy Configurations**

| Number of Power Supplies at $n$ W Each | Normal Operating Power in $N+1$ Configuration | Normal Operating Power in $N+N$ Configuration |
|----------------------------------------|-----------------------------------------------|-----------------------------------------------|
| 2                                      | 1 x ( $n$ W)                                  | 1 x ( $n$ W)                                  |
| 3                                      | 2 x ( $n$ W)                                  | 2 x ( $n$ W)                                  |
| 4                                      | 3 x ( $n$ W)                                  | 2 x ( $n$ W)                                  |
| 5 (EX8200 switches only)               | 4 x ( $n$ W)                                  | 3 x ( $n$ W)                                  |
| 6 (EX8200 switches only)               | 5 x ( $n$ W)                                  | 3 x ( $n$ W)                                  |

To compensate for the reduced normal operating power, power management on EX8200 switches allocates less power to the chassis in an  $N+N$  configuration than in an  $N+1$  configuration. This reduction in allocated chassis power allows a switch in an  $N+N$  configuration to power more line cards than it could without the reduction. For the EX8208 switch, the power allocated for the chassis is reduced to 1200 W from 1600 W; for the EX8216 switch, it is reduced to 1800 W from 2400 W.



**NOTE:** To achieve the reduction in allocated chassis power in an EX8200 switch, power management reduces the maximum fan speed to 60 percent in an  $N+N$  configuration from 80 percent in an  $N+1$  configuration. Because the maximum fan speed is reduced, it is possible that a line card that overheats would be shut down sooner in an  $N+N$  configuration than in an  $N+1$  configuration.

On EX6200 switches, the same amount of power is allocated for the chassis in  $N+N$  configurations as in  $N+1$  configurations.

Power management automatically recalculates the reserved power and normal operating power as power supplies go online or offline. For example, if you have an  $N+N$  configuration with three online 2000 W power supplies, power management allocates 2000 W to reserved power. If you bring a fourth 2000 W power supply online, power management then allocates 4000 W to reserved power. If a power supply goes offline again, power management once again allocates 2000 W to reserved power.

When power is insufficient to meet the budgeted power requirements, power management raises alarms as follows:

- A minor (yellow) alarm is raised when insufficient power exists to maintain the configured  $N+1$  or  $N+N$  power reserves, but all line cards are still receiving their base and PoE power allocations. If this condition persists for 5 minutes, the alarm becomes a major (red) alarm. Even though operation of the switch is unaffected in this condition, you should remedy it as quickly as possible because a power supply failure might cause a disruption in switch operation.

- A major (red) alarm is raised when insufficient power exists to provide all the line cards with their base and PoE power allocations. One or more PoE ports might be down or one or more line cards might be down.

Power management clears all alarms when sufficient power is available to meet normal operating and reserved power requirements.

#### Related Documentation

- [Understanding Alarm Types and Severity Levels on EX Series Switches on page 797](#)
- [Configuring the Power Priority of Line Cards \(CLI Procedure\)](#)
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 2512](#)
- [Verifying Power Configuration and Use on page 2545](#)

## VRRP Overview

- [Understanding VRRP on EX Series Switches on page 2501](#)

### Understanding VRRP on EX Series Switches

Juniper Networks EX Series Ethernet Switches support the Virtual Router Redundancy Protocol (VRRP) and VRRP for IPv6. This topic covers:

- [Overview of VRRP on EX Series Switches on page 2501](#)
- [Examples of VRRP Topologies on page 2502](#)

#### Overview of VRRP on EX Series Switches

You can configure the Virtual Router Redundancy Protocol (VRRP) or VRRP for IPv6 on Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and logical interfaces on EX Series switches. When VRRP is configured, the switches act as virtual routing platforms. VRRP enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master routing platform fails, one of the backup routing platforms becomes the new master, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup EX Series switch can take over a failed default switch within a few seconds. This is done with minimum loss of VRRP traffic and without any interaction with the hosts. Virtual Router Redundancy Protocol is not supported on management interfaces.

VRRP for IPv6 provides a much faster switchover to an alternate default routing platform than IPv6 Neighbor Discovery (ND) procedures. VRRP for IPv6 does not support the **authentication-type** or **authentication-key** statements.



**NOTE:** Do not confuse the VRRP master and backup routing platforms with the master and backup member switches of a Virtual Chassis configuration. The master and backup members of a Virtual Chassis configuration compose a single host. In a VRRP topology, one host operates as the master routing platform and another operates as the backup routing platform, as shown in [Figure 40 on page 2503](#).

Switches running VRRP dynamically elect master and backup routing platforms. You can also force assignment of master and backup routing platforms using priorities from 1 through 255, with 255 being the highest priority. In VRRP operation, the default master routing platform sends advertisements to backup routing platforms at regular intervals. The default interval is 1 second. If the backup routing platforms do not receive an advertisement for a set period, the backup routing platform with the highest priority takes over as master and begins forwarding packets.




**NOTE:** Priority 255 cannot be set for routed VLAN interfaces (RVIs).

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol*.

### Examples of VRRP Topologies

[Figure 39 on page 2502](#) illustrates a basic VRRP topology with EX Series switches. In this example, Switches A, B, and C are running VRRP and together they make up a virtual routing platform. The IP address of this virtual routing platform is 10.10.0.1 (the same address as the physical interface of Switch A).


**Figure 39: Basic VRRP on EX Series Switches**



[Figure 40 on page 2503](#) illustrates a basic VRRP topology using Virtual Chassis configurations. Switch A, Switch B, and Switch C are each composed of multiple interconnected Juniper Networks EX4200 Ethernet Switches. Each Virtual Chassis configuration operates as a single switch, which is running VRRP, and together they make

up a virtual routing platform. The IP address of this virtual routing platform is **10.10.0.1** (the same address as the physical interface of Switch A).

**Figure 40: VRRP on Virtual Chassis Switches**



Because the virtual routing platform uses the IP address of the physical interface of Switch A, Switch A is the master VRRP routing platform, while Switch B and Switch C function as backup VRRP routing platforms. Clients 1 through 3 are configured with the default gateway IP address of **10.10.0.1** as the master router. Switch A, forwards packets sent to its IP address. If the master routing platform fails, the switch configured with the higher priority becomes the master virtual routing platform and provides uninterrupted service for the LAN hosts. When Switch A recovers, it becomes the master virtual routing platform again.

**Related Documentation**

- For more information on VRRP or VRRP for IPv6, see the [Junos OS High Availability Configuration Guide](#).
- [High Availability Features for EX Series Switches Overview](#)
- [Configuring VRRP for IPv6 \(CLI Procedure\) on page 2513](#)





## CHAPTER 41

# Configuration

- [Configuration Tasks on page 2505](#)
- [Configuration Examples on page 2514](#)
- [Configuration Statements on page 2517](#)

## Configuration Tasks

---

- [Configuring Nonstop Active Routing on Switches on page 2505](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 2507](#)
- [Upgrading Software on an EX3300, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and Mixed Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 2508](#)
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 2512](#)
- [Configuring VRRP for IPv6 \(CLI Procedure\) on page 2513](#)

## Configuring Nonstop Active Routing on Switches

Nonstop active routing (NSR) provides a mechanism for transparent switchover of the Routing Engines without necessitating restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbors.

You can configure NSR on an on a Juniper Networks EX Series switch with multiple Routing Engines or an EX Series or QFX Series switch in a Virtual Chassis or Virtual Chassis Fabric configuration.

To configure nonstop active routing:

1. Enable graceful Routing Engine switchover (GRES):  

```
[edit chassis redundancy]  
user@switch# set graceful-switchover
```
2. Enable nonstop active routing (by default, nonstop active routing is disabled):  

```
[edit routing-options]  
user@switch# set nonstop-routing
```
3. Synchronize configuration changes between the Routing Engines:  

```
[edit system]  
user@switch# set commit synchronize
```

If you try to commit the nonstop active routing configuration without including the **commit synchronize** statement, the commit fails.



**NOTE:** There is no requirement to start the two Routing Engines simultaneously. If the backup Routing Engine is not up when you issue the **commit synchronize** command, the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine is inserted or comes online, its configuration is automatically synchronized with that of the master.



**BEST PRACTICE:** After a graceful Routing Engine switchover, we recommend that you issue the clear interface statistics (*interface-name* | all) command to reset the cumulative values for local statistics on the new master Routing Engine.

To disable nonstop active routing:

```
[edit routing-options]  
user@switch# delete nonstop-routing
```

**Related  
Documentation**

- [Example: Configuring Nonstop Active Routing on Switches on page 2514](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 2573](#)
- [Understanding Nonstop Active Routing on EX Series Switches on page 2486](#)
- [Nonstop Active Routing Concepts](#)

## Configuring Nonstop Bridging on Switches (CLI Procedure)



**NOTE:** This task uses switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

You can configure nonstop bridging (NSB) to provide resilience for Layer 2 protocol sessions on a Juniper Networks EX Series switch with multiple Routing Engines or an EX Series or QFX Series switch in a Virtual Chassis or Virtual Chassis Fabric configuration. Limited support for NSB is also provided on QFX5100 and EX4600 standalone switches, but NSB is enabled *only* during an ISSU.

NSB operates by synchronizing all protocol information for NSB-supported Layer 2 protocols between the master and backup Routing Engines. If the switch has a Routing Engine switchover, the NSB-supported Layer 2 protocol sessions remain active because they are already synchronized on the backup Routing Engine. The Routing Engine switchover is transparent to neighbor devices, which do not detect any changes related to the Layer 2 protocol sessions. The neighboring devices and other devices on the network do not, therefore, have to resynchronize their Layer 2 protocol states to respond to the downtime on the switch—a process that adds network overhead and risks disrupting network performance—when a Routing Engine switchover occurs when NSB is enabled.



**NOTE:** If you are using a QFX5100 or EX4600 standalone switch and you want to use ISSU, configure Graceful Routing Engine switchover (GRES), NSB and nonstop active routing (NSR). You must configure NSB, GRES, and NSR in order to run ISSU. However, GRES, NSB and NSR are enabled *only* during the upgrade. During an ISSU, the Junos OS runs in two separate virtual machines (VMs)—one VM is in the master role acting as the master Routing Engine, and the other VM is in the backup role acting as the backup Routing Engine. The Junos OS is upgraded on the backup VM. After a successful software upgrade, the backup VM then becomes the master VM, and the original master VM is no longer needed and is shut down.

To configure NSB:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch# set graceful-switchover
```

2. Enable NSB:

```
[edit protocols layer2-control]
user@switch# set nonstop-bridging
```

3. Synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch# set commit synchronize
```

If you try to commit a configuration that includes NSB without including the **commit synchronize** statement, the commit fails.



**NOTE:** There is no requirement to start the two Routing Engines simultaneously. If the backup Routing Engine is not up when you use the **commit synchronize** statement, the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine comes online, its configuration is automatically synchronized with that of the master.



**BEST PRACTICE:** After a graceful Routing Engine switchover, we recommend that you issue the clear interface statistics (*interface-name | all*) command to reset the cumulative values for local statistics on the new master Routing Engine.

#### Related Documentation

- [Performing an In-Service Software Upgrade \(ISSU\)](#)
- [Understanding Nonstop Bridging on EX Series Switches on page 2487](#)
- [Nonstop Bridging Concepts](#)
- [Understanding In-Service Software Upgrade \(ISSU\)](#)

### Upgrading Software on an EX3300, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and Mixed Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)

You can use nonstop software upgrade (NSSU) to upgrade the software running on all member switches in most EX Series Virtual Chassis with minimal traffic disruption during the upgrade.

NSSU is supported on the following Virtual Chassis platforms:

- EX3300 Virtual Chassis
- EX4200 Virtual Chassis
- EX4300 Virtual Chassis
- EX4500 Virtual Chassis
- EX4550 Virtual Chassis
- All mixed Virtual Chassis composed of EX4200, EX4500, and EX4550 switches
- EX8200 Virtual Chassis

This topic covers:

- [Preparing the Switch for Software Installation on page 2509](#)
- [Upgrading the Software Using NSSU on page 2510](#)

## Preparing the Switch for Software Installation

Before you begin software installation using NSSU:

- Ensure that the Virtual Chassis is configured correctly to support NSSU. Verify that:
  - The Virtual Chassis members are connected in a ring topology. A ring topology prevents the Virtual Chassis from splitting during an NSSU.
  - The Virtual Chassis master and backup are adjacent to each other in the ring topology. Adjacency permits the master and backup to always be in sync, even when the switches in linecard roles are rebooting.
  - The Virtual Chassis is preprovisioned so that the linecard role has been explicitly assigned to member switches acting in the linecard role. During an NSSU, the Virtual Chassis members must maintain their roles—the master and backup must maintain their master and backup roles (although mastership will change), and the other member switches must maintain their linecard roles.

For information on configuring a preprovisioned Virtual Chassis, see *Configuring an EX3300 Virtual Chassis (CLI Procedure)*, *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*, and *Configuring an EX8200 Virtual Chassis (CLI Procedure)*.

- A two-member Virtual Chassis has **no-split-detection** configured so that the Virtual Chassis does not split when an NSSU upgrades a member.
- Verify that the members are running the same version of the software:

```
user@switch> show version
```

If the Virtual Chassis members are not running the same version of the software, use the **request system software add** command to upgrade the software on the inconsistent members.

- Ensure that nonstop active routing (NSR) and graceful Routing Engine switchover (GRES) are enabled. To verify that they are enabled, you need to check only the state of nonstop active routing—if nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

To verify that nonstop active routing is enabled:

```
user@switch> show task replication
```

```
Stateful Replication: Enabled
RE mode: Master
```

| Protocol | Synchronization Status |
|----------|------------------------|
| OSPF     | Complete               |
| BGP      | Complete               |
| PIM      | Complete               |

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see [“Configuring Nonstop Active Routing on Switches” on page 2505](#) for information on how to enable it.

- For the EX4300 Virtual Chassis, you should enable the **vcp-no-hold-time** statement at the **[edit virtual-chassis]** hierarchy level before performing a software upgrade using

NSSU. If you do not enable the **vcp-no-hold-time** statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see [“Understanding Split and Merge in a Virtual Chassis” on page 5087](#)

- (Optional) Enable nonstop bridging (NSB). Enabling NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on each member to an external storage device with the [request system snapshot](#) command.

### Upgrading the Software Using NSSU

---

This procedure describes how to upgrade the software running on all Virtual Chassis members using NSSU. When the upgrade completes, all members are running the new version of the software. Because a graceful Routing Engine switchover occurs during the upgrade, the original Virtual Chassis backup is the new master.

To upgrade all members using NSSU:

1. Download the software package by following the procedure in *Downloading Software Packages from Juniper Networks*. If you are upgrading the software running on a mixed Virtual Chassis, download the software packages for both switch types.
2. Copy the software package or packages to the Virtual Chassis. We recommend that you copy the file to the **/var/tmp** directory on the master.
3. Log in to the Virtual Chassis using the console connection or the virtual management Ethernet (VME) interface. Using a console connection allows you to monitor the progress of the master switch reboot.
4. Start the NSSU:
  - On an EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4300 Virtual Chassis, EX4500 Virtual Chassis, or EX4550 Virtual Chassis, enter:

```
user@switch> request system software nonstop-upgrade  
/var/tmp/package-name.tgz
```

where **package-name.tgz** is, for example, **jinstall-ex4200-12.1R2.5-domestic-signed.tgz**.

- On a mixed Virtual Chassis, enter:

```
user@switch> request system software nonstop-upgrade set  
[/var/tmp/package-name.tgz /var/tmp/package-name.tgz]
```

where **[/var/tmp/package-name.tgz /var/tmp/package-name.tgz]** specifies the EX4200 and EX4500 software packages.

The switch displays status messages similar to the following messages as the upgrade executes:

```
Chassis ISSU Check Done
```

```

ISSU: Validating Image
ISSU: Preparing Backup RE
Installing image on other FPC's along with the backup

Checking pending install on fpc1
Pushing bundle to fpc1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc1

Checking pending install on fpc2
Pushing bundle to fpc2
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc2

Rebooting fpc1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online
  FPC 1         Online
  FPC 2         Online (ISSU)
Going to install image on master
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE

*** FINAL System shutdown message from user@switch ***

System going down IMMEDIATELY

Shutdown NOW!
[pid 9336]

```

5. Log in after the reboot of the original master switch completes. To verify that the software on all Routing Engines in the Virtual Chassis members has been upgraded, enter the following command:

```
user@switch> show version
```

6. To ensure that the resilient dual-root partitions feature operates correctly, copy the new Junos OS image into the alternate root partitions of all members:

```
user@switch> request system snapshot slice alternate all-members
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

**Related Documentation**

- [Understanding Nonstop Software Upgrade on EX Series Switches on page 2488](#)
- [Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\)](#)
- [Understanding Resilient Dual-Root Partitions on Switches on page 695](#)
- [Understanding Software Installation on EX Series Switches on page 691](#)
- [Troubleshooting Software Installation on page 785](#)
- [Junos OS Package Names on page 698](#)
- [Understanding Nonstop Software Upgrade on EX Series Switches on page 2488](#)

## Configuring Power Supply Redundancy (CLI Procedure)

By default, the power management feature in EX Series switches is configured to manage the power supplies for  $N+1$  redundancy, in which one power supply is held in reserve for backup if any one of the other power supplies is removed or fails.

You can configure power management to manage the power supplies for  $N+N$  redundancy. For example, to set up your AC power supplies for dual power feed,  $N+N$  redundancy is required. In  $N+N$  redundancy, power management allocates half of the online power supplies to normal operating power and half to redundant power. If you have an odd number of online power supplies, power management allocates one more power supply to normal operating power than to redundant power.

This topic describes how to configure power management for  $N+N$  redundancy and how to revert back to  $N+1$  redundancy if your deployment needs change.

Before you configure power management for  $N+N$  redundancy, ensure that you have sufficient power supplies to meet the power requirements of an  $N+N$  configuration. Use the [show chassis power-budget-statistics](#) command to display your current power budget.



**NOTE:** To allow more power to be available to line cards in an EX8200 switch, power management compensates for the reduced normal operating power in an  $N+N$  configuration by allocating less power to the chassis than it does in an  $N+1$  configuration. For the EX8208 switch, the power allocated to the chassis is reduced to 1200 W from 1600 W. For the EX8216 switch, it is reduced to 1800 W from 2400 W. In determining whether you have enough power for an  $N+N$  configuration, take this reduction of allocated chassis power into account.

The reduction in allocated chassis power is achieved by reducing the maximum fan speed to 60 percent in an  $N+N$  configuration from 80 percent in an  $N+1$  configuration. Because the maximum fan speed is reduced, it is possible that a line card that overheats would be shut down sooner in an  $N+N$  configuration than in an  $N+1$  configuration.

On EX6200 switches, the same amount of power is allocated for the chassis in  $N+N$  configurations as in  $N+1$  configurations.



To configure *N+N* redundancy:

```
[edit chassis]
user@switch# set psu redundancy n-plus-n
```

To revert back to *N+1* redundancy:

```
[edit chassis]
user@switch# delete chassis psu redundancy n-plus-n
```

#### Related Documentation

- [Configuring the Power Priority of Line Cards \(CLI Procedure\)](#)
- [Verifying Power Configuration and Use on page 2545](#)
- [Understanding Power Management on EX Series Switches on page 2495](#)

## Configuring VRRP for IPv6 (CLI Procedure)

By configuring the Virtual Router Redundancy Protocol (VRRP) on EX Series switches, you can enable hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. You can configure VRRP for IPv6 on Gigabit Ethernet, 10-Gigabit Ethernet, and logical interfaces.

To configure VRRP for IPv6:

1. Configure VRRP group support on interfaces:

```
[edit interfaces interface-name unit logical-unit-number family inet6 address address]
user@switch# set vrrp-inet6-group group-id priority number virtual-inet6-address address
virtual-link-local-address ipv6-address
```

You must explicitly define a virtual link local address for each VRRP for IPv6 group. Otherwise, when you attempt to commit the configuration, the commit request fails. The virtual link local address must be on the same subnet as the physical interface address.

2. If you want to configure the priority order in which this switch functioning as a backup router becomes the master router if the master router becomes nonoperational, configure a priority for this switch:

```
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
user@switch# set priority number
```

3. Specify the interval in milliseconds in which the master router sends advertisement packets to the members of the VRRP group:

```
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
user@switch# set inet6-advertise-interval milliseconds
```

4. By default, a higher-priority backup router preempts a lower-priority master router.

- To explicitly enable the master router to be preempted:

```
[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]
user@switch# set preempt
```

- To prohibit a higher-priority backup router from preempting a lower priority master router:

```
[edit interfaces interface-name unit logical-unit-number family inet6 address  
address vrrp-inet6-group group-id]  
user@switch# set no-preempt
```

**Related  
Documentation**

- [show vrrp on page 2562](#)
- [Understanding VRRP on EX Series Switches on page 2501](#)

---

## Configuration Examples

- [Example: Configuring Nonstop Active Routing on Switches on page 2514](#)

### Example: Configuring Nonstop Active Routing on Switches

Nonstop active routing (NSR) provides high availability for Routing Engines by enabling transparent switchover of the Routing Engines without necessitating restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbors.

This example describes how to configure nonstop active routing on switches with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration.

- [Requirements on page 2514](#)
- [Overview and Topology on page 2514](#)
- [Configuration on page 2515](#)
- [Verification on page 2516](#)
- [Troubleshooting on page 2516](#)

---

#### Requirements

This example uses the following hardware and software components:

- An EX Series with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration
- Junos OS Release 10.4 or later for EX Series switches
- Junos OS Release 13.2X51-D20 or later for QFX Series switches

---

#### Overview and Topology

Configure nonstop active routing on any EX Series with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration. Nonstop active routing is advantageous in networks where neighbor routing devices do not support graceful restart protocol extensions.

The topology used in this example consists of an EX8200 switch with redundant Routing Engines connected to neighbor routing devices that are not configured to support graceful restart of protocols.

### Configuration

#### CLI Quick Configuration

To quickly configure nonstop active routing, copy the following commands and paste them into the switch terminal window:

```
[edit]
set chassis redundancy graceful-switchover
set routing-options nonstop-routing
set system commit synchronize
```

#### Step-by-Step Procedure

To configure nonstop active routing on a switch:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch# set graceful-switchover
```

2. Enable nonstop active routing (by default, nonstop active routing is disabled):

```
[edit routing-options]
user@switch# set nonstop-routing
```

3. Synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch# set commit synchronize
```

If you try to commit the nonstop active routing configuration without including the **commit synchronize** statement, the commit fails.



**NOTE:** If the backup Routing Engine is down when you issue the commit, a warning is displayed and the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine comes up, its configuration is automatically synchronized with that of the master. If you subsequently insert or bring up a backup Routing Engine, it automatically synchronizes its configuration with the master Routing Engine configuration.

### Results

Check the results of the configuration:

```
[edit]
user@switch# show
chassis {
  redundancy {
    graceful-switchover;
  }
}
routing-options {
  nonstop-routing;
}
system {
  commit synchronize;
```

```
}
```

## Verification

---

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That Nonstop Active Routing Is Working Correctly on the Switch on page 2516](#)

### *Verifying That Nonstop Active Routing Is Working Correctly on the Switch*

**Purpose** Verify that nonstop active routing is enabled.

**Action** Issue the [show task replication](#) command:

```
user@switch# show task replication
Stateful Replication: Enabled
RE mode: Master
```

| Protocol | Synchronization Status |
|----------|------------------------|
| OSPF     | Complete               |
| RIP      | Complete               |
| PIM      | Complete               |
| RSVP     | Complete               |

**Meaning** This output shows that nonstop active routing (Stateful Replication) is enabled on master routing engine. If nonstop routing is not enabled, instead of the output shown above:

- On the backup routing engine the following error message is displayed: “**error: the routing subsystem is not running.**”
- On the master routing engine, the following output is displayed if nonstop routing is not enabled:

```
Stateful Replication: Disabled
RE mode: Master
```

## Troubleshooting

---

To troubleshoot nonstop active routing, perform these tasks:

- [Investigating Problems with Synchronization of Routing Engines When NSR Is Enabled on page 2516](#)

### *Investigating Problems with Synchronization of Routing Engines When NSR Is Enabled*

**Problem** A protocol loses connectivity with neighbors after a graceful Routing Engine switchover (GRES) occurs with nonstop active routing (NSR) enabled.

**Solution** Use trace options to help isolate the problem and gather troubleshooting information. Using the information gathered from trace options, you can confirm or eliminate the synchronization of the Routing Engines as the cause of the loss of connectivity for the protocol. See “[Tracing Nonstop Active Routing Synchronization Events](#)” on page 2573.

- Related Documentation**
- [Configuring Nonstop Active Routing on Switches on page 2505](#)
  - [Tracing Nonstop Active Routing Synchronization Events on page 2573](#)
  - [Understanding Nonstop Active Routing on EX Series Switches on page 2486](#)
  - *Nonstop Active Routing Concepts*

## Configuration Statements

- [\[edit chassis\] Configuration Statement Hierarchy on EX Series Switches on page 2517](#)
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2519](#)
- [\[edit protocols layer2-control\] Hierarchy Level on page 2520](#)
- [\[edit protocols vrrp\] Hierarchy Level on page 2521](#)
- [chassis on page 2522](#)
- [synchronize on page 2524](#)
- [failover \(Chassis\) on page 2525](#)
- [fpc on page 2526](#)
- [graceful-switchover on page 2527](#)
- [hold-time on page 2528](#)
- [inet6-advertise-interval on page 2528](#)
- [keepalive-time on page 2529](#)
- [n-plus-n \(Power Management\) on page 2530](#)
- [nonstop-bridging on page 2530](#)
- [nonstop-routing on page 2531](#)
- [preempt on page 2532](#)
- [priority on page 2533](#)
- [psu on page 2534](#)
- [redundancy \(Graceful Switchover\) on page 2535](#)
- [redundancy \(Power Management\) on page 2536](#)
- [traceoptions \(Routing Options\) on page 2537](#)
- [vcp-no-hold-time on page 2540](#)
- [virtual-inet6-address on page 2541](#)
- [virtual-link-local-address on page 2542](#)
- [vrrp-inet6-group on page 2543](#)

### [edit chassis] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit chassis]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.

- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit chassis\] Hierarchy Level on page 2518](#)

### [Supported Statements in the \[edit chassis\] Hierarchy Level](#)

---

The following hierarchy shows the **[edit chassis]** configuration statements supported on EX Series switches:

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
      lacp {
        link-protection non-revertive;
        system-priority system-priority-number
      }
    }
  }
  alarm {
    ethernet {
      link-down (ignore | red | yellow);
    }
    management-ethernet {
      link-down (ignore | red | yellow);
    }
  }
  container-devices {
    device-count device-count-number;
  }
  disk-partition {
    /config {
      level (full | high) {
        free-space (free-space-threshold-value | mb | percent);
      }
    }
  }
  /var {
    level (full | high) {
      free-space (free-space-threshold-value | mb | percent);
    }
  }
}
fpc slot-number {
  pic pic-number {
    no-multi-rate;
    q-pic-large-buffer (large-scale | small-scale);
  }
}
```

```

maximum-ecmp maximum-ecmp-routes;
lcd-menu {
  fpc slot-number {
    menu-item menu-name);
    disable;
  }
  pseudowire-service {
    device-count device-count-number;
  }
  psu {
    redundancy {
      n-plus-n;
    }
    redundancy {
      graceful-switchover;
    }
  }
  slow-pfe-alarm;
}

```

#### Related Documentation

- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\) on page 33](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 5116](#)
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 2512](#)
- [Configuring the Power Priority of Line Cards \(CLI Procedure\)](#)
- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade \(CLI Procedure\)](#)

### [edit interfaces] Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the **[edit interfaces]** hierarchy:

- [\[edit interfaces ae\] Configuration Statement Hierarchy on EX Series Switches on page 2700](#)
- [\[edit interfaces et\] Configuration Statement Hierarchy on EX Series Switches on page 2705](#)
- [\[edit interfaces ge\] Configuration Statement Hierarchy on EX Series Switches on page 2711](#)
- [\[edit interfaces interface-range\] Configuration Statement Hierarchy on EX Series Switches on page 2716](#)
- [\[edit interfaces irb\] Configuration Statement Hierarchy on EX Series Switches on page 2725](#)
- [\[edit interfaces lo\] Configuration Statement Hierarchy on EX Series Switches on page 360](#)
- [\[edit interfaces me\] Configuration Statement Hierarchy on EX Series Switches on page 363](#)

- [\[edit interfaces vme\] Configuration Statement Hierarchy on EX Series Switches on page 370](#)
- [\[edit interfaces xe\] Configuration Statement Hierarchy on EX Series Switches on page 2738](#)

**Related  
Documentation**

- [EX Series Switches Interfaces Overview on page 2577](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 2689](#)
- [Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\) on page 2340](#)
- [Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis \(CLI Procedure\)](#)
- [Junos OS Interfaces Fundamentals Configuration Guide](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

## [\[edit protocols layer2-control\] Hierarchy Level](#)

The following statement hierarchy can also be included at the [\[edit logical-systems logical-system-name\]](#) hierarchy level.

```
protocols {
  layer2-control {
    bpdu-block {
      disable-timeout seconds;
      interface [ interface-names ];
    }
    mac-rewrite {
      interface interface-name {
        enable-all-ifl;
        protocol {
          cdp;
          stp;
          vtp;
          pvstp;
        }
      }
    }
    nonstop-bridging;
    traceoptions {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      flag flag <disable>;
    }
  }
}
```



- Related Documentation**
- [layer2-control on page 2399](#)
  - *Notational Conventions Used in Junos OS Configuration Hierarchies*
  - *[edit protocols] Hierarchy Level*

## **[edit protocols vrrp] Hierarchy Level**

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
  vrrp {
    asymmetric-hold-time;
    delegate-processing;
    failover-delay milliseconds;
    global-advertisements-threshold advertisement-value;
    skew-timer-disable;
    startup-silent-period seconds;
    traceoptions {
      file <filename> <files number> <match regular-expression> <microsecond-stamp>
        <size maximum-file-size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
    version-3;
  }
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
  - *[edit protocols] Hierarchy Level*
  - *Junos OS Hierarchy and RFC Reference*
  - *Ethernet Interfaces Feature Guide for Routing Devices*
  - *Junos OS Network Interfaces Library for Routing Devices*

## chassis

```
Syntax  chassis {
        aggregated-devices {
            ethernet (Aggregated Devices) {
                device-count number;
            }
        }
        auto-image-upgrade;
        fpc slot {
            pic pic-number {
                sfpplus {
                    pic-mode mode;
                }
            }
            power-budget-priority priority;
        }
        lcd-menu {
            fpc slot-number {
                menu-item (menu-name | menu-option) {
                    disable;
                }
            }
        }
        nssu {
            upgrade-group group-name {
                fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
                member (NSSU Upgrade Groups) member-id {
                    fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
                }
            }
        }
        psu {
            redundancy {
                n-plus-n (Power Management);
            }
        }
        redundancy {
            graceful-switchover;
        }
    }
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure chassis-specific properties for the switch.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

**Related  
Documentation**

- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [\*Upgrading Software by Using Automatic Software Download\*](#)
- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\) on page 33](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 5116](#)
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 2512](#)
- [\*Configuring the Power Priority of Line Cards \(CLI Procedure\)\*](#)
- [\*Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade \(CLI Procedure\)\*](#)

## synchronize

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | synchronize;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | [edit system commit]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 10.4 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>         | For devices with multiple Routing Engines only. Configure the <b>commit</b> command to automatically perform a <b>commit synchronize</b> action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the <b>commit</b> command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engine. Each Routing Engine then performs a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines. |



**NOTE:** When you configure nonstop active routing (NSR), you must configure the **commit synchronize** statement. Otherwise, the commit operation fails.

---

On the TX Matrix router, synchronization only occurs between the Routing Engines within the same chassis. When synchronization is complete, the new configuration is then distributed to the Routing Engines on the T640 routers. That is, the master Routing Engine on the TX Matrix router distributes the configuration to the master Routing Engine on each T640 router. Likewise, the backup Routing Engine on the TX Matrix router distributes the configuration to the backup Routing Engine on each T640 router.

On the TX Matrix Plus router, synchronization only occurs between the Routing Engines within the switch-fabric chassis and when synchronization is complete, the new configuration is then distributed to the Routing Engines on the line-card chassis (LCC). That is, the master Routing Engine on the TX Matrix Plus router distributes the configuration to the master Routing Engine on each LCC. Likewise, the backup Routing Engine on the TX Matrix Plus router distributes the configuration to the backup Routing Engine on each LCC.

In EX Series Virtual Chassis configurations:

- On EX4200 switches in Virtual Chassis, synchronization occurs between the switch in the master role and the switch in the backup role.
- On EX8200 switches in a Virtual Chassis, synchronization occurs only between the master and backup XRE200 External Routing Engines.

|                |                                                                                             |
|----------------|---------------------------------------------------------------------------------------------|
| <b>Options</b> | <b>and-quit</b> —(Optional) Quit configuration mode if the commit synchronization succeeds. |
|                | <b>at</b> —(Optional) Time at which to activate configuration changes.                      |
|                | <b>comment</b> —(Optional) Write a message to the commit log.                               |

**force**—(Optional) Force a commit synchronization on the other Routing Engine (ignore warnings).

**scripts**—(Optional) Push scripts to the other Routing Engine.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- *Synchronizing the Routing Engine Configuration*
- *Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically*

## failover (Chassis)

**Syntax**

```
failover {
  on-disk-failure;
  on-loss-of-keepalives;
}
```

**Hierarchy Level** [edit chassis [redundancy](#)]

**Release Information** Statement introduced in Junos OS Release 9.2 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Specify conditions on the master Routing Engine that cause the backup router to take mastership.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [graceful-switchover on page 2527](#)
- *On Detection of a Hard Disk Error on the Master Routing Engine*
- *Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)*
- *High Availability Features for EX Series Switches Overview*

## fpc

---

**Syntax** `fpc slot {  
 pic pic-number {  
 sfpplus {  
 pic-mode mode;  
 }  
 tunnel-port port-number tunnel-services;  
 }  
 power (off | on);  
 power-budget-priority priority;  
}`

**Hierarchy Level** [edit [chassis](#)]

**Release Information** Statement introduced in Junos OS Release 9.4 for EX Series switches.  
Statement introduced in Junos OS Release 13.2 for the QFX Series.

**Description** Specify the port of the SFP+ uplink module for which you want to configure the operating mode, or specify the line card slot for which you want to assign a power priority.



**NOTE:** On an EX6200 switch, you cannot change the power priority of a slot containing a Switch Fabric and Routing Engine (SRE) module. Although the CLI allows you to set a power priority for the slot, your change does not take effect, and the power priority remains 0. A message is sent to the system log to inform you that changing the power priority of the slot is unsupported.

For generic routing encapsulation (GRE) tunneling, use the **fpc** statement along with the **tunnel-port** statement to specify the port on the switch that you want to convert to a GRE tunnel port.

**Options** *slot*—Number of the slot:

- 0—EX3200 and standalone EX4200, EX4500, and EX4550 switches. The FPC value refers to the switch itself.
- 0–9—EX4200, EX4500, or EX4550 switch in a Virtual Chassis configuration. The value corresponds to the switch's member ID.
- 0–3 and 6–9—EX6210 switch. The slot is a line card slot.  
4–5—The slot is a line card slot or an SRE module slot.
- 0–7—EX8208 switch. The slot is a line card slot.
- 0–15—EX8216 switch. The slot is a line card slot.

The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module (CLI Procedure)</i></li> <li>• <i>Configuring the Power Priority of Line Cards (CLI Procedure)</i></li> <li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li> </ul> |

## graceful-switchover

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | graceful-switchover;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis redundancy</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | For switches with more than one Routing Engine, including those in a Virtual Chassis or a Virtual Chassis Fabric, configure the master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | Graceful Routing Engine switchover (GRES) is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Nonstop Active Routing on Switches on page 2514</a></li> <li>• <i>Configuring Graceful Routing Engine Switchover</i></li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 5116</a></li> <li>• <a href="#">Configuring Nonstop Active Routing on Switches on page 2505</a></li> <li>• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i></li> </ul> |

## hold-time

---

|                                 |                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hold-time seconds;</code>                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <b>vrrp-inet6-group</b> <i>group-id</i> preempt] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                       |
| <b>Description</b>              | Configure the time in seconds after which a backup router with the highest priority preempts the master router.                                             |
| <b>Options</b>                  | <i>seconds</i> —Hold-time period.                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring VRRP for IPv6 (CLI Procedure) on page 2513</a></li></ul>                                    |

## inet6-advertise-interval

---

|                                 |                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>inet6-advertise-interval milliseconds;</code>                                                                                                  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <b>vrrp-inet6-group</b> <i>group-id</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                |
| <b>Description</b>              | Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv6 advertisement packets.                                                 |
| <b>Options</b>                  | <i>milliseconds</i> —Interval, in milliseconds, between advertisement packets.<br><b>Range:</b> 100 to 40,000 ms<br><b>Default:</b> 1 second         |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring VRRP for IPv6 (CLI Procedure) on page 2513</a></li></ul>                             |



## keepalive-time

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>keepalive-time <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit chassis <a href="#">redundancy</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure the time period that must elapse before the backup router takes mastership when it detects loss of the keepalive signal.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | <p>The <b>on-loss-of-keepalives</b> statement at the <b>[edit chassis redundancy failover]</b> hierarchy level must be included for failover to occur.</p> <p>When the <b>on-loss-of-keepalives</b> statement is included and graceful Routing Engine switchover <i>is not</i> configured, failover occurs after 300 seconds (5 minutes).</p> <p>When the <b>on-loss-of-keepalives</b> statement is included and graceful Routing Engine switchover <i>is</i> configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds.</p> |
| <b>Options</b>                  | <b><i>seconds</i></b> —Time before the backup router takes mastership when it detects loss of the keepalive signal. The range of values is 2 through 10,000.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">failover on page 2525</a></li> <li>• <a href="#">graceful-switchover on page 2527</a></li> <li>• <i>on-loss-of-keepalives</i></li> <li>• <i>High Availability Features for EX Series Switches Overview</i></li> </ul>                                                                                                                                                                                                                                                                                             |

## n-plus-n (Power Management)

---

|                                 |                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | n-plus-n;                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis psu redundancy</a> ]                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                              |
| <b>Description</b>              | Configure N+N power supply redundancy for power management on an EX6200 or EX8200 switch.                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Power Supply Redundancy (CLI Procedure) on page 2512</a></li></ul> |

## nonstop-bridging

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | nonstop-bridging;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols layer2-control]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | For platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 Control Protocol (L2CP) information.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Synchronizing the Routing Engine Configuration</i></li><li>• <i>Configuring Nonstop Bridging</i></li><li>• For information about configuring NSB on EX Series switches that do not support the Enhanced Layer 2 Software (ELS) CLI style, see <i>Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)</i></li><li>• For information about configuring NSB on switches that support ELS, see <a href="#">Configuring Nonstop Bridging on Switches (CLI Procedure) on page 2507</a></li></ul> |

## nonstop-routing

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | nonstop-routing;                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options]                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 10.4 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.<br>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches                                                                                                               |
| <b>Description</b>              | For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and to preserve routing protocol information.                                                                                                                                                                                                          |
| <b>Default</b>                  | disabled                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Nonstop Active Routing</a></li> <li>• <a href="#">Configuring Nonstop Active Routing on Switches on page 2505</a></li> <li>• <a href="#">Example: Configuring Nonstop Active Routing on Switches on page 2514</a></li> <li>• <a href="#">Example: Configuring Nonstop Active Routing on Switches on page 2514</a></li> </ul> |

## preempt

---

|                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (preempt   no-preempt) {<br>hold-time seconds;<br>}                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i><br>vrrp-inet6-group <i>group-id</i> ]                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure whether a backup router can preempt a master router:</p> <ul style="list-style-type: none"><li>• <b>preempt</b>—Allow the master router to be preempted.</li><li>• <b>no-preempt</b>—Prohibit the preemption of the master router.</li></ul> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring VRRP for IPv6 (CLI Procedure) on page 2513</a></li></ul>                                                                                                                                                                                          |

## priority

|                            |                                                                                                                                                                                                                                                                                                 |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <b>vrrp-inet6-group</b> <i>group-id</i> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                                                                           |
| <b>Description</b>         | Configure a switch's priority for becoming the master default routing platform. The routing platform with the highest priority within the group becomes the master.                                                                                                                             |
| <b>Options</b>             | <b><i>number</i></b> —Routing platform's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected.<br><b>Range:</b> 1 through 255<br><b>Default:</b> 100 (for backup routers)                                         |



**NOTE:** Priority 255 cannot be assigned to routed VLAN interfaces (RVIs).

|                                 |                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring VRRP for IPv6 (CLI Procedure) on page 2513</a></li> </ul> |

## psu

---

|                                 |                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>psu {<br/>    redundancy {<br/>        n-plus-n (Power Management);<br/>    }<br/>}</pre>                                                                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis</a> ]                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                             |
| <b>Description</b>              | <p>Configure <i>N+N</i> power supply redundancy for power management on an EX6200 or EX8200 switch.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Power Supply Redundancy (CLI Procedure) on page 2512</a></li></ul>                                |

## redundancy (Graceful Switchover)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> redundancy {     failover {         on-disk-failure;         on-loss-of-keepalives;     }     graceful-switchover; } </pre>                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Enable redundant Routing Engines on a Virtual Chassis with two or more member switches or on a Virtual Chassis Fabric, on a standalone EX6200 or EX8200 switch with more than one Routing Engine.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                         |
| <b>Default</b>                  | Redundancy is enabled for the Routing Engines.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">graceful-switchover on page 2527</a></li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 5116</a></li> <li>• <i>Configuring Graceful Routing Engine Switchover</i></li> <li>• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i></li> <li>• <i>High Availability Features for EX Series Switches Overview</i></li> </ul> |

## redundancy (Power Management)

---

|                                 |                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>redundancy {<br/>    n-plus-n (Power Management);<br/>}</code>                                                                                     |
| <b>Hierarchy Level</b>          | [edit <code>chassis psu</code> ]                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                    |
| <b>Description</b>              | <p>Configure N+N power supply redundancy for power management on an EX6200 or EX8200 switch.</p> <p>The remaining statement is explained separately.</p> |
| <b>Default</b>                  | N+1 power supply redundancy is configured by default.                                                                                                    |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Power Supply Redundancy (CLI Procedure) on page 2512</a></li></ul>                       |



## traceoptions (Routing Options)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>                                                                                                                                                                                                                                         |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>nsr-synchronization</b> flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p><b>nsr-synchronization</b> and <b>nsr-packet</b> flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>nsr-synchronization</b> flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p><b>nsr-synchronization</b> flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p><b>nsr-synchronization</b> flag for PIM added in Junos OS Release 9.3.</p> <p><b>nsr-synchronization</b> flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>nsr-synchronization</b> flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> |
| <b>Description</b>         | <p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>             | If you do not include this statement, no global tracing operations are performed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>             | <p><b>Values:</b></p> <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place global routing protocol tracing output in the file <b>routing-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and</p>                                                                                                                                                                                                 |

so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization-detail**—(MPLS only) Tracing operations for nonstop active routing in detail
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.


**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Tracing Global Routing Protocol Operations*
- [Tracing Nonstop Active Routing Synchronization Events on page 2573](#)

## vcp-no-hold-time

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | vcp-no-hold-time;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Disable the Virtual Chassis port (VCP) holddown timer for all VCPs in the Virtual Chassis or Virtual Chassis Fabric (VCF).</p> <p>The VCP holddown timer is an internal mechanism that delays a Virtual Chassis reconvergence for several seconds when a VCP becomes inactive. The purpose of this delay is to provide the VCP time to return online without having to reconverge the Virtual Chassis to adjust to the inactive VCP. All traffic to the VCP is dropped while the VCP is inactive. If the VCP remains down for a time that exceeds the VCP holddown timer, a Virtual Chassis reconvergence occurs.</p> <p>When this statement is enabled, the VCP holddown timer is disabled and the Virtual Chassis reconvergence occurs when a VCP becomes inactive. The period of time where traffic is dropped waiting for the VCP to return online is avoided.</p> <p>We recommend enabling this statement after a Virtual Chassis is operational. We recommend disabling this statement when you are adding or removing member switches from your Virtual Chassis.</p> <p>The VCP holddown timer cannot be viewed and is not user-configurable. You can only control whether the VCP holddown timer is enabled or disabled by configuring this statement.</p> <div> <b>NOTE:</b> For the EX4300 Virtual Chassis, you should enable the <code>vcp-no-hold-time</code> statement before performing a software upgrade using NSSU. If you do not enable the <code>vcp-no-hold-time</code> statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see <a href="#">“Understanding Split and Merge in a Virtual Chassis”</a> on page 5087</div> |
| <b>Default</b>                  | The VCP holddown timer is enabled by default on all devices that support this statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding EX4300 Virtual Chassis on page 5065</a></li><li>• <a href="#">Understanding QFX Series Virtual Chassis</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

- [Understanding EX Series Virtual Chassis Components on page 5067](#)
- [Understanding QFX Series Virtual Chassis Components](#)

## virtual-inet6-address

|                            |                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>virtual-inet6-address [addresses];</code>                                                                                                    |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> ]      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                              |
| <b>Description</b>         | Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You can configure up to eight addresses. |



**NOTE:** The address of an aggregated Ethernet interface (a LAG) or a routed VLAN interface (RVI) cannot be assigned as the virtual router address in a VRRP IPv6 group.

|                                 |                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <b>addresses</b> —Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring VRRP for IPv6 (CLI Procedure) on page 2513</a></li> </ul>                                                                                                 |

## virtual-link-local-address

---

|                                 |                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>virtual-link-local-address <i>ipv6-address</i>;</code>                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-inet6-group <i>group-id</i>]</code><br><code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure a virtual link local address for a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You must explicitly define a virtual link local address for each VRRP IPv6 group. The virtual link local address must be in the same subnet as the physical interface address.                                     |
| <b>Options</b>                  | <i>ipv6-address</i> —Virtual link local IPv6 address for VRRP for an IPv6 group.                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <code>interface</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring VRRP for IPv6 (CLI Procedure) on page 2513</a></li></ul>                                                                                                                                                                                              |

## vrrp-inet6-group

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>vrrp-inet6-group <i>group-id</i> {   inet6-advertise-interval <i>milliseconds</i>;   preempt {     hold-time <i>seconds</i>;   }   priority <i>number</i>;   virtual-inet6-address;   virtual-link-local-address }</pre>                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Configure a Virtual Router Redundancy Protocol (VRRP) IPv6 group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b><i>group-id</i></b>—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the <b>source-address-filter</b> statement. The MAC address <b>00-00-5E-00-02-{VRID}</b> is reserved for VRRP, as defined in RFC 5798. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p><b>Range:</b> 0 through 255</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring VRRP for IPv6 (CLI Procedure) on page 2513</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                           |





# Administration

- [Verification Tasks on page 2545](#)
- [Operational Commands on page 2547](#)

## Verification Tasks

---

- [Verifying Power Configuration and Use on page 2545](#)

### Verifying Power Configuration and Use

**Purpose** Verify on an EX Series switch:

- The power redundancy and line card priority settings
- The PoE power budgets for line cards that support PoE
- Whether the  $N+1$  or  $N+N$  power requirements are being met
- Whether the switch has sufficient power for a new line card or an  $N+N$  configuration

**Action** Enter the following command:

```
user@switch> show chassis power-budget-statistics
```

Example output for an EX6200 switch:

```

PSU 0      (EX6200-PWR-AC2500)      :   2500 W   Online
PSU 1      (EX6200-PWR-AC2500)      :   2500 W   Online
PSU 2      (EX6200-PWR-AC2500)      :   2500 W   Online
PSU 3      (EX6200-PWR-AC2500)      :   2500 W   Online
Total Power supplied by all Online PSUs : 10000 W
Power Redundancy Configuration         :   N+1
Power Reserved for the Chassis         :    500 W

Fan Tray Statistics
FTC 0      Base power   Power Used
          :    300 W    43.04 W

FPC Statistics
          Base power   Power Used   PoE power   Priority
FPC 1    (EX6200-48P)   :    220 W    49.47 W    1440 W     1
FPC 2    (EX6200-48P)   :    220 W    47.20 W     800 W     2
FPC 3    (EX6200-48P)   :    220 W   1493.57 W    1440 W     0
FPC 4    (EX6200-SRE64-4XS) :    100 W    51.38 W     0 W       0
FPC 5    (EX6200-SRE64-4XS) :    100 W    50.28 W     0 W       0
FPC 6    (EX6200-48P)   :    220 W    49.38 W     800 W     6
FPC 8    (EX6200-48P)   :    220 W    61.41 W    1440 W     9

```

```

FPC 9 (EX6200-48T) : 150 W 12.49 W 0 W 9

Total (non-PoE) Power allocated : 1750 W
Total Power allocated for PoE : 5920 W
Power Available (Redundant case) : 5750 W
Total Power Available : 2515 W

```

Example output for an EX8200 switch:

```

PSU 0 (EX8200-AC2K) : 1200 W Online
PSU 1 (EX8200-AC2K) : 1200 W Online
PSU 2 (EX8200-AC2K) : 1200 W Online
PSU 3 (EX8200-AC2K) : 1200 W Online
Total Power supplied by all Online PSUs : 4800 W
Power Redundancy Configuration : N+1
Power Reserved for the Chassis : 1600 W

FPC Statistics Base power PoE power Priority
FPC 0 (EX8200-48T) : 350 W 0 W 2
FPC 1 (EX8200-2XS-40P) : 387 W 300 W 0
FPC 2 (EX8200-48PL) : 267 W 350 W 15
FPC 4 (EX8200-2XS-40P) : 387 W 300 W 1
FPC 5 (EX8200-48TL) : 230 W 0 W 15
FPC 6 (EX8200-48TL) : 230 W 0 W 15

Total (non-PoE) Power allocated : 3451 W
Total Power allocated for PoE : 950 W
Power Available (Redundant case) : 149 W
Total Power Available : 510 W

```

- Meaning**
- Example output for an EX6200 switch —The online power supplies can supply a total of 10,000 W to the switch. The switch is configured for *N+1* redundancy, which means 7500 W of redundant power can be supplied. The **Power Available (Redundant case)** field shows that the switch is meeting the *N+1* power requirements, with an additional 5750 W available. This value is calculated by subtracting all power allocations except PoE power allocations from redundant power (7500 W).

The total amount of power available on the switch is 2515 W. This value is calculated by subtracting all power allocations, including PoE power allocations, from the total power (10,000 W). On a switch with PoE line cards, if **Total Power Available** is 0, some or all of the PoE line cards might not be allocated their configured PoE power budgets, which means power to some or all PoE ports might be disabled.

The power priority order of the line cards, from highest priority line card to the lowest priority line card, is 4, 5, 3, 1, 2, 6, 8, 9. Slots 4 and 5, which contain the Switch Fabric and Routing Engine (SRE) modules, always have highest priority, even if a lower-numbered slot, such as slot 3 in this example, has a priority of 0. Should two or more 2500 W power supplies fail, power management will remove or reduce the PoE power allocations from the PoE line cards in the following order to balance the power budget: 8, 6, 2, 1, and 3.

The **Power Used** values for the fan tray and line cards shows the actual power being consumed for these components at the time the command was executed. These values are for your information only; power management uses allocated power, which is based on the maximum power the component might consume, and not actual power consumed, in determining its power budget.

- Example output for an EX8200 switch—The online power supplies can supply a total of 4800 W to the switch. The switch is configured for  $N+1$  redundancy, which means 3600 W of redundant power can be supplied. The **Power Available (Redundant case)** field shows that the switch is meeting the  $N+1$  power requirements, with an additional 149 W available. This value is calculated by subtracting all power allocations except PoE power allocations from redundant power (3600 W). Because 149 W is insufficient power for a line card, another line card cannot be added to the switch while maintaining  $N+1$  redundancy.

The total amount of power available on the switch is 510 W. This value is calculated by subtracting all power allocations, including PoE power allocations, from the total power (4800 W). On a switch with PoE line cards, if **Total Power Available** is 0, some or all of the PoE line cards might not be allocated their configured PoE power budgets, which means power to some or all PoE ports might be disabled.

The power priority order of the line cards, from highest priority line card to the lowest priority line card, is 1, 4, 0, 2, 5, 6. Should one or more 1200 W power supplies fail, power management will remove or reduce the PoE power allocations from the PoE line cards in the following order to balance the power budget: 2, 4, and 1.

- Related Documentation**
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 2512](#)
  - [Configuring the Power Priority of Line Cards \(CLI Procedure\)](#)

## Operational Commands

- [request system software nonstop-upgrade](#)
- [show chassis power-budget-statistics](#)
- [show vrrp](#)

## request system software nonstop-upgrade

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>request system software nonstop-upgrade (<i>package-name</i>   set [<i>package-name</i> <i>package-name</i>])</code><br><code>&lt;no-copy&gt;</code><br><code>&lt;no-old-master-upgrade&gt;</code><br><code>&lt;reboot&gt;</code><br><code>&lt;unlink&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b> | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Option <b>set [<i>package-name</i> <i>package-name</i>]</b> added in Junos OS Release 12.1 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D20 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>         | <p>Perform a nonstop software upgrade (NSSU) on a switch with redundant Routing Engines or on a Virtual Chassis. The behavior of this command depends on which switch or Virtual Chassis it is executed on:</p> <ul style="list-style-type: none"><li>• When you execute this command on an EX3300, EX4200, EX4300, EX4500, or EX4550 Virtual Chassis or QFX3500 and QFX3600 Virtual Chassis, a fixed configuration of switches in a Virtual Chassis Fabric (QFX3500/QFX3600 and QFX5100 switches) or for a mixed Virtual Chassis Fabric composed of any combination of QFX3500/QFX3600, QFX5100, and EX4300 switches, or a mixed Virtual Chassis composed of any combination of EX4200, EX4500, and EX4550 switches, all members are upgraded. The original Virtual Chassis backup becomes the master. The original master is automatically upgraded and rebooted and rejoins the Virtual Chassis as the backup after the upgrade completes.</li><li>• When you execute this command on an EX6200 or EX8200 switch, both the backup and master Routing Engines are upgraded, with the original backup Routing Engine becoming the new master at the end of the upgrade.<br/><br/>The original master Routing Engine is automatically rebooted on an EX6200 switch.<br/><br/>The original master Routing Engine is not automatically rebooted on an EX8200 switch unless you specify the <b>reboot</b> option.</li><li>• When you execute this command on an EX8200 Virtual Chassis, all master and backup Routing Engines are upgraded in the Virtual Chassis, including the external Routing Engines. The original backup Routing Engines become the new master Routing Engines. The original master Routing Engines are not automatically rebooted, unless you specify the <b>reboot</b> option.</li></ul> <p>This command has the following requirements:</p> <ul style="list-style-type: none"><li>• All Virtual Chassis members and all Routing Engines must be running the same Junos OS release.</li><li>• Graceful Routing Engine switchover (GRES) must be enabled.</li><li>• Nonstop active routing (NSR) must be enabled.</li></ul> |



**NOTE:** Although nonstop bridging (NSB) does not have to be enabled for you to use this command, we recommend that you enable NSB. Enabling NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU. See *Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)*.

- The command must be executed from the master Routing Engine on a standalone switch or from the master on a Virtual Chassis.
- For minimal traffic disruption, you must define link aggregation groups (LAGs) such that the member links reside on different Virtual Chassis members (for EX3300, EX4200, EX4300, EX4500, EX4550, QFX3500 and QFX3600 Virtual Chassis, and mixed Virtual Chassis, and Virtual Chassis Fabric) or on different line cards (for EX6200 and EX8200 switches, and for EX8200 Virtual Chassis).
- For EX3300, EX4200, EX4300, EX4500, EX4550, QFX3500 and QFX3600 Virtual Chassis, and mixed Virtual Chassis:
  - The Virtual Chassis members must be connected in a ring topology. A ring topology prevents the Virtual Chassis from splitting during an NSSU.
  - The Virtual Chassis master and backup must be adjacent to each other in the ring topology. Adjacency permits the master and backup to always be in sync, even when the switches in linecard roles are rebooting.
  - The Virtual Chassis must be preprovisioned so that the linecard role has been explicitly assigned to member switches acting in a linecard role. During an NSSU, the Virtual Chassis members must maintain their roles—the master and backup must maintain their Routing Engine roles (although mastership will change), and the remaining switches must maintain their linecard roles.
  - A two-member Virtual Chassis must have **no-split-detection** configured so that the Virtual Chassis does not split when an NSSU upgrades a member.
- For Virtual Chassis Fabric:
  - Only two pre-provisioned members in the routing engine role are supported. If more than two routing engines are configured, a warning will be issued, and NSSU will stop.
  - The Virtual Chassis Fabric members are connected in a spine and leaf topology. A spine and leaf topology prevents the Virtual Chassis Fabric from splitting during an NSSU. Each leaf device must be connected to both spine devices.
  - The Virtual Chassis Fabric must be preprovisioned so that the line card role has been explicitly assigned to member switches acting in a line card role, and that the routing engine role has been explicitly assigned to member switches acting in a routing engine role. During an NSSU, the Virtual Chassis Fabric members must maintain their roles—the master and backup must maintain their master and backup roles (although

mastership will change), the member switches must remain their routing engine roles, and the remaining switches must maintain their linecard roles.

- A two-member Virtual Chassis Fabric must have **no-split-detection** configured so that the Virtual Chassis Fabric does not split when an NSSU upgrades a member.

**Options** *package-name*—Location from which the software package or bundle is to be installed. For example:

- */var/tmp/package-name*—For a software package or bundle that is being installed from a local directory on the switch.
- *protocol://hostname/pathname/package-name*—For a software package or bundle that is to be downloaded and installed from a remote location. Replace **protocol** with one of the following:
  - **ftp**—File Transfer Protocol.  
Use *ftp://hostname/pathname/package-name*. To specify authentication credentials, use *ftp://<username>:<password>@hostname/pathname/package-name*. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
  - **http**—Hypertext Transfer Protocol.  
Use *http://hostname/pathname/package-name*. To specify authentication credentials, use *http://<username>:<password>@hostname/pathname/package-name*. If a password is required and you omit it, you are prompted for it.
  - **scp**—Secure copy (available only for Canada and U.S. version).  
Use *scp://hostname/pathname/package-name*. To specify authentication credentials, use *scp://<username>:<password>@hostname/pathname/package-name*.



**NOTE:** The *pathname* in the protocol is the relative path to the user home directory on the remote system and not the root directory.

---

**set** [*package-name package-name*]—(Mixed Virtual Chassis only) Locations of the EX4200 and the EX4500 installation packages. These packages must be for the same Junos OS release. See the description of the *package-name* option for information about how to specify the location of the installation packages.

**no-copy**—(Optional) Install a software package or bundle, but do not save copies of package or bundle files.

**no-old-master-upgrade**—(Optional) (EX8200 switches only) Upgrade the backup Routing Engine only. After the upgrade completes, the original master Routing Engine becomes the backup Routing Engine and continues running the previous software version.

**reboot**—(Optional) (EX8200 switches and EX8200 Virtual Chassis only) When the **reboot** option is included, the original master (new backup) Routing Engines are automatically rebooted after being upgraded to the new software. When the **reboot** option is not included, you must manually reboot the original master (new backup) Routing Engines using the [request system reboot](#) command.



**NOTE:** If you do not use the **reboot** option on an EX8200 Virtual Chassis, you must establish a connection to the console port on the Switch Fabric and Routing Engine (SRE) module or Routing Engine (RE) module to perform the manual reboot of the backup Routing Engines.

**unlink**—(Optional) Remove the software package after a successful upgrade is completed.

**Required Privilege Level** maintenance

**Related Documentation**

- *show chassis nonstop-upgrade*
- [Upgrading Software on an EX3300, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and Mixed Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 2508](#)
- *Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)*
- *Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)*
- *Upgrading Software on QFX3500, QFX3600, and QFX5100 Virtual Chassis Using Nonstop Software Upgrade*
- *Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade*

**List of Sample Output**

[request system software nonstop-upgrade \(EX4200 Virtual Chassis\) on page 2551](#)  
[request system software nonstop-upgrade \(EX6200 Switch\) on page 2553](#)  
[request system software nonstop-upgrade reboot \(EX8200 Switch\) on page 2554](#)  
[request system software nonstop-upgrade no-old-master-upgrade \(EX8200 Switch\) on page 2555](#)  
[request system software nonstop-upgrade reboot \(EX8200 Virtual Chassis\) on page 2555](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### [request system software nonstop-upgrade \(EX4200 Virtual Chassis\)](#)

```
user@switch> request system software nonstop-upgrade
/var/tmp/jinstall-ex-4200-12.1R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
```

Installing image on other FPC's along with the backup

Checking pending install on fpc1

Pushing bundle to fpc1

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc1

Checking pending install on fpc2

Pushing bundle to fpc2

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc2

Checking pending install on fpc3

Pushing bundle to fpc3

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc3

Checking pending install on fpc4

Pushing bundle to fpc4

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc4

Checking pending install on fpc5

Pushing bundle to fpc5

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc5

Checking pending install on fpc6

Pushing bundle to fpc6

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc6

Checking pending install on fpc7

Pushing bundle to fpc7

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc7

Backup upgrade done

Rebooting Backup RE

Rebooting fpc1

ISSU: Backup RE Prepare Done

Waiting for Backup RE reboot

GRES operational

Initiating Chassis In-Service-Upgrade

Chassis ISSU Started

ISSU: Preparing Daemons

ISSU: Daemons Ready for ISSU

ISSU: Starting Upgrade for FRUs

ISSU: Preparing for Switchover

ISSU: Ready for Switchover

Checking In-Service-Upgrade status

| Item  | Status | Reason |
|-------|--------|--------|
| FPC 0 | Online |        |
| FPC 1 | Online |        |



```

FPC 2      Online (ISSU)
FPC 3      Online (ISSU)
FPC 4      Online (ISSU)
FPC 5      Online (ISSU)
FPC 6      Online (ISSU)
FPC 7      Online (ISSU)
Going to install image on master
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE

*** FINAL System shutdown message from root@switch ***

System going down IMMEDIATELY

Shutdown NOW!
[pid 9336]

```

#### request system software nonstop-upgrade (EX6200 Switch)

```

{master}
user@switch> request system software nonstop-upgrade
/var/tmp/jinstall-ex-6200-12.2R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re0
NOTICE: Validating configuration against
jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re0
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status

```

| Item  | Status        | Reason |
|-------|---------------|--------|
| FPC 0 | Online (ISSU) |        |
| FPC 1 | Online (ISSU) |        |
| FPC 2 | Online (ISSU) |        |
| FPC 3 | Online (ISSU) |        |
| FPC 4 | Online        |        |
| FPC 5 | Online        |        |
| FPC 6 | Online (ISSU) |        |
| FPC 7 | Online (ISSU) |        |
| FPC 8 | Online (ISSU) |        |
| FPC 9 | Online (ISSU) |        |

```
Going to install image on master
NOTICE: Validating configuration against
jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE
Trying to relinquish mastership before rebooting...
Resolving mastership...
Complete. The other routing engine becomes the master.

*** FINAL System shutdown message from user@switch ***

System going down IMMEDIATELY
```

### request system software nonstop-upgrade reboot (EX8200 Switch)

```
{master}
user@switch> request system software nonstop-upgrade reboot
/var/tmp/jinstall-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 2         Offline          Offlined by CLI command
  FPC 3         Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
[pid 2635]

*** FINAL System shutdown message from user@switch ***
System going down IMMEDIATELY
```

**request system software nonstop-upgrade no-old-master-upgrade (EX8200 Switch)**

```
{master}
user@switch> request system software nonstop-upgrade no-old-master-upgrade
/var/tmp/jinstall-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 1         Online (ISSU)
  FPC 2         Online (ISSU)
  FPC 3         Offline
  FPC 4         Online (ISSU)
  FPC 5         Online (ISSU)
  FPC 6         Online (ISSU)
  FPC 7         Online (ISSU)
  Offlined by CLI command
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
Skipping Old Master Upgrade
ISSU: IDLE
```

**request system software nonstop-upgrade reboot (EX8200 Virtual Chassis)**

```
{master:9}
user@external-routing-engine> request system software nonstop-upgrade reboot
/var/tmp/jinstall-ex-xre200-11.1-20101130.0-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing LCC Backup REs
ISSU: Preparing Backup RE
Pushing bundle /var/tmp/jinstall-ex-xre200-11.1-20101130.0-domestic-signed.tgz
to member8
-----
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
VC Backup upgrade done
Rebooting VC Backup RE

Rebooting member8
ISSU: Backup RE Prepare Done
Waiting for VC Backup RE reboot
```

```
Pushing bundle to member0-backup
Pushing bundle to member1-backup
WARNING: A reboot is required to install the software
WARNING:   Use the 'request system reboot' command immediately
WARNING: A reboot is required to install the software
WARNING:   Use the 'request system reboot' command immediately
```

```
Rebooting member0-backup
Rebooting LCC [member0-backup]
```

```
Rebooting member1-backup
Rebooting LCC [member1-backup]
ISSU: LCC Backup REs Prepare Done
GRES operational
Initiating Chassis Nonstop-Software-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking Nonstop-Upgrade status
member0:
```

| Item  | Status        | Reason |
|-------|---------------|--------|
| FPC 0 | Online (ISSU) |        |
| FPC 1 | Online (ISSU) |        |
| FPC 2 | Online (ISSU) |        |
| FPC 5 | Online (ISSU) |        |

member1:

| Item  | Status        | Reason                 |
|-------|---------------|------------------------|
| FPC 0 | Online (ISSU) |                        |
| FPC 1 | Offline       | Offlined due to config |
| FPC 2 | Online (ISSU) |                        |
| FPC 3 | Online (ISSU) |                        |
| FPC 4 | Online (ISSU) |                        |
| FPC 5 | Online (ISSU) |                        |
| FPC 7 | Online (ISSU) |                        |

member0:

| Item  | Status        | Reason |
|-------|---------------|--------|
| FPC 0 | Online (ISSU) |        |
| FPC 1 | Online (ISSU) |        |
| FPC 2 | Online (ISSU) |        |
| FPC 5 | Online (ISSU) |        |

member1:

| Item  | Status        | Reason                 |
|-------|---------------|------------------------|
| FPC 0 | Online (ISSU) |                        |
| FPC 1 | Offline       | Offlined due to config |
| FPC 2 | Online (ISSU) |                        |
| FPC 3 | Online (ISSU) |                        |
| FPC 4 | Online (ISSU) |                        |
| FPC 5 | Online (ISSU) |                        |
| FPC 7 | Online (ISSU) |                        |

```
ISSU: Upgrading Old Master RE
Pushing bundle /var/tmp/incoming-package-8200.tgz to member0-master
```

Pushing bundle /var/tmp/incoming-package-8200.tgz to member1-master

ISSU: RE switchover Done

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Rebooting ...

shutdown: [pid 2188]

Shutdown NOW!

ISSU: Old Master Upgrade Done

ISSU: IDLE

Shutdown NOW!

\*\*\* FINAL System shutdown message from root@ \*\*\*

System going down IMMEDIATELY

## show chassis power-budget-statistics

|                                 |                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show chassis power-budget-statistics</b>                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                        |
| <b>Description</b>              | Display the power budget of an EX Series switch.                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Verifying Power Configuration and Use on page 2545</a></li> <li>• <a href="#">Configuring the Power Priority of Line Cards (CLI Procedure)</a></li> <li>• <a href="#">Configuring Power Supply Redundancy (CLI Procedure) on page 2512</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show chassis power-budget-statistics (EX6200 Switch) on page 2560</a><br><a href="#">show chassis power-budget-statistics (EX8200 Switch) on page 2560</a>                                                                                                                                     |
| <b>Output Fields</b>            | Table 256 on page 2558 lists the output fields for the <b>show chassis power-budget-statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                        |

**Table 256: show chassis power-budget-statistics Output Fields**

| Field Name                                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>PSU <i>n</i> (supply type)</b>              | Capacity rating of the power supply and whether the power supply is currently operating ( <b>Online</b> ) or not ( <b>Offline</b> ). If a power supply is offline, the capacity is shown as 0 W.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Total Power supplied by all Online PSUs</b> | Total number of watts supplied by all currently operating power supplies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Power Redundancy Configuration</b>          | Configured power redundancy setting, either <i>N</i> +1 or <i>N</i> + <i>N</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Base power reserved</b>                     | Total number of watts reserved for the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Non-PoE power being consumed</b>            | The amount of power, in W, currently being consumed for PoE.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Power Reserved for the Chassis</b>          | <p>Power reserved for the chassis:</p> <ul style="list-style-type: none"> <li>• For an EX6200 switch, 500 W.</li> <li>• For an EX8208 switch: 1600 W in an <i>N</i>+1 configuration; 1200 W in an <i>N</i>+<i>N</i> configuration</li> <li>• For an EX8216 switch: 2400 W in an <i>N</i>+1 configuration; 1800 W in an <i>N</i>+<i>N</i> configuration</li> </ul> <p>The power reserved for the chassis includes the maximum power requirements for the fan tray and Switch Fabric and Routing Engine (SRE), Routing Engine (RE), and Switch Fabric (SF) modules in both base and redundant configurations.</p> |

Table 256: show chassis power-budget-statistics Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fan Tray Statistics               | <p>(EX6200 switch only) Information about the fan tray:</p> <ul style="list-style-type: none"> <li>• <b>Base power</b>—Power allocated to the fan tray in the power budget. This allocation is included in <b>Power Reserved for the Chassis</b>.</li> <li>• <b>Power Used</b>—Actual power being used by the fan tray. This value is for informational purposes only: the power budget for the switch is based on allocated power (the theoretical maximum the fan tray might use) rather than used power.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| FPC <i>n</i> ( <i>card type</i> ) | <p>Information about the line card installed in slot <i>n</i>. For EX6200 switches, information about the SRE modules in slot 4 and slot 5 is also shown.</p> <ul style="list-style-type: none"> <li>• <b>Base power</b>—For line cards without PoE ports, the total power allocated to the line card.<br/>For line cards with PoE ports, the power allocated to the line card before the PoE power budget is allocated. The base power includes 37 W of PoE power that is always allocated to line cards that support PoE.</li> <li>• <b>Power Used</b>—(EX6200 switch only) The actual power being consumed by the line card or SRE module, including PoE power. This value is for informational purposes only: the power budget for the switch is based on allocated power (the theoretical maximum the line card might use) rather than used power.</li> <li>• <b>PoE power</b>—For line cards with PoE ports, the PoE power budget allocated to the line card. This value includes the 37 W of PoE power that is always part of the base power allocation for line cards that support PoE.<br/>For line cards without PoE ports, the value is always 0 W.</li> <li>• The power priority assigned to the line card slot.</li> </ul> |
| Total (non-PoE) Power allocated   | Power budgeted for all the components in the switch, excluding the PoE power budget allocated to line cards. This value is equal to the power reserved for the chassis plus the base power allocations of all online line cards.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Total Power allocated for PoE     | The total of the PoE power budgets allocated to the line cards in the switch. This figure includes the 37 W of PoE power always included in the base allocation for each line card that supports PoE.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Total PoE power consumed          | The amount of power that has been consumed by PoE.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Total PoE power remaining         | The amount of available power remaining that can be used for PoE.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Power Available (Redundant case)  | Unused power available to the switch in the power budget, not including the power reserved for redundancy. If power is insufficient to meet the <i>N</i> +1 or <i>N</i> + <i>N</i> redundancy requirements, this value is 0. PoE power allocations are not included in the calculation of this value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 256: show chassis power-budget-statistics Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                                                               |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Total Power Available</b> | Unused power available to the switch in the power budget. This value is derived by subtracting all power allocations, including PoE power allocations, from the total power available on the switch (the <b>Total Power supplied by all Online PSUs</b> value). |

## Sample Output

### show chassis power-budget-statistics (EX6200 Switch)

```

user@switch> show chassis power-budget-statistics
PSU 0      (EX6200-PWR-AC2500)      :   2500 W   Online
PSU 1      (EX6200-PWR-AC2500)      :   2500 W   Online
PSU 2      (EX6200-PWR-AC2500)      :   2500 W   Online
PSU 3      (EX6200-PWR-AC2500)      :   2500 W   Online
Total Power supplied by all Online PSUs : 10000 W
Power Redundancy Configuration        : N+1
Power Reserved for the Chassis        :   500 W

Fan Tray Statistics
FTC 0      :   300 W   43.04 W

FPC Statistics
Base power Power Used PoE power Priority
FPC 1 (EX6200-48P) :   220 W   49.47 W   1440 W   1
FPC 2 (EX6200-48P) :   220 W   47.20 W    800 W   2
FPC 3 (EX6200-48P) :   220 W  1493.57 W   1440 W   0
FPC 4 (EX6200-SRE64-4XS) :   100 W   51.38 W    0 W   0
FPC 5 (EX6200-SRE64-4XS) :   100 W   50.28 W    0 W   0
FPC 6 (EX6200-48P) :   220 W   49.38 W    800 W   6
FPC 8 (EX6200-48P) :   220 W   61.41 W   1440 W   9
FPC 9 (EX6200-48T) :   150 W   12.49 W    0 W   9

Total (non-PoE) Power allocated :   1750 W
Total Power allocated for PoE   :   5920 W
Power Available (Redundant case) :   5750 W
Total Power Available           :   2515 W

```

### show chassis power-budget-statistics (EX8200 Switch)

```

user@switch> show chassis power-budget-statistics
PSU 0      (EX8200-AC2K)      :   2000 W   Online
PSU 1      (EX8200-AC2K)      :   2000 W   Online
PSU 2      (EX8200-AC2K)      :   2000 W   Online
PSU 3      (EX8200-AC2K)      :   2000 W   online
PSU 4      (EX8200-AC2K)      :   2000 W   Online
Total Power supplied by all Online PSUs : 10000 W
Power Redundancy Configuration        : N+1
Power Reserved for the Chassis        :   2400 W

FPC Statistics
Base power PoE power Priority
FPC 1 (EX8200-48T) :   350 W    0 W   15
FPC 5 (EX8200-2XS-40P) :   387 W   792 W   0
FPC 9 (EX8200-48PL) :   267 W   915 W  15
FPC 10 (EX8200-2XS-40T) :   350 W    0 W   1
FPC 12 (EX8200-48T) :   350 W    0 W  15

Total (non-PoE) Power allocated :   4104 W

```



|                                  |   |        |
|----------------------------------|---|--------|
| Total Power allocated for PoE    | : | 1707 W |
| Power Available (Redundant case) | : | 3896 W |
| Total Power Available            | : | 4263 W |

## show vrrp

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show vrrp &lt;brief   detail   extensive   summary&gt; &lt;interface <i>interface-name</i>&gt; &lt;track interfaces&gt;</pre>                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Display information and status about VRRP groups.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>none</b>—(Same as brief) Display brief status information about all VRRP interfaces.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display information and status about the specified VRRP interface.</p> <p><b>track interfaces</b>—(Optional) Display information and status about VRRP track interfaces.</p>                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring VRRP for IPv6 (CLI Procedure) on page 2513</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>    | <a href="#">show vrrp on page 2567</a><br><a href="#">show vrrp brief on page 2567</a><br><a href="#">show vrrp detail (IPv6) on page 2567</a><br><a href="#">show vrrp detail (Route Track) on page 2568</a><br><a href="#">show vrrp extensive on page 2568</a><br><a href="#">show vrrp interface on page 2569</a><br><a href="#">show vrrp summary on page 2570</a><br><a href="#">show vrrp track detail on page 2570</a><br><a href="#">show vrrp track summary on page 2571</a> |
| <b>Output Fields</b>            | <a href="#">Table 257 on page 2562</a> lists the output fields for the <b>show vrrp</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                               |

**Table 257: show vrrp Output Fields**

| Field Name             | Field Description                                                            | Level of Output                        |
|------------------------|------------------------------------------------------------------------------|----------------------------------------|
| <b>Interface</b>       | Name of the logical interface.                                               | <b>none, brief, extensive, summary</b> |
| <b>Interface index</b> | Physical interface index number, which reflects its initialization sequence. | <b>extensive</b>                       |
| <b>Groups</b>          | Total number of VRRP groups configured on the interface.                     | <b>extensive</b>                       |

Table 257: show vrrp Output Fields (*continued*)

| Field Name                                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output                       |
|--------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| <b>Active</b>                              | Total number of VRRP groups that are active (that is, whose interface state is either up or down).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>                      |
| <b>Interface VRRP PDU statistics</b>       | Nonerrored statistics for the logical interface: <ul style="list-style-type: none"> <li>• <b>Advertisement sent</b>—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted.</li> <li>• <b>Advertisement received</b>—Number of VRRP advertisement PDUs received by the interface.</li> <li>• <b>Packets received</b>—Number of VRRP packets received for VRRP groups on the interface.</li> <li>• <b>No group match received</b>—Number of VRRP packets received for VRRP groups that do not exist on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>                      |
| <b>Interface VRRP PDU error statistics</b> | Errored statistics for the logical interface: <ul style="list-style-type: none"> <li>• <b>Invalid IPAH next type received</b>—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets.</li> <li>• <b>Invalid VRRP ttl value received</b>—Number of packets received whose IP time-to-live (TTL) value is not 255.</li> <li>• <b>Invalid VRRP version received</b>—Number of packets received whose VRRP version is not 2.</li> <li>• <b>Invalid VRRP pdu type received</b>—Number of packets received whose VRRP PDU type is not 1.</li> <li>• <b>Invalid VRRP authentication type received</b>—Number of packets received whose VRRP authentication is not none, simple, or md5.</li> <li>• <b>Invalid VRRP IP count received</b>—Number of packets received whose VRRP IP count exceeds 8.</li> <li>• <b>Invalid VRRP checksum received</b>—Number of packets received whose VRRP checksum does not match the calculated value.</li> </ul> | <b>extensive</b>                      |
| <b>Physical interface</b>                  | Name of the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail, extensive</b>              |
| <b>Unit</b>                                | Logical unit number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels                            |
| <b>Address</b>                             | Address of the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>none, brief, detail, extensive</b> |
| <b>Index</b>                               | Physical interface index number, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail, extensive</b>              |
| <b>SNMP ifIndex</b>                        | SNMP index number for the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail, extensive</b>              |
| <b>VRRP-Traps</b>                          | Status of VRRP traps: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail, extensive</b>              |

Table 257: show vrrp Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output                 |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Type and Address</b>             | Identifier for the address and the address itself: <ul style="list-style-type: none"> <li>• <b>lcl</b>—Configured local interface address.</li> <li>• <b>mas</b>—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router.</li> <li>• <b>vip</b>—Configured virtual IP addresses.</li> </ul>                                                                                                                                                                                                                                                                                         | none, brief, summary            |
| <b>Interface state or Int state</b> | State of the physical interface: <ul style="list-style-type: none"> <li>• <b>down</b>—The device is present and the link is unavailable.</li> <li>• <b>not present</b>—The interface is configured, but no physical device is present.</li> <li>• <b>unknown</b>—The VRRP process has not had time to query the kernel about the state of the interface.</li> <li>• <b>up</b>—The device is present and the link is established.</li> </ul>                                                                                                                                                                                                            | none, brief, extensive, summary |
| <b>Group</b>                        | VRRP group number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | none, brief, extensive, summary |
| <b>State</b>                        | VRRP state: <ul style="list-style-type: none"> <li>• <b>backup</b>—The interface is acting as the backup router interface.</li> <li>• <b>bringup</b>—VRRP is just starting, and the physical device is not yet present.</li> <li>• <b>idle</b>—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established.</li> <li>• <b>initializing</b>—VRRP is initializing.</li> <li>• <b>master</b>—The interface is acting as the master router interface.</li> <li>• <b>transition</b>—The interface is changing between being the backup and being the master router.</li> </ul> | extensive                       |
| <b>Priority</b>                     | Configured VRRP priority for the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | detail, extensive               |
| <b>Advertisement interval</b>       | Configured VRRP advertisement interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | detail, extensive               |
| <b>Authentication type</b>          | Configured VRRP authentication type: <b>none</b> , <b>simple</b> , or <b>md5</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | detail, extensive               |
| <b>Preempt</b>                      | Whether preemption is allowed on the interface: <b>yes</b> or <b>no</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | detail, extensive               |
| <b>Accept-data mode</b>             | Whether the interface is configured to accept packets destined for the virtual IP address: <b>yes</b> or <b>no</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | detail, extensive               |
| <b>VIP count</b>                    | Number of virtual IP addresses that have been configured on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | detail, extensive               |
| <b>VIP</b>                          | List of virtual IP addresses configured on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | detail, extensive               |
| <b>Advertisement timer</b>          | Time until the advertisement timer expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | detail, extensive               |

Table 257: show vrrp Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                              | Level of Output   |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| Master router               | IP address of the interface that is acting as the master. If the VRRP interface is down, the output is <b>N/A</b> .                                                                                                                                                                                            | detail, extensive |
| Virtual router uptime       | Time that the virtual router has been up.                                                                                                                                                                                                                                                                      | detail, extensive |
| Master router uptime        | Time that the master router has been up.                                                                                                                                                                                                                                                                       | detail, extensive |
| Virtual MAC                 | MAC address associated with the virtual IP address.                                                                                                                                                                                                                                                            | detail, extensive |
| Tracking                    | Whether tracking is <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                                                        | detail, extensive |
| Current priority            | Current operational priority for being the VRRP master.                                                                                                                                                                                                                                                        | detail, extensive |
| Configured priority         | Configured base priority for being the VRRP master.                                                                                                                                                                                                                                                            | detail, extensive |
| Priority hold-time          | Minimum time interval, in seconds, between successive changes to the current priority. <b>Disabled</b> indicates no minimum interval.                                                                                                                                                                          | detail, extensive |
| Remaining-time              | ( <b>track</b> option only) Displays the time remaining in the priority hold-time interval.                                                                                                                                                                                                                    | detail            |
| Interface tracking          | Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.                                                                                                                                                                                    | detail extensive  |
| Interface/Tracked interface | Name of the tracked interface.                                                                                                                                                                                                                                                                                 | detail extensive  |
| Int state/Interface state   | Current operational state of the tracked interface: <b>up</b> or <b>down</b> .                                                                                                                                                                                                                                 | detail, extensive |
| Int speed/Speed             | Current operational speed, in bits per second, of the tracked interface.                                                                                                                                                                                                                                       | detail, extensive |
| Incurred priority cost      | Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.                                                                                                                                       | detail, extensive |
| Threshold                   | Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred.<br><br>An entry of <b>down</b> means that the corresponding priority cost is incurred when the interface is down. | detail, extensive |
| Route tracking              | Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.                                                                                                                                                                                            | detail, extensive |
| Route count                 | The number of routes being tracked.                                                                                                                                                                                                                                                                            | detail, extensive |

Table 257: show vrrp Output Fields (*continued*)

| Field Name                               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output          |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>Route</b>                             | The IP address of the route being tracked.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail, extensive</b> |
| <b>VRF name</b>                          | The VPN routing and forwarding (VRF) routing instance that the tracked route is in.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail, extensive</b> |
| <b>Route state</b>                       | The state of the route being tracked: <b>up</b> , <b>down</b> , or <b>unknown</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail, extensive</b> |
| <b>Priority cost</b>                     | Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail, extensive</b> |
| <b>Active</b>                            | Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail, extensive</b> |
| <b>Group VRRP PDU statistics</b>         | Number of VRRP advertisements sent and received by the group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>extensive</b>         |
| <b>Group VRRP PDU error statistics</b>   | <p>Errored statistics for the VRRP group:</p> <ul style="list-style-type: none"> <li>• <b>Bad authentication type received</b>—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be <b>none</b>, <b>simple</b>, or <b>md5</b> and must be the same for all routers in the VRRP group.</li> <li>• <b>Bad password received</b>—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group</li> <li>• <b>Bad MD5 digest received</b>—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router.</li> <li>• <b>Bad advertisement timer received</b>—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group.</li> <li>• <b>Bad VIP count received</b>—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance.</li> <li>• <b>Bad VIPADDR received</b>—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance.</li> </ul> | <b>extensive</b>         |
| <b>Group state transition statistics</b> | <p>State transition statistics for the VRRP group:</p> <ul style="list-style-type: none"> <li>• <b>Idle to master transitions</b>—Number of times that the VRRP instance transitioned from the idle state to the master state.</li> <li>• <b>Idle to backup transitions</b>—Number of times that the VRRP instance transitioned from the idle state to the backup state.</li> <li>• <b>Backup to master transitions</b>—Number of times that the VRRP instance transitioned from the backup state to the master state.</li> <li>• <b>Master to backup transitions</b>—Number of times that the VRRP instance transitioned from the master state to the backup state.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>extensive</b>         |
| <b>Vlan-id</b>                           | ID of Vlan                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>            |

Table 257: show vrrp Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| VR state   | VRRP information: <ul style="list-style-type: none"> <li>• <b>backup</b>—The interface is acting as the backup router interface.</li> <li>• <b>bringup</b>—VRRP is just starting, and the physical device is not yet present.</li> <li>• <b>idle</b>—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established.</li> <li>• <b>initializing</b>—VRRP is initializing.</li> <li>• <b>master</b>—The interface is acting as the master router interface.</li> <li>• <b>transition</b>—The interface is changing between being the backup and being the master router.</li> </ul> | none, brief     |
| Timer      | VRRP timer information: <ul style="list-style-type: none"> <li>• <b>A</b>—Time, in seconds, until the advertisement timer expires.</li> <li>• <b>D</b>—Time, in seconds, until the Master is Dead timer expires.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                  | none, brief     |

## Sample Output

### show vrrp

```

user@host> show vrrp
Interface      State      Group  VR state  Timer  Type  Address
ge-0/0/0.121   up         1      master    A 1.052 1c1  gec0::12:1:1:1
                                     vip  ge80::12:1:1:99
                                     vip  gec0::12:1:1:99
ge-0/0/2.131   up         1      master    A 0.364 1c1  gec0::13:1:1:1
                                     vip  ge80::13:1:1:99
                                     vip  gec0::13:1:1:99

```

### show vrrp brief

The output for the **show vrrp brief** command is identical to that for the **show vrrp** command. For sample output, see [show vrrp on page 2567](#).

### show vrrp detail (IPv6)

```

user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 1.121s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

```

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled  
 Interface state: up, Group: 1, State: master  
 Priority: 200, Advertisement interval: 1, Authentication type: none  
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,  
 gec0::13:1:1:99  
 Advertisement timer: 0.327s, Master router: ge80::13:1:1:1  
 Virtual router uptime: 00:03:47, Master router uptime: 00:03:41  
 Virtual MAC: 00:00:5e:00:02:01  
 Tracking: disabled

### show vrrp detail (Route Track)

user@host> show vrrp detail

Physical interface: ge-1/1/0, Unit: 0, Address: 30.30.30.30/24  
 Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled  
 Interface state: up, Group: 100, State: master  
 Priority: 150, Advertisement interval: 1, Authentication type: none  
 Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100  
 Advertisement timer: 1.218s, Master router: 30.30.30.30  
 Virtual router uptime: 00:04:28, Master router uptime: 00:00:13  
 Virtual MAC: 00:00:5e:00:01:64  
 Tracking: enabled  
 Current priority: 150, Configured priority: 150  
 Priority hold-time: disabled  
 Interface tracking: disabled  
 Route tracking: enabled, Route count: 1

| Route           | VRF name | Route state | Priority cost |
|-----------------|----------|-------------|---------------|
| 192.168.40.0/22 | default  | up          | 30            |

### show vrrp extensive

user@host> show vrrp extensive

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1

#### Interface VRRP PDU statistics

|                         |   |     |
|-------------------------|---|-----|
| Advertisement sent      | : | 188 |
| Advertisement received  | : | 0   |
| Packets received        | : | 0   |
| No group match received | : | 0   |

#### Interface VRRP PDU error statistics

|                                           |   |   |
|-------------------------------------------|---|---|
| Invalid IPAH next type received           | : | 0 |
| Invalid VRRP TTL value received           | : | 0 |
| Invalid VRRP version received             | : | 0 |
| Invalid VRRP PDU type received            | : | 0 |
| Invalid VRRP authentication type received | : | 0 |
| Invalid VRRP IP count received            | : | 0 |
| Invalid VRRP checksum received            | : | 0 |

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled  
 Interface state: up, Group: 1, State: master  
 Priority: 200, Advertisement interval: 1, Authentication type: none  
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,  
 gec0::12:1:1:99  
 Advertisement timer: 1.034s, Master router: ge80::12:1:1:1  
 Virtual router uptime: 00:04:04, Master router uptime: 00:03:58  
 Virtual MAC: 00:00:5e:00:02:01  
 Tracking: disabled  
 Group VRRP PDU statistics



```

    Advertisement sent           :          188
    Advertisement received       :           0
Group VRRP PDU error statistics
    Bad authentication type received:         0
    Bad password received          :           0
    Bad MD5 digest received        :           0
    Bad advertisement timer received:         0
    Bad VIP count received         :           0
    Bad VIPADDR received          :           0
Group state transition statistics
    Idle to master transitions      :           0
    Idle to backup transitions     :           1
    Backup to master transitions   :           1
    Master to backup transitions   :           0

Interface: ge-0/0/2.131, Interface index: 69, Groups: 1, Active : 1
Interface VRRP PDU statistics
    Advertisement sent           :          186
    Advertisement received       :           0
    Packets received             :           0
    No group match received      :           0
Interface VRRP PDU error statistics
    Invalid IPAH next type received :           0
    Invalid VRRP TTL value received :           0
    Invalid VRRP version received  :           0
    Invalid VRRP PDU type received :           0
    Invalid VRRP authentication type received:         0
    Invalid VRRP IP count received :           0
    Invalid VRRP checksum received :           0

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
Advertisement timer: 0.396s, Master router: ge80::13:1:1:1
Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
    Advertisement sent           :          186
    Advertisement received       :           0
Group VRRP PDU error statistics
    Bad authentication type received:         0
    Bad password received          :           0
    Bad MD5 digest received        :           0
    Bad advertisement timer received:         0
    Bad VIP count received         :           0
    Bad VIPADDR received          :           0
Group state transition statistics
    Idle to master transitions      :           0
    Idle to backup transitions     :           1
    Backup to master transitions   :           1
    Master to backup transitions   :           0

```

#### show vrrp interface

```
user@host> show vrrp interface
```

```

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
Interface VRRP PDU statistics
  Advertisement sent           :          205
  Advertisement received       :           0
  Packets received             :           0
  No group match received      :           0
Interface VRRP PDU error statistics
  Invalid IPAH next type received :           0
  Invalid VRRP TTL value received :           0
  Invalid VRRP version received  :           0
  Invalid VRRP PDU type received :           0
  Invalid VRRP authentication type received:           0
  Invalid VRRP IP count received :           0
  Invalid VRRP checksum received :           0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 0.789s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:04:26, Master router uptime: 00:04:20
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
  Advertisement sent           :          205
  Advertisement received       :           0
Group VRRP PDU error statistics
  Bad authentication type received:           0
  Bad password received         :           0
  Bad MD5 digest received       :           0
  Bad advertisement timer received:           0
  Bad VIP count received        :           0
  Bad VIPADDR received          :           0
Group state transition statistics
  Idle to master transitions     :           0
  Idle to backup transitions     :           1
  Backup to master transitions   :           1
  Master to backup transitions   :           0

```

### show vrrp summary

```

user@host> show vrrp summary

```

| Interface  | State | Group | VR state | Type | Address     |
|------------|-------|-------|----------|------|-------------|
| ge-4/1/0.0 | up    | 1     | backup   | lcl  | 10.57.0.2   |
|            |       |       |          | vip  | 10.57.0.100 |

### show vrrp track detail

```

user@host> show vrrp track detail
Tracked interface: ae1.211
State: up, Speed: 400m
Incurred priority cost: 0

```

| Threshold | Priority cost | Active |
|-----------|---------------|--------|
| 400m      | 10            |        |
| 300m      | 60            |        |
| 200m      | 110           |        |
| 100m      | 160           |        |
| down      | 190           |        |

```
Tracking VRRP interface: ae0.210, Group: 1
VR State: master
Current priority: 200, Configured priority: 200
Priority hold-time: disabled,    Remaining-time: 50.351
```

#### show vrrp track summary

```
user@host> show vrrp track summary
```

| Track if | State | Speed | VRRP if | Group | VR State | Current priority |
|----------|-------|-------|---------|-------|----------|------------------|
| ae1.211  | up    | 400m  | ae0.210 | 1     | master   | 200              |



# Troubleshooting Procedures

- [Tracing Nonstop Active Routing Synchronization Events on page 2573](#)

## Tracing Nonstop Active Routing Synchronization Events

---

To track the progress of nonstop active routing synchronization between Routing Engines, you can configure nonstop active routing trace options flags for each supported protocol and for BFD sessions and record these operations to a log file.

To configure nonstop active routing trace options for supported routing protocols, include the **nsr-synchronization** statement at the **[edit protocols protocol-name traceoptions flag]** hierarchy level and optionally specify one or more of the **detail**, **disable**, **receive**, and **send** options:

```
[edit protocols]
bgp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
isis {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
ldp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
mpls {
  traceoptions {
    flag nsr-synchronization;
    flag nsr-synchronization-detail;
  }
}
msdp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
(ospf | ospf3) {
```

```
traceoptions {  
    flag nsr-synchronization <detail> <disable> <receive> <send>;  
}  
}  
rip {  
    traceoptions {  
        flag nsr-synchronization <detail> <disable> <receive> <send>;  
    }  
}  
ripng {  
    traceoptions {  
        flag nsr-synchronization <detail> <disable> <receive> <send>;  
    }  
}  
pim {  
    traceoptions {  
        flag nsr-synchronization <detail> <disable> <receive> <send>;  
    }  
}
```

To configure nonstop active routing trace options for BFD sessions, include the **nsr-synchronization** and **nsr-packet** statements at the **[edit protocols bfd traceoptions flag]** hierarchy level.

```
[edit protocols]  
bfd {  
    traceoptions {  
        flag nsr-synchronization;  
        flag nsr-packet;  
    }  
}
```

To trace the Layer 2 VPN signaling state replicated from routes advertised by BGP, include the **nsr-synchronization** statement at the **[edit routing-options traceoptions flag]** hierarchy level. This flag also traces the label and logical interface association that VPLS receives from the kernel replication state.

```
[edit routing-options]  
traceoptions {  
    flag nsr-synchronization;  
}
```

**Related  
Documentation**

- [Configuring Nonstop Active Routing](#)
- [Configuring Nonstop Active Routing on Switches on page 2505](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 2514](#)
- [Configuring Nonstop Active Routing](#)

## PART 14

# Interfaces

- [Overview on page 2577](#)
- [Configuration on page 2615](#)
- [Administration on page 2835](#)
- [Troubleshooting Procedures on page 2931](#)





## CHAPTER 44

# Overview

- [Interfaces Overview on page 2577](#)

## Interfaces Overview

---

- [EX Series Switches Interfaces Overview on page 2577](#)
- [Understanding Interface Naming Conventions on EX Series Switches on page 2580](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2582](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2585](#)
- [Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces on page 2590](#)
- [Understanding Local Link Bias on page 2590](#)
- [Understanding Layer 3 Subinterfaces on page 2591](#)
- [Understanding Unicast RPF on page 2592](#)
- [Understanding IP Directed Broadcast for EX Series Switches on page 2596](#)
- [Understanding Interface Ranges on EX Series Switches on page 2598](#)
- [Understanding Multichassis Link Aggregation on page 2599](#)
- [802.1Q VLANs Overview on page 2614](#)

## EX Series Switches Interfaces Overview

Juniper Networks EX Series Ethernet Switches have two types of interfaces: network interfaces and special interfaces. This topic provides brief information about these interfaces. For additional information, see the *[Junos OS Interfaces Fundamentals Configuration Guide](#)*.

For information about interface-naming conventions on EX Series switches, see “[Understanding Interface Naming Conventions on EX Series Switches](#)” on page 2580.

This topic describes:

- [Network Interfaces on page 2578](#)
- [Special Interfaces on page 2578](#)

## Network Interfaces

Network interfaces connect to the network and carry network traffic. [Table 258 on page 2578](#) lists the types of network interfaces supported on EX Series switches.

**Table 258: Network Interface Types and Purposes**

| Type                                 | Purpose                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aggregated Ethernet interfaces       | All EX Series switches allow you to group Ethernet interfaces at the physical layer to form a single link layer interface, also known as a <i>link aggregation group (LAG)</i> or <i>bundle</i> . These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.                                                                                                |
| LAN access interfaces                | Use these EX Series switch interfaces to connect a personal computer, laptop, file server, or printer to the network. When you power on an EX Series switch and use the factory-default configuration, the software automatically configures interfaces in access mode for each of the network ports. The default configuration also enables autonegotiation for both speed and link mode.       |
| Power over Ethernet (PoE) interfaces | EX Series switches provide PoE network ports with various switch models. These ports can be used to connect voice over IP (VoIP) telephones, wireless access points, video cameras, and point-of-sale devices to safely receive power from the same access ports that are used to connect personal computers to the network. PoE interfaces are enabled by default in the factory configuration. |
| Trunk interfaces                     | EX Series access switches can be connected to a distribution switch or customer-edge (CE) switches or routers. To use a port for this type of connection, you must explicitly configure the network interface for trunk mode. The interfaces from the distribution switch or CE switch to the access switches must also be configured for trunk mode.                                            |

## Special Interfaces

[Table 259 on page 2578](#) lists the types of special interfaces supported on EX Series switches.

**Table 259: Special Interface Types and Purposes**

| Type                 | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Console port         | Each EX Series switch has a serial port, labeled <b>CON</b> or <b>CONSOLE</b> , for connecting tty-type terminals to the switch using standard PC-type tty cables. The console port does not have a physical address or IP address associated with it. However, it is an interface in the sense that it provides access to the switch. On an EX3300 Virtual Chassis, an EX4200 Virtual Chassis, or an EX4500 Virtual Chassis, you can access the master and configure all members of the Virtual Chassis through any member's console port. For more information about the console port in a Virtual Chassis, see <a href="#">"Understanding Global Management of a Virtual Chassis" on page 5080</a> . |
| Loopback             | All EX Series switches have this software-only virtual interface that is always up. The loopback interface provides a stable and consistent interface and IP address on the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Management interface | The Juniper Networks Junos operating system (Junos OS) for EX Series switches automatically creates the switch's management Ethernet interface, <b>me0</b> . The management Ethernet interface provides an out-of-band method for connecting to the switch. To use <b>me0</b> as a management port, you must configure its logical port, <b>me0.0</b> , with a valid IP address. You can connect to the management interface over the network using utilities such as SSH or Telnet. SNMP can use the management interface to gather statistics from the switch. (The management interface <b>me0</b> is analogous to the <b>fxp0</b> interfaces on routers running Junos OS.)                          |

Table 259: Special Interface Types and Purposes (*continued*)

| Type                                                                           | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Integrated Routing and Bridging (IRB) Interface or Routed VLAN Interface (RVI) | <p>EX Series switches use an integrated routing and bridging (IRB) interface or Routed VLAN Interface (RVI) to route traffic from one broadcast domain to another and to perform other Layer 3 functions such as traffic engineering. These functions are typically performed by a router interface in a traditional network.</p> <p>The IRB interface or RVI functions as a logical router, eliminating the need for having both a switch and a router. These interfaces must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed from.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Virtual Chassis port (VCP) interfaces                                          | <p>Virtual Chassis ports (VCPs) are used to interconnect switches in a Virtual Chassis:</p> <ul style="list-style-type: none"> <li>EX3300 switches—Port 2 and port 3 of the SFP+ uplink ports are preconfigured as VCPs and can be used to interconnect up to six EX3300 switches in an EX3300 Virtual Chassis. See <a href="#">“Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)”</a> on page 5109.</li> <li>EX4200 and EX4500 switches—Each EX4200 switch or each EX4500 switch with a Virtual Chassis module installed has two dedicated VCPs on its rear panel. These ports can be used to interconnect up to ten EX4200 switches in an EX4200 Virtual Chassis, up to ten EX4500 switches in an EX4500 Virtual Chassis, and up to ten switches in a mixed EX4200 and EX4500 Virtual Chassis. When you power on switches that are interconnected in this manner, the software automatically configures the VCP interfaces for the dedicated ports that have been interconnected. These VCP interfaces are not configurable or modifiable. See <i>Understanding the High-Speed Interconnection of the Dedicated Virtual Chassis Ports Connecting EX4200, EX4500, and EX4550 Member Switches</i>.</li> </ul> <p>You can also interconnect EX4200 and EX4500 switches by using uplink module ports. Using uplink ports allows you to connect switches over longer distances than you can by using the dedicated VCPs. To use the uplink ports as VCPs, you must explicitly configure the uplink module ports on the members you want to connect as VCPs. See <a href="#">“Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)”</a> on page 5109 or <a href="#">Setting an Uplink Port as a Virtual Chassis Port on an EX4500 or EX4550 Switch (CLI Procedure)</a>.</p> <ul style="list-style-type: none"> <li>EX4300 switches—All QSFP+ ports are configured as VCPs, by default. See <a href="#">“Understanding EX4300 Virtual Chassis”</a> on page 5065.</li> </ul> <p>You can also interconnect EX4300 switches into a Virtual Chassis by using SFP+ uplink module ports as VCPs. Using uplink ports as VCPs allows you to connect switches over longer distances than you can by using the QSFP+ ports as VCPs. To use the uplink ports as VCPs, you must explicitly configure the uplink module ports on the members you want to connect as VCPs. See <a href="#">“Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)”</a> on page 5109.</p> <ul style="list-style-type: none"> <li>EX8200 switches—EX8200 switches can be connected to an XRE200 External Routing Engine to create an EX8200 Virtual Chassis. The XRE200 External Routing Engine has dedicated VCPs that connect to ports on the internal Routing Engines of the EX8200 switches and can connect to another XRE200 External Routing Engine for redundancy. These ports require no configuration.</li> </ul> <p>You can also connect two members of an EX8200 Virtual Chassis so that they can exchange Virtual Chassis Control Protocol (VCCP) traffic. To do so, you explicitly configure network ports on the EX8200 switches as VCPs. See <i>Understanding Virtual Chassis Ports in an EX8200 Virtual Chassis</i>.</p> |
| Virtual management Ethernet (VME) interface                                    | <p>EX3300, EX4200, EX4300, and EX4500 switches have a VME interface. This is a logical interface that is used for Virtual Chassis configurations and allows you to manage all the members of the Virtual Chassis through the master. For more information about the VME interface, see <a href="#">“Understanding Global Management of a Virtual Chassis”</a> on page 5080.</p> <p>EX8200 switches do not use a VME interface. An EX8200 Virtual Chassis is managed through the management Ethernet (<b>me0</b>) interface on the XRE200 External Routing Engine.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Related Documentation**

- [EX2200 Switches Hardware Overview](#)
- [EX3200 Switches Hardware Overview](#)
- [EX3300 Switches Hardware Overview](#)
- [EX4200 Switches Hardware Overview](#)
- [EX4300 Switches Hardware Overview](#)
- [EX4500 Switches Hardware Overview](#)
- [EX6210 Switch Hardware Overview](#)
- [EX8208 Switch Hardware Overview](#)
- [EX8216 Switch Hardware Overview](#)
- [XRE200 External Routing Engine Hardware Overview](#)
- [Understanding PoE on EX Series Switches on page 4423](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2582](#)
- [Understanding Layer 3 Subinterfaces on page 2591](#)

## Understanding Interface Naming Conventions on EX Series Switches

Juniper Networks EX Series Ethernet Switches use a naming convention for defining the interfaces that is similar to that of other platforms running under Juniper Networks Junos operating system (Junos OS). This topic provides brief information about the naming conventions used for interfaces on EX Series switches. For additional information, see the [Junos OS Network Interfaces Configuration Guide](#).

This topic describes:

- [Physical Part of an Interface Name on page 2580](#)
- [Logical Part of an Interface Name on page 2582](#)
- [Wildcard Characters in Interface Names on page 2582](#)

### Physical Part of an Interface Name

Network interfaces in Junos OS are specified as follows:

*type-fpc / pic / port*

EX Series switches apply this convention as follows:

- *type*—EX Series interfaces use the following media types:
  - **ge**—Gigabit Ethernet interface
  - **xe**—10 Gigabit Ethernet interface
  - **et**—40 Gigabit Ethernet interface
- *fpc*—Flexible PIC Concentrator. EX Series interfaces use the following convention for the FPC number in interface names:

- On an EX2200 switch, an EX3200 switch, a standalone EX3300 switch, a standalone EX4200 switch, a standalone EX4300 switch, a standalone EX4500, and a standalone EX4550 switch, FPC refers to the switch itself. The FPC number is **0** by default on these switches.
- On an EX3300 Virtual Chassis, an EX4200 Virtual Chassis, an EX4300 Virtual Chassis, an EX4500 Virtual Chassis, an EX4550 Virtual Chassis, or a mixed Virtual Chassis, the FPC number indicates the member ID of the switch in the Virtual Chassis.
- On an EX6200 switch and a standalone EX8200 switch, the FPC number indicates the slot number of the line card that contains the physical interface. On an EX6200 switch, the FPC number also indicates the slot number of the Switch Fabric and Routing Engine (SRE) module that contains the uplink port.
- On an EX8200 Virtual Chassis, the FPC number indicates the slot number of the line card on the Virtual Chassis. The line card slots on Virtual Chassis member 0 are numbered 0 through 15; on member 1, they are numbered 16 through 31, and so on.
- *pic*—EX Series interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:
  - On EX2200, EX3200, EX3300, EX4200, EX4500 switch, and EX4550 switches, the PIC number is **0** for all built-in interfaces (interfaces that are not uplink ports).
  - On EX2200, EX3200, EX3300, and EX4200 switches, the PIC number is **1** for uplink ports.
  - On EX4300 switches, the PIC number is **0** for built-in network ports, **1** for built-in QSFP+ ports (located on the rear panel of the switch), and **2** for uplink module ports.
  - On EX4500 switches, the PIC number is **1** for ports on the left-hand uplink module and **2** for ports on the right-hand uplink module.
  - On EX4550 switches, the PIC number is **1** for ports in the expansion module or Virtual Chassis module installed in the module slot on the front panel of the switch and **2** for those in the expansion module or Virtual Chassis module installed in the module slot on the rear panel of the switch.
  - On EX6200 and EX8200 switches, the PIC number is always **0**.
- *port*—EX Series interfaces use the following convention for port numbers:
  - On EX2200, EX3200, EX3300, EX4200, EX4300, EX4500, and EX4550 switches, built-in network ports are numbered from left to right. On models that have two rows of ports, the ports on the top row start with **0** followed by the remaining even-numbered ports, and the ports on the bottom row start with **1** followed by the remaining odd-numbered ports.
  - Uplink ports in EX2200, EX3200, EX3300, EX4200, EX4300, EX4500, and EX4550 switches are labeled from left to right, starting with **0**.
  - On EX6200 and EX8200 switches, the network ports are numbered from left to right on each line card. On line cards that have two rows of ports, the ports on the top row

start with **0** followed by the remaining even-numbered ports, and the ports on the bottom row start with **1** followed by the remaining odd-numbered ports.

- Uplink ports on an SRE module in an EX6200 switch are labeled from left to right, starting with **0**.

---

### Logical Part of an Interface Name

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16384. In the virtual part of the name, a period (.) separates the port and logical unit numbers: *type-fpc/pic/port.logical-unit-number*. For example, if you issue the **show ethernet-switching interfaces** command on a system with a default VLAN, the resulting display shows the logical interfaces associated with the VLAN:

| Interface   | State | VLAN members    | Blocking  |
|-------------|-------|-----------------|-----------|
| ge-0/0/0.0  | down  | remote-analyzer | unblocked |
| ge-0/0/1.0  | down  | default         | unblocked |
| ge-0/0/10.0 | down  | default         | unblocked |

---

### Wildcard Characters in Interface Names

In the **show interfaces** and **clear interfaces** commands, you can use wildcard characters in the *interface-name* option to specify groups of interface names without having to type each name individually. You must enclose all wildcard characters except the asterisk (\*) in quotation marks (" ").

#### Related Documentation

- [EX Series Switches Interfaces Overview on page 2577](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)

## Understanding Aggregated Ethernet Interfaces and LACP

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single link layer interface, also known as a *link aggregation group (LAG)* or *bundle*.

Aggregating multiple links between physical interfaces creates a single logical point-to-point trunk link or a LAG. The LAG balances traffic across the member links within an aggregated Ethernet bundle and effectively increases the uplink bandwidth. Another advantage of link aggregation is increased availability, because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.

Link Aggregation Control Protocol (LACP), a component of IEEE 802.3ad, provides additional functionality for LAGs.

This topic describes:

- [Link Aggregation Group \(LAG\) on page 2583](#)
- [Link Aggregation Control Protocol \(LACP\) on page 2584](#)

### Link Aggregation Group (LAG)

You configure a LAG by specifying the link number as a physical device and then associating a set of interfaces (ports) with the link. All the interfaces must have the same speed and be in full-duplex mode. Juniper Networks Junos operating system (Junos OS) for EX Series Ethernet Switches assigns a unique ID and port priority to each interface. The ID and priority are not configurable.

The number of interfaces that can be grouped into a LAG and the total number of LAGs supported on a switch varies according to switch model. [Table 260 on page 2583](#) lists the EX Series switches and the maximum number of interfaces per LAG and the maximum number of LAGs they support. MX Series devices can support up to 64 LAGs.

**Table 260: Maximum Interfaces per LAG and Maximum LAGs per Switch**

| Switch                                                             | Maximum Interfaces per LAG | Maximum LAGs |
|--------------------------------------------------------------------|----------------------------|--------------|
| EX2200                                                             | 8                          | 32           |
| EX3200                                                             | 8                          | 32           |
| EX3300 and EX3300 Virtual Chassis                                  | 8                          | 111          |
| EX4200 and EX4200 Virtual Chassis                                  | 8                          | 111          |
| EX4300 and EX4300 Virtual Chassis                                  | 16                         | 112          |
| EX4500, EX4500 Virtual Chassis, EX4550, and EX4550 Virtual Chassis | 8                          | 111          |
| EX6200                                                             | 8                          | 111          |
| EX8200                                                             | 12                         | 255          |
| EX8200 Virtual Chassis                                             | 12                         | 239          |

When configuring LAGs, consider the following guidelines:

- You must configure the LAG on both sides of the link.
- You must set the interfaces on either side of the link to the same speed.
- You can configure and apply firewall filters on a LAG.
- You can optionally configure LACP for link negotiation.
- You can optionally configure LACP for link protection.

You can combine physical Ethernet ports belonging to different member switches of a Virtual Chassis configuration to form a LAG. See [“Understanding EX Series Virtual Chassis Port Link Aggregation” on page 5083](#) and [Understanding Link Aggregation in an EX8200 Virtual Chassis](#).



**NOTE:** The interfaces that are included within a LAG are sometimes referred to as *member interfaces*. Do not confuse this term with *member switches*, which refers to switches that are interconnected as a Virtual Chassis. It is possible to create a LAG that is composed of member interfaces that are located in different member switches of a Virtual Chassis.

A LAG hashing algorithm determines how traffic entering a LAG is placed onto the bundle's member links. The LAG hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle. You can configure the fields used by the LAG hashing algorithm on some EX series switches. See [“Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\)” on page 2686](#).

A LAG creates a single logical point-to-point connection. A typical deployment for a LAG would be to aggregate trunk links between an access switch and a distribution switch or customer edge (CE) router.

### Link Aggregation Control Protocol (LACP)

---

When LACP is configured, it detects misconfigurations on the local end or the remote end of the link. Thus, LACP can help prevent communication failure:

- When LACP is not enabled, a local LAG might attempt to transmit packets to a remote single interface, which causes the communication to fail.
- When LACP is enabled, a local LAG cannot transmit packets unless a LAG with LACP is also configured on the remote end of the link.

By default, Ethernet links do not exchange LACP protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit LACP PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when the Ethernet link receives them from the remote end. The transmitting link is known as the *actor* and the receiving link is known as the *partner*.

In a scenario where a dual-homed server is deployed with a switch, the network interface cards form a LAG with the switch. During a server upgrade, the server might not be able to exchange LACP PDUs. In such a situation, you can configure an interface to be in the **up** state even if no PDUs are exchanged. Use the **force-up** statement to configure an interface when the peer has limited LACP capability. The interface selects the associated LAG by default, whether the switch and peer are both in active or passive mode. When PDUs are not received, the partner is considered to be working in the passive mode. Therefore, LACP PDU transmissions are controlled by the transmitting link.

If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In such a scenario, do not



configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

**Related Documentation**

- [Understanding EX Series Virtual Chassis Port Link Aggregation on page 5083](#)
- [Understanding Link Aggregation in an EX8200 Virtual Chassis](#)
- [Understanding Redundant Trunk Links on page 2276](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 2671](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\) on page 2672](#)
- [Junos OS Network Interfaces Configuration Guide](#)

## Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic

Juniper Networks EX Series and QFX Series use a hashing algorithm to determine how to forward traffic over a link aggregation group (LAG) bundle or to the next-hop device when equal-cost multipath (ECMP) is enabled.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. You can configure some of the fields that are used by the hashing algorithm.

This topic contains the following sections:

- [Understanding the Hashing Algorithm on page 2585](#)
- [IP \(IPv4 and IPv6\) on page 2586](#)
- [MPLS on page 2587](#)
- [MAC-in-MAC Packet Hashing on page 2588](#)
- [Layer 2 Header Hashing on page 2589](#)

### Understanding the Hashing Algorithm

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

For ECMP, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. The packet fields used by the hashing algorithm varies by the packet's EtherType and, in some instances, by the configuration on the switch. The hashing algorithm recognizes the following EtherTypes:

- IP (IPv4 and IPv6)
- MPLS
- MAC-in-MAC

Traffic that is not recognized as belonging to any of these EtherTypes is hashed based on the Layer 2 header. IP and MPLS traffic are also hashed based on the Layer 2 header when a user configures the hash mode as Layer 2 header.

You can configure some fields that are used by the hashing algorithm to make traffic forwarding decisions. You cannot, however, configure how certain values within a header are used by the hashing algorithm.

Note the following points regarding the hashing algorithm:

- The fields selected for hashing are based on the packet type only. The fields are not based on any other parameters, including forwarding decision (bridged or routed) or egress LAG bundle configuration (Layer 2 or Layer 3).
- The same fields are used for hashing unicast and multicast packets. Unicast and multicast packets are, however, hashed differently.
- The same fields are used by the hashing algorithm to hash ECMP and LAG traffic, but the hashing algorithm hashes ECMP and LAG traffic differently. The different hashing ensures that traffic is not polarized when a LAG bundle is part of the ECMP next-hop path.
- The same fields are used for hashing regardless of whether the switch is or is not participating in a mixed or non-mixed Virtual Chassis or Virtual Chassis Fabric (VCF).

The fields used for hashing by each EtherType as well as the fields used by the Layer 2 header are discussed in the following sections.

### IP (IPv4 and IPv6)

---

Payload fields in IPv4 and IPv6 packets are used by the hashing algorithm when IPv4 or IPv6 packets need to be placed onto a member link in a LAG bundle or sent to the next-hop device when ECMP is enabled.

The hash mode is set to Layer 2 payload field, by default. IPv4 and IPv6 payload fields are used for hashing when the hash mode is set to Layer 2 payload.

If the hash mode is configured to Layer 2 header, IPv4, IPv6, and MPLS packets are hashed using the Layer 2 header fields. If you want incoming IPv4, IPv6, and MPLS packets hashed by the source MAC address, destination MAC address, or EtherType fields, you must set the hash mode to Layer 2 header.

[Table 261 on page 2587](#) displays the IPv4 and IPv6 payload fields that are used by the hashing algorithm, by default.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.
- (configurable)—Field can be configured to be used or not used by the hashing algorithm.

Table 261: IPv4 and IPv6 Hashing Fields

| Fields                      | EX4300              |                     | QFX5100             |                     |
|-----------------------------|---------------------|---------------------|---------------------|---------------------|
|                             | LAG                 | ECMP                | LAG                 | ECMP                |
| Source MAC                  | X                   | X                   | X                   | X                   |
| Destination MAC             | X                   | X                   | X                   | X                   |
| EtherType                   | X                   | X                   | X                   | X                   |
| VLAN ID                     | X<br>(configurable) | X<br>(configurable) | X<br>(configurable) | X<br>(configurable) |
| Source IP or IPv6           | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) |
| Destination IP or IPv6      | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) |
| Protocol (IPv4 only)        | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) |
| Next header (IPv6 only)     | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) |
| Layer 4 Source Port         | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) |
| Layer 4 Destination Port    | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) | ✓<br>(configurable) |
| IPv6 Flow label (IPv6 only) | X                   | X                   | X                   | X                   |

## MPLS

The hashing algorithm hashes MPLS packets using the source IP, destination IP, MPLS label 0, MPLS label 1, and MPLS label 2 fields.

Table 262 on page 2588 displays the MPLS payload fields that are used by the hashing algorithm, by default:

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.

The fields used by the hashing algorithm for MPLS packet hashing are not user-configurable.

The source IP and destination IP fields are not always used for hashing. For non-terminated MPLS packets, the payload is checked if the bottom of stack (BoS) flag is seen in the packet. If the payload is IPv4 or IPv6, then the IP source address and IP destination address fields are used for hashing along with the MPLS labels. If the BoS flag is not seen in the packet, only the MPLS labels are used for hashing.

**Table 262: MPLS Hashing Fields**

| Field                          | EX4300 | QFX5100 |
|--------------------------------|--------|---------|
| Source MAC                     | X      | X       |
| Destination MAC                | X      | X       |
| EtherType                      | X      | X       |
| VLAN ID                        | X      | X       |
| Source IP                      | ✓      | ✓       |
| Destination IP                 | ✓      | ✓       |
| Protocol (for IPv4 packets)    | X      | X       |
| Next header (for IPv6 packets) | X      | X       |
| Layer 4 Source Port            | X      | X       |
| Layer 4 Destination Port       | X      | X       |
| IPv6 Flow lab                  | X      | X       |
| MPLS label 0                   | ✓      | ✓       |
| MPLS label 1                   | ✓      | ✓       |
| MPLS label 2                   | ✓      | ✓       |

### MAC-in-MAC Packet Hashing

Packets using the MAC-in-MAC EtherType are hashed by the hashing algorithm using the Layer 2 payload source MAC, Layer 2 payload destination MAC, and Layer 2 payload EtherType fields. See [Table 263 on page 2589](#).

Hashing using the fields in the MAC-in-MAC EtherType packet is first supported on EX4300 switches in Release 13.2X51-D20. Hashing using the fields in the MAC-in-MAC EtherType is not supported on earlier releases.

The fields used by the hashing algorithm for MAC-in-MAC hashing are not user-configurable.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.

**Table 263: MAC-in-MAC Hashing Fields**

| Field                           | EX4300 | QFX5100 |
|---------------------------------|--------|---------|
| Layer 2 Payload Source MAC      | ✓      | ✓       |
| Layer 2 Payload Destination MAC | ✓      | ✓       |
| Layer 2 Payload EtherType       | ✓      | ✓       |
| Layer 2 Payload Outer VLAN      | X      | X       |

#### Layer 2 Header Hashing

Layer 2 header fields are used by the hashing algorithm when a packet's EtherType is not recognized as IP (IPv4 or IPv6), MPLS, or MAC-in-MAC. The Layer 2 header fields are also used for hashing IPv4, IPv6, and MPLS traffic instead of the payload fields when the hash mode is set to Layer 2 header.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.
- (configurable)—Field can be configured to be used or not used by the hashing algorithm.

**Table 264: Layer 2 Header Hashing Fields**

| Field           | EX4300              | QFX5100             |
|-----------------|---------------------|---------------------|
| Source MAC      | ✓<br>(configurable) | ✓<br>(configurable) |
| Destination MAC | ✓<br>(configurable) | ✓<br>(configurable) |
| EtherType       | ✓<br>(configurable) | ✓<br>(configurable) |
| VLAN ID         | X<br>(configurable) | X<br>(configurable) |

#### Related Documentation

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2686](#)

## Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces

Energy Efficient Ethernet (EEE), an Institute of Electrical and Electronics Engineers (IEEE) 802.3az standard, reduces the power consumption of physical layer devices (PHYs) during periods of low link utilization. EEE saves energy by putting part of the transmission circuit into low power mode when the link is idle.

An Ethernet link consumes power even when a link is idle. EEE provides a method to utilize power in such a way that Ethernet links use power only during data transmission. EEE specifies a signaling protocol, Low Power Idle (LPI) for achieving the power saving during the idle time of Ethernet links. EEE allows PHYs to exchange LPI indications to signal the transition to low power mode when there is no traffic. LPI indicates when a link can go idle and when the link needs to resume after a predefined delay without impacting data transmission.

The following copper PHYs are standardized by IEEE 802.3az:

- 100BASE-T
- 1000BASE-T
- 10GBASE-T


### Related Documentation

- [Configuring Energy Efficient Ethernet on Interfaces \(CLI Procedure\) on page 2684](#)

## Understanding Local Link Bias


Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis or VCF using a different member link in the LAG bundle. The traffic flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is enabled is illustrated in [Figure 41 on page 2590](#).

Figure 41: Egress Traffic Flow with Local Link Bias



When local link bias is disabled, egress traffic exiting a Virtual Chassis or VCF on a LAG bundle can be forwarded out of any member link in the LAG bundle. Traffic forwarding decisions are made by an internal algorithm that attempts to load-balance traffic between the member links in the bundle. VCP bandwidth is frequently consumed by egress traffic when local link bias is disabled because the egress traffic traverses the VCPs to reach the destination egress member link in the LAG bundle. The traffic flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is disabled is illustrated in Figure 42 on page 2591.

Figure 42: Egress Traffic Flow without Local Link Bias



Local link bias is configured in a LAG bundle. A Virtual Chassis or VCF that has multiple LAG bundles can contain bundles that have and have not enabled local link bias. Local link bias only impacts the forwarding of unicast traffic exiting a Virtual Chassis or VCF; ingress traffic handling is not impacted by the local link bias setting. Egress multicast, unknown unicast, and broadcast traffic exiting a Virtual Chassis or VCF over a LAG bundle is not impacted by the local link bias setting and is always load-balanced among the member links. Local link bias is disabled, by default.

You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced across the member links in the LAG bundle as it exits the Virtual Chassis or VCF.

#### Related Documentation

- [Configuring Local Link Bias \(CLI Procedure\) on page 2686](#)

## Understanding Layer 3 Subinterfaces

A Layer 3 subinterface is a logical division of a physical interface that operates at the network level and therefore can receive and forward 802.1Q VLAN tags. You can use Layer 3 subinterfaces to route traffic among multiple VLANs along a single trunk line that connects a Juniper Networks EX Series Ethernet Switch to a Layer 2 switch. Only one physical connection is required between the switches. This topology is often called a *router on a stick* or a *one-armed router* when the Layer 3 device is a router.

To create Layer 3 subinterfaces on an EX Series switch, you enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface.

You can partition one physical interface into up to 4094 different subinterfaces, one for each VLAN. We recommend that you use the VLAN ID as the subinterface number when you configure the subinterface. Juniper Networks Junos operating system (Junos OS) reserves VLAN IDs 0 and 4095.

VLAN tagging places the VLAN ID in the frame header, allowing each physical interface to handle multiple VLANs. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. Junos OS on EX Series switches supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces. Double-tagging is not supported.

**Related  
Documentation**

- [EX Series Switches Interfaces Overview on page 2577](#)
- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- [Junos OS Ethernet Interfaces Configuration Guide](#)

## Understanding Unicast RPF

Unicast reverse-path forwarding (RPF) helps protect the switch against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled. It also helps ensure that traffic arriving on ingress interfaces comes from a network source that the receiving interface can reach.

When you enable unicast RPF, the switch forwards a packet only if the receiving interface is the best return path to the packet's unicast source address. This is known as strict mode unicast RPF.



**NOTE:** On Juniper Networks EX3200, EX4200, and EX4300 Ethernet Switches, the switch applies unicast RPF *globally* to all interfaces when unicast RPF is configured on any interface. For additional information, see [“Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches” on page 2595](#).

This topic covers:

- [Unicast RPF for Switches Overview on page 2593](#)
- [Unicast RPF Implementation on page 2593](#)
- [When to Enable Unicast RPF on page 2594](#)
- [When Not to Enable Unicast RPF on page 2595](#)
- [Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches on page 2595](#)



---

## Unicast RPF for Switches Overview

---

Unicast RPF functions as an ingress filter that reduces the forwarding of IP packets that might be spoofing an address. By default, unicast RPF is disabled on the switch interfaces.

The type of unicast RPF provided on the switches—that is, strict mode unicast RPF is especially useful on untrusted interfaces. An untrusted interface is an interface where untrusted users or processes can place packets on the network segment.

The switch supports only the active paths method of determining the best return path back to a unicast source address. The active paths method looks up the best reverse path entry in the forwarding table. It does not consider alternate routes specified using routing-protocol-specific methods when determining the best return path.

If the forwarding table lists the receiving interface as the interface to use to forward the packet back to its unicast source, it is the best return path interface.

Use strict mode unicast RPF only on symmetrically routed interfaces. (For information about symmetrically routed interfaces, see [“When to Enable Unicast RPF” on page 2594.](#))

For more information about strict unicast RPF, see RFC 3704, *Ingress Filtering for Multihomed Networks* at <http://www.ietf.org/rfc/rfc3704.txt>.

---

## Unicast RPF Implementation

---

This section includes:

- [Unicast RPF Packet Filtering on page 2593](#)
- [Bootstrap Protocol \(BOOTP\) and DHCP Requests on page 2593](#)
- [Default Route Handling on page 2594](#)

### ***Unicast RPF Packet Filtering***

When you enable unicast RPF on the switch, the switch handles traffic in the following manner:

- If the switch receives a packet on the interface that is the best return path to the unicast source address of that packet, the switch forwards the packet.
- If the best return path from the switch to the packet's unicast source address is not the receiving interface, the switch discards the packet.
- If the switch receives a packet that has a source IP address that does not have a routing entry in the forwarding table, the switch discards the packet.

### ***Bootstrap Protocol (BOOTP) and DHCP Requests***

Bootstrap protocol (BOOTP) and DHCP request packets are sent with a broadcast MAC address and therefore the switch does not perform unicast RPF checks on them. The switch forwards all BOOTP packets and DHCP request packets without performing unicast RPF checks.

### Default Route Handling

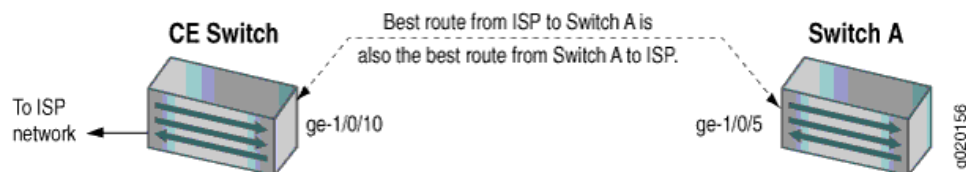
If the best return path to the source is the default route (0.0.0.0) and the default route points to **reject**, the switch discards the packets. If the default route points to a valid network interface, the switch performs a normal unicast RPF check on the packets.

### When to Enable Unicast RPF

Enable unicast RPF when you want to ensure that traffic arriving on a network interface comes from a source that resides on a network that that interface can reach. You can enable unicast RPF on untrusted interfaces to filter spoofed packets. For example, a common application for unicast RPF is to help defend an enterprise network from DoS/DDoS attacks coming from the Internet.

Enable unicast RPF only on symmetrically routed interfaces. A symmetrically routed interface uses the same route in both directions between the source and the destination, as shown in [Figure 43 on page 2594](#). Symmetrical routing means that if an interface receives a packet, the switch uses the same interface to send a reply to the packet source (the receiving interface matches the forwarding-table entry for the best return path to the source).

Figure 43: Symmetrically Routed Interfaces



Enabling unicast RPF on asymmetrically routed interfaces (where different interfaces receive a packet and reply to its source) results in packets from legitimate sources being filtered (discarded) because the best return path is not the same interface that received the packet.

The following switch interfaces are most likely to be symmetrically routed and thus are candidates for unicast RPF enabling:

- The service provider edge to a customer
- The customer edge to a service provider
- A single access point out of the network (usually on the network perimeter)
- A terminal network that has only one link



**NOTE:** Because unicast RPF is enabled globally on EX3200, EX4200, and EX4300 switches, ensure that *all* interfaces are symmetrically routed before you enable unicast RPF on these switches. Enabling unicast RPF on asymmetrically routed interfaces results in packets from legitimate sources being filtered.



**TIP:** Enabling unicast RPF as close as possible to the traffic source stops spoofed traffic before it can proliferate or reach interfaces that do not have unicast RPF enabled.


### When Not to Enable Unicast RPF

Typically, you will not enable unicast RPF if:

- Switch interfaces are multihomed.
- Switch interfaces are trusted interfaces.
- BGP is carrying prefixes and some of those prefixes are not advertised or are not accepted by the ISP under its policy. (The effect in this case is the same as filtering an interface by using an incomplete access list.)
- Switch interfaces face the network core. Core-facing interfaces are usually asymmetrically routed.

An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, as shown in [Figure 44 on page 2595](#). This means that if an interface receives a packet, that interface does not match the forwarding table entry as the best return path back to the source. If the receiving interface is not the best return path to the source of a packet, unicast RPF causes the switch to discard the packet even though it comes from a valid source.

**Figure 44: Asymmetrically Routed Interfaces**



**NOTE:** Do not enable unicast RPF on EX3200, EX4200, and EX4300 switches if any switch interfaces are asymmetrically routed, because unicast RPF is enabled globally on all interfaces of these switches. All switch interfaces must be symmetrically routed for you to enable unicast RPF without the risk of the switch discarding traffic that you want to forward.

### Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches

On EX3200, EX4200, and EX4300 switches, the switch implements unicast RPF on a global basis. You cannot enable unicast RPF on a per-interface basis. Unicast RPF is globally disabled by default.

- When you enable unicast RPF on any interface, it is automatically enabled on all switch interfaces, including link aggregation groups (LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs).
- When you disable unicast RPF on the interface (or interfaces) on which you enabled unicast RPF, it is automatically disabled on all switch interfaces.



**NOTE:** You must explicitly disable unicast RPF on every interface on which it was explicitly enabled or unicast RPF remains enabled on all switch interfaces.

QFX switches and EX3200 and EX4200 switches do not perform unicast RPF filtering on equal-cost multipath (ECMP) traffic. The unicast RPF check examines only one best return path to the packet source, but ECMP traffic employs an address block consisting of multiple paths. Using unicast RPF to filter ECMP traffic on these switches can result in the switch discarding packets that you want to forward because the unicast RPF filter does not examine the entire ECMP address block.

**Related  
Documentation**

- [Example: Configuring Unicast RPF on an EX Series Switch](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 2689](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 2691](#)

## Understanding IP Directed Broadcast for EX Series Switches

IP directed broadcast helps you implement remote administration tasks such as backups and wake-on-LAN (WOL) application tasks by sending broadcast packets targeted at the hosts in a specified destination subnet. IP directed broadcast packets traverse the network in the same way as unicast IP packets until they reach the destination subnet. When they reach the destination subnet and IP directed broadcast is enabled on the receiving switch, the switch translates (*explodes*) the IP directed broadcast packet into a broadcast that floods the packet on the target subnet. All hosts on the target subnet receive the IP directed broadcast packet.

This topic covers:

- [IP Directed Broadcast for EX Series Switches Overview on page 2596](#)
- [IP Directed Broadcast Implementation for EX Series Switches on page 2597](#)
- [When to Enable IP Directed Broadcast on page 2597](#)
- [When Not to Enable IP Directed Broadcast on page 2597](#)

### IP Directed Broadcast for EX Series Switches Overview

---

IP directed broadcast packets have a destination IP address that is a valid broadcast address for the subnet that is the target of the directed broadcast (the target subnet). The intent of an IP directed broadcast is to flood the target subnet with the broadcast packets without broadcasting to the entire network. IP directed broadcast packets cannot originate from the target subnet.

When you send an IP directed broadcast packet, as it travels to the target subnet, the network forwards it in the same way as it forwards a unicast packet. When the packet reaches a switch that is directly connected to the target subnet, the switch checks to see whether IP directed broadcast is enabled on the interface that is directly connected to the target subnet:

- If IP directed broadcast is enabled on that interface, the switch broadcasts the packet on that subnet by rewriting the destination IP address as the configured broadcast IP address for the subnet. The switch converts the packet to a link-layer broadcast packet that every host on the network processes.
- If IP directed broadcast is disabled on the interface that is directly connected to the target subnet, the switch drops the packet.

### IP Directed Broadcast Implementation for EX Series Switches

---

You configure IP directed broadcast on a per-subnet basis by enabling IP directed broadcast on the Layer 3 interface of the subnet's VLAN. When the switch that is connected to that subnet receives a packet that has the subnet's broadcast IP address as the destination address, the switch broadcasts the packet to all hosts on the subnet.

By default, IP directed broadcast is disabled.

### When to Enable IP Directed Broadcast

---

IP directed broadcast is disabled by default. Enable IP directed broadcast when you want to perform remote management or administration services such as backups or WOL tasks on hosts in a subnet that does not have a direct connection to the Internet.

Enabling IP directed broadcast on a subnet affects only the hosts within that subnet. Only packets received on the subnet's Layer 3 interface that have the subnet's broadcast IP address as the destination address are flooded on the subnet.

### When Not to Enable IP Directed Broadcast

---

Typically, you do not enable IP directed broadcast on subnets that have direct connections to the Internet. Disabling IP directed broadcast on a subnet's Layer 3 interface affects only that subnet. If you disable IP directed broadcast on a subnet and a packet that has the broadcast IP address of that subnet arrives at the switch, the switch drops the broadcast packet.

If a subnet has a direct connection to the Internet, enabling IP directed broadcast on it increases the network's susceptibility to denial-of-service (DoS) attacks.

For example, a malicious attacker can spoof a source IP address (use a source IP address that is not the actual source of the transmission to deceive a network into identifying the attacker as a legitimate source) and send IP directed broadcasts containing Internet Control Message Protocol (ICMP) echo (ping) packets. When the hosts on the network with IP directed broadcast enabled receive the ICMP echo packets, they all send replies to the victim that has the spoofed source IP address. This creates a flood of ping replies in a DoS attack that can overwhelm the spoofed source address; this is known as a *smurf* attack. Another common DoS attack on exposed networks with IP directed broadcast

enabled is a *fraggle* attack, which is similar to a smurf attack except that the malicious packet is a User Datagram Protocol (UDP) echo packet instead of an ICMP echo packet.

**Related  
Documentation**

- *Example: Configuring IP Directed Broadcast on an EX Series Switch*
- *Configuring IP Directed Broadcast (CLI Procedure)*
- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 2691](#)

## Understanding Interface Ranges on EX Series Switches



**NOTE:** This concept uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Understanding Interface Ranges on EX Series Switches*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

You can use the interface ranges to group interfaces of the same type that share a common configuration profile. This helps reduce the time and effort in configuring interfaces on Juniper Networks EX Series Ethernet Switches. The configurations common to all the interfaces can be included in the interface range definition.

The interface range definition contains the name of the interface range defined, the names of the individual member interfaces that do not fall in a series of interfaces, a range of interfaces defined in the member range, and the configuration statements common to all the interfaces. An interface range defined with member ranges and individual members but without any common configurations, is also a valid definition.



**NOTE:** The interface range definition is supported only for Gigabit, 10-Gigabit, 40-Gigabit, and Fast Ethernet interfaces.

The common configurations defined in the interface range will be overridden by the local configuration.

The defined interface ranges can be used at places where the **interface** node is used in the following configuration hierarchies:

- **forwarding-options analyzer *name* input egress interface**
- **forwarding-options analyzer *name* input ingress interface**
- **poe interface**
- **protocols dot1x authenticator interface**
- **protocols igmp interface**
- **protocols isis interface**
- **protocols layer2-control bpdu-block interface**

- `protocols link-management peer name lmp-control-channel`
- `protocols link-management te-link name interface`
- `protocols lldp interface`
- `protocols lldp-med interface`
- `protocols mstp interface`
- `protocols oam ethernet link-fault-management interface`
- `protocols ospf area area-id interface`
- `protocols pim interface`
- `protocols router-advertisement interface`
- `protocols router-discovery interface`
- `protocols rsvp interface`
- `protocols sflow interfaces`
- `protocols vstp vlan vlan-id interface`
- `switch-options redundant-trunk-group group-name interface`
- `switch-options voip interface`

#### Related Documentation

- [Interface Ranges on page 2651](#)
- [EX Series Switches Interfaces Overview on page 2577](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 2689](#)
- [interface-range on page 2789](#)

## Understanding Multichassis Link Aggregation

Layer 2 networks are increasing in scale mainly because of technologies such as virtualization. Protocol and control mechanisms that limit the disastrous effects of a topology loop in the network have become necessary. Spanning Tree Protocol (STP) is the primary solution to this problem because it provides a loop-free Layer 2 environment. STP has gone through a number of enhancements and extensions, and although it scales to very large network environments, it still provides only one active path from one device to another, regardless of the number of actual connections that might exist in the network. Although STP is a robust and scalable solution to redundancy in a Layer 2 network, the single logical link creates two problems: At least half of the available system bandwidth is off-limits to data traffic, and network topology changes occur. Rapid Spanning Tree Protocol (RSTP) reduces the overhead of the rediscovery process and allows a Layer 2 network to reconverge faster, but the delay is still high.

Link aggregation (IEEE 802.3ad) solves some of these problems by enabling users to use more than one link connection between switches. All physical connections are considered one logical connection. The problem with standard link aggregation is that the connections are point to point.

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two MC-LAG peers (for example, EX9200 switches). An MC-LAG provides redundancy and load balancing between the two MC-LAG peers, multihoming support, and a loop-free Layer 2 network without running STP.

On one end of an MC-LAG, there is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to have an MC-LAG configured. On the other side of the MC-LAG, there are two MC-LAG peers. Each of the MC-LAG peers has one or more physical links connected to a single client device.

This topic describes the following:

- [MC-LAG Configuration Guidelines and Functional Behavior on page 2600](#)
- [Active-Active and Active-Standby Mode on page 2602](#)
- [ICCP and ICL on page 2602](#)
- [LACP on page 2603](#)
- [Data Traffic Forwarding Rules on page 2603](#)
- [Multichassis Link Protection on page 2604](#)
- [Failure Handling During a Split-Brain State on page 2604](#)
- [Load-Balancing on page 2606](#)
- [Layer 2 Unicast Features Supported on page 2606](#)
- [Layer 2 Multicast Features Supported on page 2607](#)
- [IGMP Snooping on an Active-Active MC-LAG on page 2607](#)
- [Layer 3 Unicast Features Supported on page 2608](#)
- [VRRP Active-Standby Support on page 2608](#)
- [MAC Address Management on page 2609](#)
- [MAC Address Synchronization on page 2609](#)
- [Address Resolution Protocol Active-Active MC-LAG Support Methodology on page 2610](#)
- [DHCP Relay with Option 82 on page 2611](#)
- [Layer 3 Multicast on page 2612](#)
- [MC-LAG Upgrade Guidelines on page 2613](#)

---

### MC-LAG Configuration Guidelines and Functional Behavior

---

When you configure MC-LAGs, we recommend that you follow certain guidelines to ensure that you obtain optimum benefit from the MC-LAG feature.

[Table 265 on page 2601](#) provides best practice configuration guidelines for MC-LAGs and [Table 266 on page 2602](#) describes important functional behavior for MC-LAGs.



Table 265: MC-LAG Configuration Guidelines

- We recommend that you use separate ports and choose different FPCs for the interchassis link (ICL) and Inter-Chassis Control Protocol (ICCP) interfaces.
- We recommend that you configure the backup liveness detection feature to implement faster failover of data traffic during an MC-LAG peer reboot. Configure the **backup-liveness-detection** statement on the management interface (fxp0) only.

**NOTE:** On EX9200 switches, the **backup-liveness-detection** statement was added in Junos OS Release 13.2R1.

- The following two methods can be used to enable Layer 3 functionality across an MC-LAG. We recommend that you use the Virtual Router Redundancy Protocol (VRRP) over integrated routing and bridging (IRB) interfaces method. Use media access control (MAC) address synchronization only when you cannot configure VRRP over IRB.
  1. Configure different IP addresses on IRB interfaces and run VRRP over the IRB interfaces. The virtual IP address is the gateway IP address for the MC-LAG clients.
  2. Configure the MAC address synchronization feature using the **set vlans *vlan-name* mcae-mac-synchronize** command and configure the same IP address on each of the IRBs on the MC-LAG peers. This IP address will be the gateway IP address for the MC-LAG clients.
- If you are using the VRRP over IRB method to enable Layer 3 functionality, you must configure static ARP entries for the IRB interface of the remote MC-LAG peer to allow routing protocols to run over the IRB interfaces.
- You must configure the ICL interface as a router-facing interface (by configuring the **multicast-router-interface** statement) for multicast forwarding to work in an MC-LAG environment.
- You must configure the **multichassis-lag-replicate-state** statement for Internet Group Management Protocol (IGMP) snooping to work properly in an MC-LAG environment.
- You must enable Protocol Independent Multicast (PIM) on the IRB interface to avoid multicast duplication.
- If you are using Layer 3 multicast, configure the IP address on the active MC-LAG peer with a high IP address or a high designated router priority.

**NOTE:** Use this configuration guideline only if you can ensure that the ICCP will not go down unless the router or switch is down.

- You can configure the **prefer-status-control-active** statement with the **mc-ae status-control standby** configuration to prevent the LACP MC-LAG system ID from reverting to the default Link Aggregation Control Protocol (LACP) system ID on ICCP failure. You must also configure the **hold-time down** value (at the **[edit interfaces *interface-name*]** hierarchy level) for the ICL with the **mc-ae status-control standby** configuration to be higher than the ICCP Bidirectional Forwarding Detection (BFD) timeout. This configuration prevents data traffic loss by ensuring that when the router or switch with the **mc-ae status-control active** configuration goes down, the router or switch with the **mc-ae status-control standby** configuration does not go into standby mode.

To make the **prefer-status-control-active** configuration work with the **mc-ae status-control standby** configuration when an ICL logical interface is configured on an aggregated Ethernet interface, you must either configure the **lACP periodic interval** statement at the **[edit interfaces *interface-name* aggregated-ether-options]** hierarchy level as slow or configure the **detection-time threshold** statement at the **[edit protocols iccp peer *liveness-detection*]** hierarchy level as less than 3 seconds.

**NOTE:** On EX9200 switches, the **prefer-status-control-active** statement was added in Junos OS Release 13.2R1.

We recommend that you configure the ICCP liveness-detection interval to be at least 8 seconds, to allow GRES to work seamlessly. By default, ICCP liveness detection uses multihop BFD, which runs in centralized mode. If you have configured ICCP connectivity through a dedicated physical interface rather than through an IRB interface, then you can configure single-hop BFD, and the restriction on the liveness-detection interval does not apply.

Table 265: MC-LAG Configuration Guidelines (*continued*)

We recommend that you configure only one redundancy group between MC-LAG nodes. The redundancy group represents the domain of high availability between the MC-LAG nodes. One redundancy group is sufficient between a pair of MC-LAG nodes.

Table 266: MC-LAG Functional Behavior

- STP is not supported on the ICL or MC-LAG interfaces.
- Load-balancing of network traffic between MC-LAG peers is 100 percent local bias.
- Load-balancing of network traffic between multiple LAG members in a local MC-LAG node is achieved through a standard LAG hashing algorithm.
- MAC learning is disabled on the ICL. Consequently, source MAC addresses cannot be learned locally on the ICL. However, MAC addresses from a remote MC-LAG node can be installed on the ICL interface. For example, the MAC address for a single-homed client on a remote MC-LAG node can be installed on the ICL interface of the local MC-LAG node.
- Dynamic Address Resolution Protocol (ARP) resolution over the ICL interface is not supported. Consequently, incoming ARP replies on the ICL are discarded. However, ARP entries can be populated on the ICL interface through ICCP exchanges from a remote MC-LAG peer.
- For EX9200 switches, ARP entries that were learned remotely will be purged and then learned again during graceful Routing Engine switchover (GRES).
- Usually, a VRRP backup node does not forward incoming packets. However, when VRRP over IRB is configured in an MC-LAG active-active environment, both the VRRP master and the VRRP backup forward Layer 3 traffic arriving on the MC-AE interface.
- If you are using the MAC address synchronization method (by configuring the **set vlans *vlan-name* mcae-mac-synchronize** command) to enable Layer 3 functionality, running routing protocols over the IRB interface is not supported and gratuitous ARP requests are not sent when the MAC address on the IRB interface changes.
- Access port security features (for example, DHCP snooping, dynamic ARP inspection (DAI), and IP source guard) are not supported on the ICL or MC-LAG interfaces.

### Active-Active and Active-Standby Mode

In active-active mode, all member links are active on the MC-LAG. In this mode, media access control (MAC) addresses learned on one MC-LAG peer are propagated to the other MC-LAG peer.

In active-standby mode, only one of the MC-LAG peers is active at any given time. The other MC-LAG peer is in backup (standby) mode. The active MC-LAG peer uses Link Aggregation Control Protocol (LACP) to advertise to client devices that its child link is available for forwarding data traffic.

### ICCP and ICL

The MC-LAG peers use Inter-Chassis Control Protocol (ICCP) to exchange control information and coordinate with each other to ensure that data traffic is forwarded properly. ICCP replicates control traffic and forwarding states across the MC-LAG peers and communicates the operational state of the MC-LAG members. Because ICCP uses TCP/IP to communicate between the peers, the two peers must be connected to each

other. ICCP messages exchange MC-LAG configuration parameters and ensure that both peers use the correct LACP parameters.

The interchassis link (ICL), also known as the interchassis link-protection link (ICL-PL), is used to forward data traffic across the MC-LAG peers. This link provides redundancy when a link failure (for example, an MC-LAG trunk failure) occurs on one of the active links. The ICL can be a single physical Ethernet interface or an aggregated Ethernet interface. You can configure only one ICL between the two MC-LAG peers, although you can configure multiple MC-LAGs between them.

## LACP

LACP is a subcomponent of the IEEE 802.3ad standard. LACP is used to discover multiple links from a client device connected to an MC-LAG peer. LACP must be configured on all member links for an MC-LAG to work correctly.

## Data Traffic Forwarding Rules

In active-active MC-LAG topologies, network interfaces are categorized into three interface types, as follows:

- S-Links—Single-homed link (S-Link) terminating on an MC-LAG peer device.
- MC-Links—MC-LAG links.
- ICL—Interchassis data link.

Depending on the incoming and outgoing interface types, some constraints are added to the Layer 2 forwarding rules for MC-LAG configurations. The following data traffic forwarding rules apply:



**NOTE:** If only one MC-LAG member link is in the up state, it is considered an S-Link.

- When an MC-LAG network receives a packet from a local MC-Link or S-Link, the packet is forwarded to other local interfaces, including S-Links and MC-Links based on the normal Layer 2 forwarding rules and on the configuration of the **mesh-group** and **no-local-switching** statements. If MC-Links and S-Links are in the same mesh group and their **no-local-switching** statements are enabled, the received packets are only forwarded upstream and not sent to MC-Links and S-Links.
- The following circumstances determine whether or not an ICL receives a packet from a local MC-Link or S-Link:
  - If the peer MC-LAG network device has S-Links or MC-LAGs that do not reside on the local MC-LAG network device.
  - Whether or not interfaces on two peering MC-LAG network devices are allowed to talk to each other.

- When an MC-LAG network receives a packet from the ICL, the packet is forwarded to all local S-Links and active MC-LAGs that do not exist in the MC-LAG network from which the packet was sent.

### Multichassis Link Protection

Multichassis link protection provides link protection between the two MC-LAG peers that host an MC-LAG. If the ICCP connection is up and the ICL comes up, the peer configured as standby brings up the multichassis aggregated Ethernet (MC-AE) interfaces shared with the peer. Multichassis protection must be configured on each MC-LAG peer that is hosting an MC-LAG.

### Failure Handling During a Split-Brain State

Configuring ICCP adjacency over aggregated links with child links on multiple FPCs mitigates the possibility of a split-brain state. A split-brain occurs when ICCP adjacency is lost between the MC-LAG peers. To work around this problem, enable backup liveness detection. With backup liveness detection enabled, the MC-LAG peers establish an out-of-band channel through the management network in addition to the ICCP channel.

During a split-brain state, both active and standby peers change LACP system IDs. Because both MC-LAG peers change the LACP system ID, the CE device accepts the LACP system ID of the first link that comes up and brings down other links carrying different LACP system IDs. When the ICCP connection is active, both of the MC-LAG peers use the configured LACP system ID. If the LACP system ID is changed during failures, the server that is connected over the MC-LAG removes these links from the aggregated Ethernet bundle.

When the ICL is operationally down and the ICCP connection is active, the LACP state of the links with status control configured as standby is set to the standby state. When the LACP state of the links is changed to standby, the server that is connected over the MC-LAG makes these links inactive and does not use them for sending data.

Table 267 on page 2604 describes the different ICCP failure scenarios. The dash means that the item is not applicable.

**Table 267: MC-LAG Failure Scenarios**

| ICCP Connection Status | ICL Status | Backup Liveness Peer Status | Action on Multichassis Aggregated Ethernet (MC-AE) Interface with Status Set to Standby |
|------------------------|------------|-----------------------------|-----------------------------------------------------------------------------------------|
| Down                   | Down or Up | Not configured              | LACP system ID is changed to default value.                                             |
| Down                   | Down or Up | Active                      | LACP system ID is changed to default value.                                             |
| Down                   | Down or Up | Inactive                    | No change in LACP system ID.                                                            |
| Up                     | Down       | —                           | LACP state is set to standby. MUX state moves to waiting state.                         |

Recovery from the split-brain state occurs automatically when the ICCP adjacency comes up between the MC-LAG peers.

If only one physical link is available for ICCP, then ICCP might go down due to link failure or FPC failure, while the node is still up—a split-brain state results. If you do not set a special configuration to avoid this situation, the mc-ae interfaces change the LACP system ID to their local defaults, thus ensuring that only one link (the first) comes up from the downstream device. A convergence delay results from the mc-ae LACP state changes on both active and standby nodes.

To avoid this problem of the split-brain state and resultant convergence delays, configure *one* of the following two options:

- Enable backup liveness detection on the management (fxp0) interface:

For example:

```
[edit ]
user@switch# set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip
10.207.64.233
```

When you configure **backup-liveness-detection**, an out-of-band channel is established between the nodes, through the management network, to test the liveness of the Routing Engine. When both ICCP and backup liveness detection fail, the remote node is considered down, so the LACP system ID is not changed on the local node.

- Configure **prefer-status-control-active** on the active node:

For example:

```
[edit ]
user@switch# set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1 events
iccp-peer-down prefer-status-control-active
```

When you configure **prefer-status-control-active**, if ICCP goes down and backup liveness detection is up, the LACP system ID is not changed. Thus if ICCP alone fails, the LACP system ID is not changed on the active node but it is changed on the standby node.



**NOTE:** Configure the **master-only** statement on the IP address of the fxp0 interface for backup liveness detection, on both the master and backup Routing Engines, to ensure that the connection is not reset during GRES in the remote peer

For example, on the master Routing Engine:

```
user@switch-re1 > show configuration interfaces fxp0 | display inheritance no-comments
unit 0 {
  family inet {
    address 10.8.2.31/24;
    address 10.8.2.33/24 {
      master-only;
    }
  }
}
```

```
}
```

For example, on the backup Routing Engine:

```
user@switch1-re1 > show configuration interfaces fxp0 | display inheritance no-comments
unit 0 {
  family inet {
    address 10.8.2.32/24;
    address 10.8.2.33/24 {
      master-only;
    }
  }
}
```

The master Routing Engine services both 10.8.2.31 and 10.8.2.33. Configure 10.8.2.33 in a backup-liveness-detection configuration on the peer node.

For example, on the backup Routing Engine:

```
user@switch2 > show configuration protocols iccp
local-ip-addr 2.2.2.2;
peer 1.1.1.1 {
  session-establishment-hold-time 50;
  redundancy-group-id-list 1;
  backup-liveness-detection {
    backup-peer-ip 10.8.2.33;
  }
  liveness-detection {
    minimum-interval 500;
    multiplier 3;
    single-hop;
  }
}
```

---

## Load-Balancing

Load-balancing of network traffic between MC-LAG peers is 100 percent local bias. Load-balancing of network traffic between multiple LAG members in a local MC-LAG node is achieved through a standard LAG hashing algorithm.

---

## Layer 2 Unicast Features Supported

The following Layer 2 unicast features are supported:

- Layer 2 unicast: learning and aging
  - Learned MAC addresses are propagated across MC-LAG peers for all of the VLANs that are spawned across the peers.
  - Aging of MAC addresses occurs when the MAC address is not seen on both of the peers.

- MAC addresses learned on single-homed links are propagated across all of the VLANs that have MC-LAG links as members.



**NOTE:** MAC learning is disabled on the ICL. Consequently, source MAC addresses cannot be learned locally on the ICL. However, MAC addresses from a remote MC-LAG node can be installed on the ICL interface. For example, the MAC address for a single-homed client on a remote MC-LAG node can be installed on the ICL interface of the local MC-LAG node.

### Layer 2 Multicast Features Supported

The following Layer 2 multicast features are supported:

- Layer 2 multicast: unknown unicast and IGMP snooping
  - Flooding happens on all links across peers if both peers have virtual LAN membership. Only one of the peers forwards traffic on a given MC-LAG link.
  - Known and unknown multicast packets are forwarded across the peers by adding the ICL port as a multicast router port.
  - IGMP membership learned on MC-LAG links is propagated across peers.
  - During an MC-LAG peer reboot, known multicast traffic is flooded until the IGMP snooping state is synced with the peer.

### IGMP Snooping on an Active-Active MC-LAG

IGMP snooping controls multicast traffic in a switched network. When IGMP snooping is not enabled, the Layer 2 device broadcasts multicast traffic out of all of its ports, even if the hosts on the network do not want the multicast traffic. With IGMP snooping enabled, a Layer 2 device monitors the IGMP join and leave messages sent from each connected host to a multicast router. This enables the Layer 2 device to keep track of the multicast groups and associated member ports. The Layer 2 device uses this information to make intelligent decisions and to forward multicast traffic to only the intended destination hosts. IGMP uses Protocol Independent Multicast (PIM) to route the multicast traffic. PIM uses distribution trees to determine which traffic is forwarded.

In an active-active MC-LAG configuration, IGMP snooping replicates the Layer 2 multicast routes so that each MC-LAG peer has the same routes. If a device is connected to an MC-LAG peer by way of a single-homed interface, IGMP snooping replicates join message to its IGMP snooping peer. If a multicast source is connected to an MC-LAG by way of a Layer 3 device, the Layer 3 device passes this information to the IRB that is configured on the MC-LAG. The first hop designated router is responsible for sending the register and register-stop messages for the multicast group. The last hop designated router is responsible for sending PIM join and leave messages toward the rendezvous point and source for the multicast group. The routing device with the smallest preference metric forwards traffic on transit LANs.

You must configure the ICL interface as a router-facing interface (by configuring the **multicast-router-interface** statement) for multicast forwarding to work in an MC-LAG environment. For the scenario in which traffic arrives by way of a Layer 3 interface, PIM and IGMP must be enabled on the IRB interface configured on the MC-LAG peers. You must enable PIM on the IRB interface to avoid multicast duplication.

### Layer 3 Unicast Features Supported

---

To provide Layer 3 routing functions to downstream clients, the MC-LAG network peers must be configured to provide the same gateway address to the downstream clients. To the upstream routers, the MC-LAG network peers could be viewed as either equal-cost multi path (ECMP) or two routes with different preference values.

The following two methods can be used to enable Layer 3 functionality across an MC-LAG. We recommend that you use the VRRP over IRB method. Use MAC address synchronization only when you cannot configure VRRP over IRB.

1. Configure different IP addresses on IRB interfaces and run Virtual Router Redundancy Protocol (VRRP) over the IRB interfaces. The virtual IP address is the gateway IP address for the MC-LAG clients.
2. Configure the MAC address synchronization feature using the **mcae-mac-synchronize** statement and configure the same IP address on each of the IRBs on the MC-LAG peers. This IP address will be the gateway IP address for the MC-LAG clients.

Layer 3 unicast feature support includes the following:

- Address Resolution Protocol (ARP) synchronization enables ARP resolution on both of the MC-LAG peers.
- DHCP Relay with option 82 enables option 82 on the MC-LAG peers. Option 82 provides information about the network location of DHCP clients. The DHCP server uses this information to implement IP addresses or other parameters for the client.

### VRRP Active-Standby Support

---

Juniper Networks Junos operating system (Junos OS) supports active-active MC-LAGs by using VRRP in active-standby mode. VRRP in active-standby mode enables Layer 3 routing over the multichassis aggregated Ethernet (MC-AE) interfaces on the MC-LAG peers. In this mode, the MC-LAG peers act as virtual routers. The peers share the virtual IP address that corresponds to the default route configured on the host or server connected to the MC-LAG. This virtual IP address (of the IRB interface) maps to either of the VRRP MAC addresses or to the logical interfaces of the MC-LAG peers. The host or server uses the VRRP MAC address to send any Layer 3 upstream packets. At any time, one of the VRRP devices is the master (active), and the other is a backup (standby). Usually, a VRRP backup node does not forward incoming packets. However, when VRRP over IRB is configured in an MC-LAG active-active environment, both the VRRP master and the VRRP backup forward Layer 3 traffic arriving on the MC-AE interface. If the master fails, all the traffic shifts to the MC-AE link on the backup.





**NOTE:** You must configure VRRP on both MC-LAG peers for both the active and standby members to accept and route packets. Additionally, you must configure the VRRP backup device to send and receive ARP requests.

Routing protocols run on the primary IP address of the IRB interface, and both of the MC-LAG peers run routing protocols independently. The routing protocols use the primary IP address of the IRB interface and the IRB MAC address to communicate with the MC-LAG peers. The IRB MAC address of each MC-LAG peer is replicated on the other MC-LAG peer and is installed as a MAC address that has been learned on the ICL.



**NOTE:** If you are using the VRRP over IRB method to enable Layer 3 functionality, you must configure static ARP entries for the IRB interface of the remote MC-LAG peer to allow routing protocols to run over the IRB interfaces.

### MAC Address Management

If an MC-LAG is configured to be active-active, upstream and downstream traffic could go through different MC-LAG peer devices. Because the MAC address is learned only on one of the MC-LAG peers, traffic in the reverse direction could be going through the other MC-LAG peer and flooding the network unnecessarily. Also, a single-homed client's MAC address is learned only on the MC-LAG peer that it is attached to. If a client attached to the peer MC-LAG network device needs to communicate with that single-homed client, then traffic would be flooded on the peer MC-LAG network device. To avoid unnecessary flooding, whenever a MAC address is learned on one of the MC-LAG peers, the address is replicated to the other MC-LAG peer. The following conditions are applied when MAC address replication is performed:

- MAC addresses learned on an MC-LAG of one MC-LAG peer must be replicated as learned on the same MC-LAG of the other MC-LAG peer.
- MAC addresses learned on single-homed customer edge (CE) clients of one MC-LAG peer must be replicated as learned on the ICL interface of the other MC-LAG peer.
- MAC address learning on an ICL is disabled from the data path. It depends on software to install MAC addresses replicated through ICCP.

### MAC Aging

MAC aging support in Junos OS extends aggregated Ethernet logic for a specified MC-LAG. A MAC address in software is not deleted until all Packet Forwarding Engines have deleted the MAC address.

### MAC Address Synchronization

MAC address synchronization enables MC-LAG peers to forward Layer 3 packets arriving on MC-AE interfaces with either its own IRB MAC address or its peer's IRB MAC address. Each MC-LAG peer installs its own IRB MAC address as well as the peer's IRB MAC address in the hardware. Each MC-LAG peer treats the packet as if it were its own packet. If MAC

address synchronization is not enabled, the IRB MAC address is installed on the MC-LAG peer as if it was learned on the ICL.



**NOTE:** If you are using the MAC address synchronization to enable Layer 3 functionality, running routing protocols over the IRB interface is not supported. If you need routing capability, configure both VRRP and routing protocols on each MC-LAG peer.

Control packets destined for a particular MC-LAG peer that arrive on an MC-AE interface of its MC-LAG peer are not forwarded on the ICL interface. Additionally, using the gateway IP address as a source address when you issue either a ping, traceroute, telnet, or FTP request is not supported.



**NOTE:** Gratuitous ARP requests are not sent when the MAC address on the IRB interface changes.

To enable the MAC address synchronization feature, issue the **set vlan *vlan-name* mcae-mac-synchronize** command on each MC-LAG peer. Configure the same IP address on both MC-LAG peers. This IP address is used as the default gateway for the MC-LAG servers or hosts.

### Address Resolution Protocol Active-Active MC-LAG Support Methodology

ARP maps IP addresses to MAC addresses. Junos OS uses ARP response packet snooping to support active-active MC-LAGs, providing easy synchronization without the need to maintain any specific state. Without synchronization, if one MC-LAG peer sends an ARP request, and the other MC-LAG peer receives the response, ARP resolution is not successful. With synchronization, the MC-LAG peers synchronize the ARP resolutions by sniffing the packet at the MC-LAG peer receiving the ARP response and replicating this to the other MC-LAG peer. This ensures that the entries in ARP tables on the MC-LAG peers are consistent.

When one of the MC-LAG peers restarts, the ARP destinations on its MC-LAG peer are synchronized. Because the ARP destinations are already resolved, its MC-LAG peer can forward Layer 3 packets out of the MC-AE interface.



**NOTE:** In some cases, ARP messages received by one MC-LAG peer are replicated to the other MC-LAG peer through ICCP. This optimization feature is applicable only for ARP replies, not ARP requests, received by the MC-LAG peers.



**NOTE:** Dynamic ARP resolution over the ICL interface is not supported. Consequently, incoming ARP replies on the ICL are discarded. However, ARP entries can be populated on the ICL interface through ICCP exchanges from a remote MC-LAG peer.



**NOTE:** During Graceful Routing Engine Switchover (GRES), ARP entries that were learned remotely will be purged and then learned again.

### DHCP Relay with Option 82



**NOTE:** DHCP relay is not supported with MAC address synchronization. If DHCP relay is required, configure VRRP over IRB for Layer 3 functionality.

DHCP relay with option 82 provides information about the network location of DHCP clients. The DHCP server uses this information to implement IP addresses or other parameters for the client. With DHCP relay enabled, DHCP request packets might take the path to the DHCP server through either of the MC-LAG peers. Because the MC-LAG peers have different host names, chassis MAC addresses, and interface names, you need to observe these requirements when you configure DHCP relay with option 82:

- Use the interface description instead of the interface name.
- Do not use the hostname as part of the circuit ID or remote ID strings.
- Do not use the chassis MAC address as part of the remote ID string.
- Do not enable the vendor ID.
- If the ICL interface receives DHCP request packets, the packets are dropped to avoid duplicate packets in the network.

A counter called *Due to received on ICL interface* has been added to the **show helper statistics** command, which tracks the packets that the ICL interface drops.

An example of the CLI output follows:

```
user@switch> show helper statistics
BOOTP:
  Received packets: 6
  Forwarded packets: 0
  Dropped packets: 6
    Due to no interface in DHCP Relay database: 0
    Due to no matching routing instance: 0
    Due to an error during packet read: 0
    Due to an error during packet send: 0
    Due to invalid server address: 0
    Due to no valid local address: 0
    Due to no route to server/client: 0
    Due to received on ICL interface: 6
```

The output shows that six packets received on the ICL interface have been dropped.

## Layer 3 Multicast

---

- [PIM Operation with Normal Mode Designated Router Election on page 2612](#)
- [PIM Operation with Dual DR Mode on page 2612](#)

Protocol Independent Multicast (PIM) and Internet Group Management Protocol (IGMP) provide support for Layer 3 multicast. In addition to the standard mode of PIM operation, there is a special mode called PIM dual DR (designated router). PIM dual DR minimizes multicast traffic loss in case of failures.

### ***PIM Operation with Normal Mode Designated Router Election***

In normal mode designated router election, the IRB interfaces on both of the MC-LAG peers are configured with PIM enabled. In this mode, one of the MC-LAG peers becomes the designated router through the PIM designated router election mechanism. The elected designated router maintains the rendezvous-point tree (RPT) and shortest-path tree (SPT) so it can receive data from the source device. The elected designated router participates in periodic PIM join and prune activities toward the rendezvous point (RP) or the source.

The trigger for initiating these join and prune activities is the IGMP membership reports that are received from interested receivers. IGMP reports received over MC-AE interfaces (potentially hashing on either of the MC-LAG peers) and single-homed links are synchronized to the MC-LAG peer through ICCP.

Both MC-LAG peers receive traffic on their incoming interface (IIF). The non-designated router receives traffic by way of the ICL interface, which acts as a multicast router (mrouter) interface.

If the designated router fails, the non-designated router has to build the entire forwarding tree (RPT and SPT), which can cause multicast traffic loss.

### ***PIM Operation with Dual DR Mode***

In this mode, both of the MC-LAG peers act as designated routers (active and backup) and send periodic join and prune messages upstream towards the RP, or source, and eventually join the RPT or SPT.

The primary MC-LAG peer forwards the multicast traffic to the receiver devices even if the standby MC-LAG peer has a smaller preference metric.

The standby MC-LAG peer also joins the forwarding tree and receives the multicast data. The standby MC-LAG peer drops the data because it has an empty outgoing interface list (OIL). When the standby MC-LAG peer detects the primary MC-LAG peer failure, it adds the receiver VLAN to the OIL, and starts to forward the multicast traffic.

To enable a multicast dual DR, issue the **set protocols pim interface interface-name dual-dr** command on the VLAN interfaces of each MC-LAG peer.

## MC-LAG Upgrade Guidelines

Upgrade the MC-LAG peers according to the following guidelines.



**NOTE:** After a reboot, the MC-AE interfaces come up immediately and might start receiving packets from the server. If routing protocols are enabled, and the routing adjacencies have not been formed, packets might be dropped.

To prevent this scenario, issue the `set interfaces interface-name aggregated-ether-options mc-ae init-delay-time time` to set a time by which the routing adjacencies are formed.

1. Make sure that both of the MC-LAG peers (node1 and node2) are in the active-active state using the following command on any one of the MC-LAG peers:

```
user@switch> show interfaces mc-ae id 1
Member Link           : ae0
Current State Machine's State: mcae active state
Local Status          : active<<<<<<<
Local State           : up
Peer Status           : active<<<<<<<
Peer State            : up
  Logical Interface    : ae0.0
  Topology Type        : bridge
  Local State          : up
  Peer State           : up
  Peer Ip/MCP/State    : 20.1.1.2 ae2.0 up
```

2. Upgrade node1 of the MC-LAG.

When node1 is upgraded it is rebooted, and all traffic is sent across the available LAG interfaces of node2, which is still up. The amount of traffic lost depends on how quickly the neighbor devices detect the link loss and rehash the flows of the LAG.

3. Verify that node1 is running the software you just installed. Issue the `show version` command.
4. Make sure that both nodes of the MC-LAG (node1 and node2) are in the active-active state after the reboot of node1.
5. Upgrade node2 of the MC-LAG.

Repeat step 1 through step 3 to upgrade node2.

### Related Documentation

- [Configuring Multichassis Link Aggregation on page 2681](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using VRRP on EX9200 Switches](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on EX9200 Switches](#)

## 802.1Q VLANs Overview

For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, the Junos OS supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or bridging domain.

### Related Documentation

- *Configuring Dynamic 802.1Q VLANs*
- *802.1Q VLAN IDs and Ethernet Interface Types*
- *Enabling VLAN Tagging*
- *Binding VLAN IDs to Logical Interfaces*
- *Configuring VLAN Encapsulation*
- *Configuring Extended VLAN Encapsulation*
- *Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs*
- *Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface*
- *Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance*
- *Specifying the Interface Over Which VPN Traffic Travels to the CE Router*
- *Specifying the Interface to Handle Traffic for a CCC*
- *Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface*
- *Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance*
- *Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit*
- *Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface*
- *Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface*
- *Configuring a Logical Interface for Access Mode*
- *Configuring a Logical Interface for Trunk Mode*
- *Configuring the VLAN ID List for a Trunk Interface*
- *Configuring a Trunk Interface on a Bridge Network*
- *Ethernet Interfaces Feature Guide for Routing Devices*

# Configuration

- [Configuration Tasks on page 2615](#)
- [Configuration Statements on page 2694](#)

## Configuration Tasks

---

- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 2619](#)
- [Port Role Configuration with the J-Web Interface \(with CLI References\) on page 2625](#)
- [Adding a Logical Unit Description to the Configuration on page 2629](#)
- [Disabling a Physical Interface on page 2630](#)
- [Disabling a Logical Interface on page 2631](#)
- [Configuring Flow Control on page 2632](#)
- [Configuring the Interface Address on page 2632](#)
- [Configuring the Interface Bandwidth on page 2637](#)
- [Configuring the Media MTU on page 2638](#)
- [Setting the Protocol MTU on page 2651](#)
- [Interface Ranges on page 2651](#)
- [Configuring Accounting for the Physical Interface on page 2660](#)
- [Configuring Accounting for the Logical Interface on page 2661](#)
- [Configuring Ethernet Loopback Capability on page 2662](#)
- [Configuring Gratuitous ARP on page 2663](#)
- [Configuring Static ARP Table Entries on page 2664](#)
- [Disabling the Transmission of Redirect Messages on an Interface on page 2665](#)
- [Configuring Restricted and Unrestricted Proxy ARP on page 2665](#)
- [Enabling or Disabling SNMP Notifications on Logical Interfaces on page 2666](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\) on page 2668](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 2671](#)

- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\) on page 2672](#)
- [Configuring Aggregated Ethernet Link Protection on page 2676](#)
- [Configuring Aggregated Ethernet Link Speed on page 2678](#)
- [Configuring Aggregated Ethernet Minimum Links on page 2679](#)
- [Configuring Multichassis Link Aggregation on page 2681](#)
- [Configuring Energy Efficient Ethernet on Interfaces \(CLI Procedure\) on page 2684](#)
- [Configuring Local Link Bias \(CLI Procedure\) on page 2686](#)
- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2686](#)
- [Configuring Tagged Aggregated Ethernet Interfaces on page 2688](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 2689](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 2689](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 2691](#)
- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 2691](#)
- [Tracing Operations of an Individual Router or Switch Interface on page 2692](#)
- [Tracing Operations of the Interface Process on page 2693](#)

## Configuring Gigabit Ethernet Interfaces (CLI Procedure)

---



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

---

An Ethernet interface must be configured for optimal performance in a high-traffic network. EX Series switches include a factory default configuration that:

- Enables all the network interfaces on the switch
- Sets a default interface mode (access)
- Sets default link settings
- Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching** (except on EX8200 switches and Virtual Chassis)
- Specifies Rapid Spanning Tree Protocol (RSTP) and Link Layer Discovery Protocol (LLDP)

This topic describes:

- [Configuring VLAN Options and Interface Mode on page 2617](#)
- [Configuring the Link Settings on page 2617](#)
- [Configuring the IP Options on page 2618](#)



### Configuring VLAN Options and Interface Mode

By default, when you boot a switch and use the factory default configuration, or when you boot the switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode and accept only untagged packets from the VLAN named **default**. You can optionally configure another VLAN and use that instead of **default**. You can also configure a port to accept untagged packets from the user-configured VLAN. For details on this concept (native VLAN), see [“Understanding Bridging and VLANs on EX Series Switches” on page 2245](#)

If you are connecting either a desktop phone, wireless access point or a security camera to a Power over Ethernet (PoE) port, you can configure some parameters for the PoE interface. PoE interfaces are enabled by default. For detailed information about PoE settings, see [“Configuring PoE on EX Series Switches \(CLI Procedure\)” on page 4440](#).

If you are connecting a device to other switches and to routers on the LAN, you need to assign the interface to a logical port and configure the logical port as a trunk port. See [“Port Role Configuration with the J-Web Interface \(with CLI References\)” on page 2625](#) for more information about port configuration.

If you are connecting to a server that contains virtual machines and a VEPA for packet aggregation from those virtual machines, configure the port as a tagged-access port. See [“Understanding Bridging and VLANs on EX Series Switches” on page 2245](#) for more information about tagged access.

To configure a 1-Gigabit, 10-Gigabit, or 40-Gigabit Ethernet interface for trunk port mode:

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

### Configuring the Link Settings

EX Series switches include a factory default configuration that enables interfaces with the link settings provided in [Table 268 on page 2617](#).

**Table 268: Factory Default Configuration Link Settings for EX Series Switches**

| Ethernet Interface                        | Autonegotiation | Flow Control | Link Mode                                       | Link Speed                                        |
|-------------------------------------------|-----------------|--------------|-------------------------------------------------|---------------------------------------------------|
| 1 gigabit                                 | Enabled         | Enabled      | Autonegotiation<br>(full duplex or half duplex) | Autonegotiation<br>(10 Mbps, 100 Mbps, or 1 Gbps) |
| 10 gigabit<br>(using a DAC cable)         | Enabled         | Enabled      | Full duplex                                     | 10 Gbps                                           |
| 10 gigabit<br>(using a fiber-optic cable) | Disabled        | Enabled      | Full duplex                                     | 10 Gbps                                           |
| 40 gigabit<br>(using a DAC cable)         | Enabled         | Enabled      | Full duplex                                     | 40 Gbps                                           |

Table 268: Factory Default Configuration Link Settings for EX Series Switches (*continued*)

| Ethernet Interface                        | Autonegotiation | Flow Control | Link Mode   | Link Speed |
|-------------------------------------------|-----------------|--------------|-------------|------------|
| 40 gigabit<br>(using a fiber-optic cable) | Disabled        | Enabled      | Full duplex | 40 Gbps    |



**NOTE:** On EX4300 switches, the interfaces operate in full duplex mode only.

To configure the link mode and speed settings for a 1-Gigabit, 10-Gigabit, or 40-Gigabit Ethernet interface:

```
[edit]
user@switch# set interfaces interface-name
```

To configure additional link settings for a 1-Gigabit, 10-Gigabit, or 40-Gigabit Ethernet interface:

```
[edit]
user@switch# set interfaces interface-name ether-options
```

For detailed information about the FPC, PIC, and port numbers used for EX Series switches, see [“Understanding Interface Naming Conventions on EX Series Switches” on page 2580](#).

Configurable link settings include:

- **802.3ad**—Specify an aggregated Ethernet bundle. See [“Configuring Aggregated Ethernet Links \(CLI Procedure\)” on page 2667](#).
- **auto-negotiation**—Enable or disable autonegotiation of flow control, link mode, and speed.
- **flow-control**—Enable or disable flow control.
- **link-mode**—Specify full duplex, half duplex, or autonegotiation. On EX4300 switches, the interfaces operate in full duplex mode only.
- **loopback**—Enable or disable loopback mode.
- **speed**—Specify 10 Mbps, 100 Mbps, 1 Gbps, or autonegotiation.

### Configuring the IP Options

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

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

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

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



**NOTE:** Access interfaces on EX4300 switches are set to family ethernet-switching by default. You might have to delete this or any other user-configured family setting before changing the setting to family inet or family inet6.

#### Related Documentation

- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 2619](#)
- [Monitoring Interface Status and Traffic on page 2835](#)
- [show interfaces ge- on page 2878](#)
- [show interfaces xe- on page 2911](#)
- [Understanding Interface Naming Conventions on EX Series Switches on page 2580](#)

## Configuring Gigabit Ethernet Interfaces (J-Web Procedure)

You can configure specific properties on your Ethernet interface to ensure optimal performance of your network in a high-traffic environment.

To configure properties on a Gigabit Ethernet interface, a 10-Gigabit Ethernet interface, and a 40-Gigabit Ethernet interface on an EX Series switch:

1. Select **Interfaces > Ports**.

The page that is displayed lists Gigabit Ethernet, 10-Gigabit Ethernet interfaces, and 40-Gigabit Ethernet interfaces, and their link statuses.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See “[Using the Commit Options to Commit Configuration Changes \(J-Web Procedure\)](#)” on page 628 for details about all commit options.

2. Select the interface you want to configure. For an EX8200 Virtual Chassis configuration, select the member and the FPC slot if the interface you want to configure is not listed under **Ports** in the top table on the page.

Details for the selected interface, such as administrative status, link status, speed, duplex, and flow control, are displayed in the **Details of port** table on the page.



**NOTE:** You can select multiple interfaces and modify their settings at the same time. However, while doing this, you cannot modify the IP address or enable or disable the administrative status of the selected interfaces.



**NOTE:** In the J-Web interface, you cannot configure interface ranges and interface groups.

3. Click **Edit** and select the set of options you want to configure first:

- Port Role—Enables you to assign a profile for the selected interface.



**NOTE:** When you select a particular port role, preconfigured port security parameters are set for the VLAN that the interface belongs to. For example, if you select the port role **Desktop**, the port security options **examine-dhcp** and **arp-inspection** are enabled on the VLAN that the interface belongs to. If there are interfaces in the VLAN that have static IP addresses, those interfaces might lose connectivity because those static IP addresses might not be present in the DHCP pool. Therefore, when you select a port role, ensure that the corresponding port security settings for the VLAN are applicable to the interface.

For basic information about port security features such as DHCP snooping (CLI option **examine-dhcp**) or dynamic ARP inspection (DAI) (CLI option **arp-inspection**), see “[Configuring Port Security \(J-Web Procedure\)](#)” on page 4541. For detailed descriptions of port security features, see the Port Security topics in the EX Series documentation at <http://www.juniper.net/techpubs/>.

Click **Details** to view the configuration parameters for the selected port role.

- VLAN—Enables you to configure VLAN options for the selected interface.
  - Link—Enables you to modify the following link options for the selected interface:
    - Speed
    - MTU
    - Autonegotiation
    - Flow Control
    - Duplex
    - Media Type
  - IP—Enables you to configure an IP address for the interface.
4. Configure the interface by configuring options in the selected option set. See [Table 269 on page 2621](#) for details of the options.
5. Repeat Steps 3 and 4 for the remaining option sets that you want to configure for the interface.



**NOTE:** To enable or disable the administrative status of a selected interface, click **Enable Port** or **Disable Port**.

Table 269: Port Edit Options

| Field                 | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port Role Options     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Port Role             | <p>Specifies a profile (role) to assign to the interface.</p> <p><b>NOTE:</b> After a port role is configured on the interface, you cannot specify VLAN options or IP options.</p> <p><b>NOTE:</b> Port roles are not supported by the et interfaces (40-Gigabit Ethernet interfaces) on EX4300 and EX4550 switches.</p> <p><b>NOTE:</b> Only the following port roles can be applied on EX8200 switch interfaces:</p> <ul style="list-style-type: none"> <li>• Default</li> <li>• Layer 2 uplink</li> <li>• Routed uplink</li> </ul> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Default               | <p>Applies the default role.</p> <p>The interface family is set to <b>ethernet-switching</b>, port mode is set to <b>access</b>, and RSTP is enabled.</p>                                                                                                                                                                                                                                                                                                                                                                             | <ol style="list-style-type: none"> <li>1. Click <b>Details</b> to view CLI commands for this role.</li> <li>2. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                            |
| Desktop               | <p>Applies the desktop role.</p> <p>The interface family is set to <b>ethernet-switching</b>, port mode is set to <b>access</b>, RSTP is enabled with the <b>edge</b> and <b>point-to-point</b> options, and port security parameters (MAC limit =1; dynamic ARP inspection and DHCP snooping enabled) are set.</p>                                                                                                                                                                                                                   | <ol style="list-style-type: none"> <li>1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface.</li> <li>2. Click <b>Details</b> to view CLI commands for this role.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                        |
| Desktop and Phone     | <p>Applies the desktop and phone role.</p> <p>The interface family is set to <b>ethernet-switching</b>, port mode is set to <b>access</b>, port security parameters (MAC limit =1; dynamic ARP Inspection and DHCP snooping enabled) are set, and recommended class-of-service (CoS) parameters are specified for forwarding classes, schedulers, and classifiers. See <a href="#">Table 270 on page 2624</a> for more CoS information.</p>                                                                                           | <ol style="list-style-type: none"> <li>1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface.<br/><br/>You can also select an existing VoIP VLAN configuration or a new VoIP VLAN configuration to be associated with the interface.<br/><br/><b>NOTE:</b> VoIP is not supported on EX8200 switches.</li> <li>2. Click <b>Details</b> to view CLI commands for this role.</li> <li>3. Click <b>OK</b>.</li> </ol> |
| Wireless Access Point | <p>Applies the wireless access point role.</p> <p>The interface family is set to <b>ethernet-switching</b>, port mode is set to <b>access</b>, and RSTP is enabled with the <b>edge</b> and <b>point-to-point</b> options.</p>                                                                                                                                                                                                                                                                                                        | <ol style="list-style-type: none"> <li>1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface. Type the VLAN ID for a new VLAN.</li> <li>2. Click <b>Details</b> to view CLI commands for this role.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                       |

Table 269: Port Edit Options (*continued*)

| Field                                                                                                      | Function                                                                                                                                                                                                                                            | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Routed Uplink                                                                                              | <p>Applies the routed uplink role.</p> <p>The interface family is set to <b>inet</b>, and recommended CoS parameters are set for schedulers and classifiers. See <a href="#">Table 270 on page 2624</a> for more CoS information.</p>               | <p>To specify an IPv4 address:</p> <ol style="list-style-type: none"> <li>1. Select the <b>IPv4 address</b> check box.</li> <li>2. Type an IP address—for example: <b>10.10.10.10</b>.</li> <li>3. Enter the subnet mask or address prefix. For example, 24 bits represents <b>255.255.255.0</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> <p>To specify an IPv6 address:</p> <ol style="list-style-type: none"> <li>1. Select the <b>IPv6 address</b> check box.</li> <li>2. Type an IP address—for example: <b>2001:ab8:85a3::8a2e:370:7334</b>.</li> <li>3. Enter the subnet mask or address prefix.</li> <li>4. Click <b>OK</b>.</li> </ol> <p><b>NOTE:</b> IPv6 is not supported on EX2200 VC switches.</p> |
| Layer 2 Uplink                                                                                             | <p>Applies the Layer 2 uplink role.</p> <p>The interface family is set to <b>ethernet-switching</b>, port mode is set to <b>trunk</b>, RSTP is enabled with the <b>point-to-point</b> option, and trusted DHCP is configured for port security.</p> | <ol style="list-style-type: none"> <li>1. For this port role, you can select a VLAN member and associate a native VLAN with the interface.</li> <li>2. Click <b>Details</b> to view CLI commands for this role.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| None                                                                                                       | Specifies that no port role is configured for the selected interface.                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>NOTE:</b> For an EX8200 switch, dynamic ARP inspection and DHCP snooping parameters are not configured. |                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| VLAN Options                                                                                               |                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

Table 269: Port Edit Options (*continued*)

| Field                   | Function                                                                                                                                                                                           | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port Mode               | Specifies the mode of operation for the interface: trunk or access.                                                                                                                                | <p>If you select <b>Trunk</b>, you can:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b> to add a VLAN member.</li> <li>2. Select the VLAN and click <b>OK</b>.</li> <li>3. (Optional) Associate a native VLAN with the interface.</li> <li>4. Click <b>OK</b>.</li> </ol> <p>If you select <b>Access</b>, you can:</p> <ol style="list-style-type: none"> <li>1. Select the VLAN member to be associated with the interface.</li> <li>2. (Optional) Associate a VoIP VLAN with the interface. Only a VLAN with a VLAN ID can be associated as a VoIP VLAN.</li> </ol> <p><b>NOTE:</b> VoIP is not supported on EX8200 switches.</p> <ol style="list-style-type: none"> <li>3. Click <b>OK</b>.</li> </ol> |
| Link Options            |                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| MTU (bytes)             | Specifies the maximum transmission unit size (MTU) for the interface.                                                                                                                              | Type a value from <b>256</b> through <b>9216</b> . The default MTU size for Gigabit Ethernet interfaces is <b>1514</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Speed                   | Specifies the speed for the mode.                                                                                                                                                                  | <p>Select one of the following values: <b>10 Mbps</b>, <b>100 Mbps</b>, <b>1000 Mbps</b>, or <b>Auto-Negotiation</b>.</p> <p><b>NOTE:</b> EX4300 switches supports <b>Auto-Negotiation 10M-100M</b> apart from the values mentioned above.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Duplex                  | Specifies the link mode.                                                                                                                                                                           | <p>Select one: <b>automatic</b>, <b>half</b>, or <b>full</b>.</p> <p><b>NOTE:</b> Link mode <b>half</b> is not supported on EX4300 switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Description             | <p>Describes the link.</p> <p><b>NOTE:</b> If the interface is part of a link aggregation group (LAG), only the <b>Description</b> option is enabled. Other Port Edit options are unavailable.</p> | Enter a brief description for the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Enable Auto Negotiation | Enables or disables autonegotiation.                                                                                                                                                               | Select the check box to enable autonegotiation, or clear the check box to disable it. By default, autonegotiation is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Enable Flow Control     | Enables or disables flow control.                                                                                                                                                                  | Select the check box to enable flow control to regulate the amount of traffic sent out of the interface, or clear the check box to disable flow control and permit unrestricted traffic. Flow control is enabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 269: Port Edit Options (*continued*)

| Field        | Function                                                                                                                                                          | Your Action                                                                                                                                                                                                                                                                                                                                                                   |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Media Type   | Specifies the media type selected.                                                                                                                                | Select the check box to enable the media type. Then select <b>Copper</b> or <b>Fiber</b> .                                                                                                                                                                                                                                                                                    |
| IP Options   |                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                               |
| IPv4 Address | Specifies an IPv4 address for the interface.<br><br><i>NOTE:</i> If the IPv4 Address check box is cleared, the interface still belongs to the <b>inet</b> family. | <ol style="list-style-type: none"> <li>1. Select the <b>IPv4 address</b> check box to specify an IPv4 address.</li> <li>2. Type an IP address—for example: <b>10.10.10.10</b>.</li> <li>3. Enter the subnet mask or address prefix. For example, 24 bits represents <b>255.255.255.0</b>.</li> <li>4. Click <b>OK</b>.</li> </ol>                                             |
| IPv6 Address | Specifies an IPv6 address for the interface.<br><br><i>NOTE:</i> If the IPv6 Address check box is cleared, the interface still belongs to the <b>inet</b> family. | <ol style="list-style-type: none"> <li>1. Select the <b>IPv6 address</b> check box to specify an IPv6 address.</li> <li>2. Type an IP address—for example: <b>2001:ab8:85a3::8a2e:370:7334</b>.</li> <li>3. Enter the subnet mask or address prefix.</li> <li>4. Click <b>OK</b>.</li> </ol> <p><i>NOTE:</i> IPv6 address is not supported on EX2200 and EX4500 switches.</p> |

Table 270: Recommended CoS Settings for Port Roles

| CoS Parameter         | Recommended Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding Classes    | <p>There are four forwarding classes:</p> <ul style="list-style-type: none"> <li>• <b>voice</b>—Queue number is set to 7.</li> <li>• <b>expedited-forwarding</b>—Queue number is set to 5.</li> <li>• <b>assured-forwarding</b>—Queue number is set to 1.</li> <li>• <b>best-effort</b>—Queue number is set to 0.</li> </ul>                                                                                                                                                                                                                                                     |
| Schedulers            | <p>The schedulers and their settings are:</p> <ul style="list-style-type: none"> <li>• Strict-priority—Transmission rate is set to 10 percent and buffer size to 5 percent.</li> <li>• Expedited-scheduler—Transmission rate is set to 30 percent, buffer size to 30 percent, and priority to <b>low</b>.</li> <li>• Assured-scheduler—Transmission rate is set to 25 percent, buffer size to 25 percent, and priority to <b>low</b>.</li> <li>• Best-effort scheduler—Transmission rate is set to 35 percent, buffer size to 40 percent, and priority to <b>low</b>.</li> </ul> |
| Scheduler maps        | When a desktop and phone, routed uplink, or Layer 2 uplink role is applied on an interface, the forwarding classes and schedulers are mapped using the scheduler map.                                                                                                                                                                                                                                                                                                                                                                                                            |
| ieee-802.1 classifier | Imports the default <b>ieee-802.1</b> classifier configuration and sets the loss priority to <b>low</b> for the code point 101 for the <b>voice</b> forwarding class.                                                                                                                                                                                                                                                                                                                                                                                                            |



Table 270: Recommended CoS Settings for Port Roles (*continued*)

| CoS Parameter                | Recommended Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dscp classifier              | Imports the default <b>dscp</b> classifier configuration and sets the loss priority to <b>low</b> for the code point 101110 for the <b>voice</b> forwarding class.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure)</a></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li> <li>• <a href="#">Monitoring Interface Status and Traffic on page 2835</a></li> <li>• <a href="#">EX Series Switches Interfaces Overview on page 2577</a></li> <li>• <a href="#">Junos OS CoS for EX Series Switches Overview on page 2042</a></li> <li>• <a href="#">Understanding Interface Naming Conventions on EX Series Switches on page 2580</a></li> </ul> |

### Port Role Configuration with the J-Web Interface (with CLI References)

When you configure Gigabit Ethernet interface properties with the J-Web interface (Configure > Interfaces) you can optionally select pre-configured port roles for those interfaces. When you select a role from the **Port Role** field and apply it to a port, the J-Web interface modifies the switch configuration using CLI commands. [Table 271 on page 2625](#) lists the CLI commands applied for each port role.



**NOTE:** If there is an existing port role configuration, it is cleared before the new port role configuration is applied.

Table 271: Port Role Configuration Summary

| Configuration Description                                 | CLI Commands                                                                          |
|-----------------------------------------------------------|---------------------------------------------------------------------------------------|
| Default Port Role                                         |                                                                                       |
| Set the port role to <b>Default</b> .                     | <code>set interfaces <i>interface</i> apply-macro juniper-port-profile Default</code> |
| Set port family to <b>ethernet-switching</b> .            | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching</code>         |
| Set port mode to <b>access</b> .                          | <code>port-mode access</code>                                                         |
| Enable RSTP if redundant trunk groups are not configured. | <code>delete protocols rstp interface <i>interface</i> disable</code>                 |
| Disable RSTP if redundant trunk groups are configured.    | <code>set protocols rstp interface <i>interface</i> disable</code>                    |
| Desktop Port Role                                         |                                                                                       |
| Set the port role to desktop.                             | <code>set interfaces <i>interface</i> apply-macro juniper-port-profile Desktop</code> |

Table 271: Port Role Configuration Summary (*continued*)

| Configuration Description                                                                                                                                                            | CLI Commands                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Set VLAN if new VLAN is specified.                                                                                                                                                   | <code>set vlans &lt;vlan name&gt; vlan-id &lt;vlan-id&gt;</code>                                                                                                                                                                                                                                             |
| Set port family to <b>ethernet-switching</b> .<br>Set Port Mode to <b>Access</b> .                                                                                                   | <code>set interfaces interface unit 0 family ethernet-switching<br/>port-mode access</code>                                                                                                                                                                                                                  |
| Set VLAN if new VLAN is specified.                                                                                                                                                   | <code>set interfaces interface unit 0 family ethernet-switching<br/>vlan members vlan-members</code>                                                                                                                                                                                                         |
| Set port security parameters.                                                                                                                                                        | <code>set ethernet-switching-options secure-access-port vlan<br/>MacTest arp-inspection</code>                                                                                                                                                                                                               |
| Set RSTP protocol with <b>edge</b> option.                                                                                                                                           | <code>set protocols rstp interface interface edge</code>                                                                                                                                                                                                                                                     |
| RSTP protocol is disabled if redundant trunk groups are configured.                                                                                                                  | <code>set protocols rstp interface interface disable</code>                                                                                                                                                                                                                                                  |
| Desktop and Phone Port Role                                                                                                                                                          |                                                                                                                                                                                                                                                                                                              |
| Set the port role to desktop and phone.                                                                                                                                              | <code>set interfaces interface apply-macro juniper-port-profile<br/>Desktop and Phone</code>                                                                                                                                                                                                                 |
| Set data VLAN if new VLAN is specified.<br>Set voice VLAN if new voice VLAN is specified.                                                                                            | <code>set vlans vlan-name vlan-id vlan id</code>                                                                                                                                                                                                                                                             |
| Set port family to <b>ethernet-switching</b> .<br>Set Port Mode to <b>access</b> .                                                                                                   | <code>set interfaces interface unit 0 family ethernet-switching<br/>port-mode access</code>                                                                                                                                                                                                                  |
| Set data VLAN on port stanza.                                                                                                                                                        | <code>set interfaces interface unit 0 family ethernet-switching<br/>vlan members vlan-members</code>                                                                                                                                                                                                         |
| Set port security parameters.                                                                                                                                                        | <code>set ethernet-switching-options secure-access-port vlan<br/>MacTest arp-inspection</code>                                                                                                                                                                                                               |
| Set VOIP VLAN.                                                                                                                                                                       | <code>set ethernet-switching-options voip interface interface.0<br/>vlan vlan vlan name</code>                                                                                                                                                                                                               |
| Set class of service parameters<br><b>SCHEDULER_MAP=juniper-port-profile-map</b><br><b>IEEE_CLASSIFIER=juniper-ieee-classifier</b><br><b>DSCP_CLASSIFIER=juniper-dscp-classifier</b> | <code>set class-of-service interfaces interfaces scheduler-map<br/>juniper-port-profile-map<br/>set class-of-service interfaces interface unit 0<br/>classifiers ieee-802.1 juniper_ieee_classifier<br/>set class-of-service interfaces interface unit 0 classifiers<br/>dscp juniper-dscp-classifier</code> |
| Set CoS Configuration                                                                                                                                                                | Refer <a href="#">Table 272 on page 2628</a> for details.                                                                                                                                                                                                                                                    |
| Wireless Access Point Port Role                                                                                                                                                      |                                                                                                                                                                                                                                                                                                              |
| Set the port role to wireless access point.                                                                                                                                          | <code>set interfaces interface apply-macro juniper-port-profile<br/>Wireless Access Point</code>                                                                                                                                                                                                             |
| Set VLAN on VLANs stanza.                                                                                                                                                            | <code>set vlans vlan name vlan-id vlan-id</code>                                                                                                                                                                                                                                                             |

Table 271: Port Role Configuration Summary (*continued*)

| Configuration Description                                                                                                                                                               | CLI Commands                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Set port family to <b>ethernet-switching</b><br>Set port mode to <b>Access</b> .                                                                                                        | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching port-mode access</code>                                                                                                                                                                                                                                                |
| Set VLAN on port stanza.                                                                                                                                                                | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching vlan members <i>vlan-members</i></code>                                                                                                                                                                                                                                |
| Set RSTP protocol with edge option.                                                                                                                                                     | <code>set protocols rstp interface <i>interface</i> edge</code>                                                                                                                                                                                                                                                                               |
| RSTP protocol is disabled if redundant trunk groups are configured.                                                                                                                     | <code>set protocols rstp interface <i>interface</i> disable</code>                                                                                                                                                                                                                                                                            |
| Routed Uplink Port Role                                                                                                                                                                 |                                                                                                                                                                                                                                                                                                                                               |
| Set the port role to Routed Uplink.                                                                                                                                                     | <code>set interfaces <i>interface</i> apply-macro juniper-port-profile Routed Uplink</code>                                                                                                                                                                                                                                                   |
| Set port family to <b>inet</b> .<br>Set IP address on the port.                                                                                                                         | <code>set interfaces <i>interface</i> unit 0 family inet address <i>ipaddress</i></code>                                                                                                                                                                                                                                                      |
| Set class-of-service parameters<br><b>SCHEDULER_MAP=</b> juniper-port-profile-map<br><b>IEEE_CLASSIFIER=</b> juniper-ieee-classifier<br><b>DSCP_CLASSIFIER=</b> juniper-dscp-classifier | <code>set class-of-service interfaces <i>interfaces</i> scheduler-map juniper-port-profile-map</code><br><code>set class-of-service interfaces <i>interface</i> unit 0 classifiers ieee-802.1 juniper_ieee_classifier</code><br><code>set class-of-service interfaces <i>interface</i> unit 0 classifiers dscp juniper-dscp-classifier</code> |
| Set CoS configuration                                                                                                                                                                   | Refer <a href="#">Table 272 on page 2628</a> for details.                                                                                                                                                                                                                                                                                     |
| Layer 2 Uplink Port Role                                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                               |
| Set the port role to <b>Layer 2 Uplink</b> .                                                                                                                                            | <code>set interfaces <i>interface</i> apply-macro juniper-port-profile Layer2 Uplink</code>                                                                                                                                                                                                                                                   |
| Set port family to <b>ethernet-switching</b><br>Set port mode to <b>trunk</b> .                                                                                                         | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching port-mode trunk</code>                                                                                                                                                                                                                                                 |
| Set Native VLAN name.                                                                                                                                                                   | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching native-vlan-id <i>vlan-name</i></code>                                                                                                                                                                                                                                 |
| Set the port as part of all valid VLANs; "valid" refers to all VLANs except native VLAN and voice VLANs.                                                                                | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching vlan members <i>vlan-members</i></code>                                                                                                                                                                                                                                |
| Set port security parameter.                                                                                                                                                            | <code>set ethernet-switching-options secure-access-port dhcp-trusted</code>                                                                                                                                                                                                                                                                   |
| Set RSTP protocol with point-to-point option.                                                                                                                                           | <code>set protocols rstp interface <i>interface</i> mode point-to-point</code>                                                                                                                                                                                                                                                                |
| Disable RSTP if redundant trunk groups are configured.                                                                                                                                  | <code>set protocols rstp interface <i>interface</i> disable</code>                                                                                                                                                                                                                                                                            |

Table 271: Port Role Configuration Summary (*continued*)

| Configuration Description                                                                                                                                                                            | CLI Commands                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Set class-of-service parameters.<br><br><b>SCHEDULER_MAP=</b> juniper-port-profile-map<br><br><b>IEEE_CLASSIFIER=</b> juniper_ieee_classifier<br><br><b>DSCP_CLASSIFIER=</b> juniper_dscp_classifier | <code>set class-of-service interfaces <i>interfaces</i> scheduler-map juniper-port-profile-map</code><br><code>set class-of-service interfaces <i>interface</i> unit 0 classifiers ieee-802.1 juniper_ieee_classifier</code><br><code>set class-of-service interfaces <i>interface</i> unit 0 classifiers dscp juniper-dscp-classifier</code> |
| Set CoS configuration                                                                                                                                                                                | Refer to <a href="#">Table 272 on page 2628</a> for details.                                                                                                                                                                                                                                                                                  |

[Table 272 on page 2628](#) lists the CLI commands for the recommended CoS settings that are committed when the CoS configuration is set.

Table 272: Recommended CoS Settings for Port Roles

| CoS Parameter                    | CLI Command                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding Classes               |                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>voice</b>                     | <code>set class-of-service forwarding-classes class voice queue-num 7</code>                                                                                                                                                                                                                                                                                                              |
| <b>expedited-forwarding</b>      | <code>set class-of-service forwarding-classes class expedited-forwarding queue-num 5</code>                                                                                                                                                                                                                                                                                               |
| <b>assured-forwarding</b>        | <code>set class-of-service forwarding-classes class assured-forwarding queue-num 1</code>                                                                                                                                                                                                                                                                                                 |
| <b>best-effort</b>               | <code>set class-of-service forwarding-classes class best-effort queue-num 0</code>                                                                                                                                                                                                                                                                                                        |
| Schedulers                       |                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>strict-priority-scheduler</b> | <p>The CLI commands are:</p> <ul style="list-style-type: none"> <li><code>set class-of-service schedulers strict-priority-scheduler transmit-rate percent 10</code></li> <li><code>set class-of-service schedulers strict-priority-scheduler buffer-size percent 5</code></li> <li><code>set class-of-service schedulers strict-priority-scheduler priority strict-high</code></li> </ul> |
| <b>expedited-scheduler</b>       | <p>The CLI commands are:</p> <ul style="list-style-type: none"> <li><code>set class-of-service schedulers expedited-scheduler transmit-rate percent 30</code></li> <li><code>set class-of-service schedulers expedited-scheduler buffer-size percent 30</code></li> <li><code>set class-of-service schedulers expedited-scheduler priority low</code></li> </ul>                          |
| <b>assured-scheduler</b>         | <p>The CLI commands are:</p> <ul style="list-style-type: none"> <li><code>set class-of-service schedulers assured-scheduler transmit-rate percent 25</code></li> <li><code>set class-of-service schedulers strict-priority-scheduler buffer-size percent 25</code></li> <li><code>set class-of-service schedulers strict-priority-scheduler priority low</code></li> </ul>                |

Table 272: Recommended CoS Settings for Port Roles (*continued*)

| CoS Parameter                | CLI Command                                                                                                                                                                                                                                                                                                                 |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>best-effort-scheduler</b> | <p>The CLI commands are:</p> <pre>set class-of-service schedulers best-effort-scheduler transmit-rate percent 35 set class-of-service schedulers best-effort-scheduler buffer-size percent 40 set class-of-service schedulers best-effort-scheduler priority low</pre>                                                      |
| <b>Classifiers</b>           | <p>The classifiers are:</p> <pre>set class-of-service classifiers ieee-802.1 juniper_ieee_classifier import default forwarding-class voice loss-priority low code-points 101 set class-of-service classifiers dscp juniper_dscp_classifier import default forwarding-class voice loss-priority low code-points 101110</pre> |

**Related Documentation**

- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 2619](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)

## Adding a Logical Unit Description to the Configuration

You can include a text description of each logical unit in the configuration file. Any descriptive text you include is displayed in the output of the **show interfaces** commands, and is also exposed in the **ifAlias** Management Information Base (MIB) object. It has no impact on the interface's configuration. To add a text description, include the **description** statement:

```
description text;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

The description can be a single line of text. If the text contains spaces, enclose it in quotation marks.



**NOTE:** You can configure the extended DHCP relay to include the interface description in the option 82 Agent Circuit ID suboption. See [“Using DHCP Relay Agent Option 82 Information” on page 1464](#) in the *Junos OS Subscriber Management and Services Library*.

For information about describing physical interfaces, see *Configuring Interface Description*.

## Disabling a Physical Interface

You can disable a physical interface, marking it as being down, without removing the interface configuration statements from the configuration. To do this, include the **disable** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
disable;
```



**CAUTION:** Dynamic subscribers and logical interfaces use physical interfaces for connection to the network. The Junos OS allows you to set the interface to disable and commit the change while dynamic subscribers and logical interfaces are still active. This action results in the loss of all subscriber connections on the interface. Use care when disabling interfaces.



**NOTE:** On the router, when you use the **disable** statement at the **edit interfaces** hierarchy level, depending on the PIC type, the interface might or might not turn off the laser. Older PIC transceivers do not support turning off the laser, but newer Gigabit Ethernet PICs with SFP and XFP transceivers do support it and the laser will be turned off when the interface is disabled.

Table 273: Effect of set interfaces disable <interface\_name> on T series PICs

| PIC Model Number  | PIC Description                                                               | Type of PIC | Behaviour             |
|-------------------|-------------------------------------------------------------------------------|-------------|-----------------------|
| PF-12XGE-SFPP     | 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (T4000 Router)                      | 5           | Tx laser disabled     |
| PF-24XGE-SFPP     | 10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (T4000 Router) | 5           | Tx laser disabled     |
| PF-1CGE-CFP       | 100-Gigabit Ethernet PIC with CFP (T4000 Router)                              | 5           | Tx laser disabled     |
| PD-4XGE-XFP       | 10-Gigabit Ethernet, 4-port LAN/WAN XFP                                       | 4           | Tx laser disabled     |
| PD-5-10XGE-SFPP   | 10-Gigabit LAN/WAN with SFP+                                                  | 4           | Tx laser disabled     |
| PD-1XLE-CFP       | 40-Gigabit with CFP                                                           | 4           | Tx laser disabled     |
| PD-1CE-CFP-FPC4   | 100-Gigabit with CFP                                                          | 4           | Tx laser disabled     |
| PD-TUNNEL         | 40-Gigabit Tunnel Services                                                    | 4           | NA                    |
| PD-4OC192-SON-XFP | OC192/STM64, 4-port XFP                                                       | 4           | Tx laser not disabled |
| PD-1OC768-SON-SR  | OC768c/STM256, 1-port                                                         | 4           | Tx laser not disabled |



**WARNING:** Do not stare into the laser beam or view it directly with optical instruments even if the interface has been disabled.

### Example: Disabling a Physical Interface

Sample interface configuration:

```
[edit interfaces]
user@host# show
ge-0/3/2 {
  unit 0 {
    description CE2-to-PE1;
    family inet {
      address 20.1.1.6/24;
    }
  }
}
```

Disabling the interface:

```
[edit interfaces]
user@host# set ge-0/3/2 disable
```

Verifying the interface configuration:

```
[edit interfaces]
user@host# show
ge-0/3/2 {
  disable; # Interface is marked as disabled.
  unit 0 {
    description CE2-to-PE1;
    family inet {
      address 20.1.1.6/24;
    }
  }
}
```

### Disabling a Logical Interface

You can unconfigure a logical interface, effectively disabling that interface, without removing the logical interface configuration statements from the configuration. To do this, include the **disable** statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

When an interface is disabled, a route (pointing to the reserved target “**REJECT**”) with the IP address of the interface and a 32-bit subnet mask is installed in the routing table. See *Routing Protocols*.

## Configuring Flow Control

By default, the router or switch imposes flow control to regulate the amount of traffic sent out on a Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interface. Flow control is not supported on the 4-port Fast Ethernet PIC. This is useful if the remote side of the connection is a Fast Ethernet or Gigabit Ethernet switch.

You can disable flow control if you want the router or switch to permit unrestricted traffic. To disable flow control, include the **no-flow-control** statement:

```
no-flow-control;
```

To explicitly reinstate flow control, include the **flow-control** statement:

```
flow-control;
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* ether-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* ggether-options]



**NOTE:** On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.

### Related Documentation

- [flow-control on page 2775](#)
- [Ethernet Interfaces Overview](#)
- [EX Series Switches Interfaces Overview on page 2577](#)
- [Ethernet Interfaces Feature Guide for Routing Devices](#)

## Configuring the Interface Address

You assign an address to an interface by specifying the address when configuring the protocol family. For the **inet** or **inet6** family, configure the interface IP address. For the **iso** family, configure one or more addresses for the loopback interface. For the **ccc**, **ethernet-switching**, **tcc**, **mpls**, **tnp**, and **vpls** families, you never configure an address.



**NOTE:** The point-to-point (PPP) address is taken from the loopback interface address that has the primary attribute. When the loopback interface is configured as an unnumbered interface, it takes the primary address from the donor interface.



To assign an address to an interface, include the **address** statement:

```
address address {
  broadcast address;
  destination address;
  destination-profile name;
  eui-64;
  preferred;
  primary;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

In the **address** statement, specify the network address of the interface.

For each address, you can optionally configure one or more of the following:

- Broadcast address for the interface subnet—Specify this in the **broadcast** statement; this applies only to Ethernet interfaces, such as the management interface **fxp0**, **em0**, or **me0** the Fast Ethernet interface, and the Gigabit Ethernet interface.
- Address of the remote side of the connection (for point-to-point interfaces only)—Specify this in the **destination** statement.
- PPP properties to the remote end—Specify this in the **destination-profile** statement. You define the profile at the [edit access group-profile *name* ppp] hierarchy level (for point-to-point interfaces only).
- Whether the router or switch automatically generates the host number portion of interface addresses—The **eui-64** statement applies only to interfaces that carry IPv6 traffic, in which the prefix length of the address is 64 bits or less, and the low-order 64 bits of the address are zero. This option does not apply to the loopback interface (**lo0**) because IPv6 addresses configured on the loopback interface must have a 128-bit prefix length.
- Whether this address is the preferred address—Each subnet on an interface has a preferred local address. If you configure more than one address on the same subnet, the preferred local address is chosen by default as the source address when you originate packets to destinations on the subnet.

By default, the preferred address is the lowest-numbered address on the subnet. To override the default and explicitly configure the preferred address, include the **preferred** statement when configuring the address.

- Whether this address is the primary address—Each interface has a primary local address. If an interface has more than one address, the primary local address is used by default as the source address when you send packets from an interface where the destination provides no information about the subnet (for example, some **ping** commands).

By default, the primary address on an interface is the lowest-numbered non-127 (in other words, non-loopback) preferred address on the interface. To override the default and explicitly configure the preferred address, include the **primary** statement when configuring the address.



**NOTE:** If you configure a duplicate IP address on an interface, even when the earlier interface with that IP address is disabled, a **Warning** message is added to the syslog and not displayed on the screen. Do not configure the same IP address of a disabled interface on another interface.

- [Configuring Interface IPv4 Addresses on page 2634](#)
- [Configuring Interface IPv6 Addresses on page 2637](#)

---

### Configuring Interface IPv4 Addresses

---

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

To configure an IPv4 address on routers and switches running Junos OS, use the **edit interface *interface-name* unit *number* family inet address *a.b.c.d/nn*** statement at the **[edit interfaces]** hierarchy level.



**NOTE:** Juniper Networks routers and switches support /31 destination prefixes when used in point-to-point Ethernet configurations; however, they are not supported by many other devices, such as hosts, hubs, routers, or switches. You must determine if the peer system also supports /31 destination prefixes before configuration.

---

### *Operational Behavior of Interfaces when the Same IPv4 Address is Assigned to Them*

You can configure the same IPv4 address on multiple physical interfaces. When you assign the same IPv4 address to multiple physical interfaces, the operational behavior of those interfaces differs, depending on whether they are implicitly or explicitly point-to-point.



**NOTE:** By default, all interfaces are assumed to be point-to-point (PPP) interfaces. For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, you can explicitly configure an interface to be a point-to-point connection.



**NOTE:** If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration, the remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
}
```

The following examples show the sample configuration of assigning the same IPv4 address to implicitly and explicitly point-to-point interfaces, and their corresponding **show interfaces terse** command outputs to see their operational status.

#### Configuring same IPv4 address on implicitly PPP interfaces:

```
[edit]
user@host# show
ge-0/1/0 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
ge-3/0/1 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
```

The sample output shown below for the above configuration reveals that only **ge-0/1/0.0** was assigned the same IPv4 address **200.1.1.1/24** and its **link** state was **up**, while **ge-3/0/1.0** was not assigned the IPv4 address, though its **link** state was **up**, which means that it will be operational only when it gets a unique IPv4 address other than **200.1.1.1/24**.

```
user@host> show interfaces terse ge*
Interface           Admin Link Proto  Local          Remote
ge-0/1/0            up    up    up
ge-0/1/0.0          up    up    inet    200.1.1.1/24
                    multiservice
ge-0/1/1            up    down
ge-3/0/0            up    down
ge-3/0/1            up    up
ge-3/0/1.0          up    up    inet
                    multiservice
```

#### Configuring same IPv4 address on explicitly PPP interfaces:

```
[edit]
user@host# show
so-0/0/0 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
so-0/0/3 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
```

The sample output shown below for the above configuration reveals that both **so-0/0/0.0** and **so-0/0/3.0** were assigned the same IPv4 address **200.1.1.1/24** and that their **link** states were **down**, which means that to make them operational at least one of them will have to be configured with a unique IPv4 address other than **200.1.1.1/24**.

```
user@host> show interfaces terse so*
Interface           Admin Link Proto  Local          Remote
so-0/0/0            up    up
so-0/0/0.0          up    down inet    200.1.1.1/24
so-0/0/1            up    up
so-0/0/2            up    down
so-0/0/3            up    up
so-0/0/3.0          up    down inet    200.1.1.1/24
so-1/1/0            up    down
so-1/1/1            up    down
so-1/1/2            up    up
so-1/1/3            up    up
so-2/0/0            up    up
so-2/0/1            up    up
so-2/0/2            up    up
so-2/0/3            up    down
```

## Configuring Interface IPv6 Addresses



**NOTE:** IPv6 is not currently supported for the QFX Series.

You represent IP version 6 (IPv6) addresses in hexadecimal notation using a colon-separated list of 16-bit values.

You assign a 128-bit IPv6 address to an interface by including the **address** statement:

```
address aaaa:bbbb:...:zzzz/nn;
```



**NOTE:** You cannot configure a subnet zero IPv6 address because RFC 2461 reserves the subnet-zero address for anycast addresses, and Junos OS complies with the RFC.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet6]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6]

The double colon (::) represents all bits set to 0, as shown in the following example:

```
interfaces fe-0/0/1 {
  unit 0 {
    family inet6 {
      address fec0:1:1::2/64;
    }
  }
}
```



**NOTE:** You must manually configure the router or switch advertisement and advertise the default prefix for autoconfiguration to work on a specific interface.

### Related Documentation

- [Configuring IPCP Options](#)
- [Configuring Default, Primary, and Preferred Addresses and Interfaces](#)

## Configuring the Interface Bandwidth

By default, the Junos OS uses the physical interface's speed for the MIB-II object, **ifSpeed**. You can configure the logical unit to populate the **ifSpeed** variable by configuring a bandwidth value for the logical interface. The **bandwidth** statement sets an informational-only parameter; you cannot adjust the actual bandwidth of an interface with this statement.



**NOTE:** We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the **bandwidth** statement affects how the interface cost is calculated for a dynamic routing protocol, such as OSPF. By default, the interface cost for a dynamic routing protocol is calculated using the following formula:

$$\text{cost} = \text{reference-bandwidth} / \text{bandwidth},$$

where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the **bandwidth** statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

To configure the bandwidth value for a logical interface, include the **bandwidth** statement:

**bandwidth** *rate*;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

**rate** is the peak rate, in bps or cps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation **c**; values expressed in cells per second are converted to bits per second using the formula 1 cps = 384 bps. The value can be any positive integer. The **bandwidth** statement is valid for all logical interfaces, except multilink interfaces.

## Configuring the Media MTU

The media maximum transmission unit (MTU) is the largest data unit that can be forwarded without fragmentation.

This topic contains the following sections:

- [Media MTU Overview on page 2639](#)
- [How to Configure the Media MTU on page 2640](#)
- [Encapsulation Overhead by Encapsulation Type on page 2641](#)
- [Media MTU Sizes by Interface Type for M5 and M7i Routers with CFEB, M10 and M10i Routers with CFEB, and M20 and M40 Routers on page 2642](#)
- [Media MTU Sizes by Interface Type for M40e Routers on page 2642](#)
- [Media MTU Sizes by Interface Type for M160 Routers on page 2644](#)
- [Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers on page 2644](#)
- [Media MTU Sizes by Interface Type for MX Series Routers on page 2645](#)

- [Media MTU Sizes by Interface Type for T320 Routers on page 2646](#)
- [Media MTU Sizes by Interface Type for T640 Platforms on page 2646](#)
- [Media MTU Sizes by Interface Type for J2300 Platforms on page 2647](#)
- [Media MTU Sizes by Interface Type for J4300 and J6300 Platforms on page 2647](#)
- [Media MTU Sizes by Interface Type for J4350 and J6350 Platforms on page 2648](#)
- [Media MTU Sizes by Interface Type for EX Series Switches and ACX Series Routers on page 2650](#)
- [Media MTU Sizes by Interface Type for PTX Series Packet Transport Routers on page 2650](#)

---

### Media MTU Overview

The default media MTU size used on a physical interface depends on the encapsulation used on that interface. In some cases, the default IP Protocol MTU depends on whether the protocol used is IP version 4 (IPv4) or International Organization for Standardization (ISO).

The default media MTU is calculated as follows:

**Default media MTU = Default IP MTU + encapsulation overhead**

When you are configuring point-to-point connections, the MTU sizes on both sides of the connections must be the same. Also, when you are configuring point-to-multipoint connections, all interfaces in the subnet must use the same MTU size. For details about encapsulation overhead, see [“Encapsulation Overhead by Encapsulation Type” on page 2641](#).



**NOTE:** The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the media MTU. For example, the media MTU for a Gigabit Ethernet Version 2 interface is specified as 1514 bytes, but the largest possible frame size is actually 1518 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

The physical MTU for Ethernet interfaces does not include the 4-byte frame check sequence (FCS) field of the Ethernet frame.

A SONET/SDH interface operating in concatenated mode has a “c” added to the rate descriptor. For example, a concatenated OC48 interface is referred to as OC48c.

If you do not configure an MPLS MTU, the Junos OS derives the MPLS MTU from the physical interface MTU. From this value, the software subtracts the encapsulation-specific overhead and space for the maximum number of labels that might be pushed in the Packet Forwarding Engine. Currently, the software provides for three labels of four bytes each, for a total of 12 bytes.

In other words, the formula used to determine the MPLS MTU is the following:

$$\text{MPLS MTU} = \text{physical interface MTU} - \text{encapsulation overhead} - 12$$

If you configure an MTU value by including the `mtu` statement at the `[edit interfaces interface-name unit logical-unit-number family mpls]` hierarchy level, the configured value is used.

---

### How to Configure the Media MTU

---

To modify the default media MTU size for a physical interface, include the `mtu` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
  mtu bytes;
```

If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead.



**NOTE:** Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

You configure the protocol MTU by including the `mtu` statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family family]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family]`



Because tunnel services interfaces are considered logical interfaces, you cannot configure the MTU setting for the physical interface. This means you cannot include the **mtu** statement at the **[edit interfaces *interface-name*]** hierarchy level for the following interface types: generic routing encapsulation (**gr-**), IP-IP (**ip-**), loopback (**lo-**), link services (**ls-**), multilink services (**ml-**), and multicast (**pe-**, **pd-**). You can, however, configure the protocol MTU on tunnel interfaces, as described in [“Setting the Protocol MTU” on page 2651](#).

### Encapsulation Overhead by Encapsulation Type

**Table 274: Encapsulation Overhead by Encapsulation Type**

| Interface Encapsulation                                                     | Encapsulation Overhead (Bytes) |
|-----------------------------------------------------------------------------|--------------------------------|
| 802.1Q/Ethernet 802.3                                                       | 21                             |
| 802.1Q/Ethernet Subnetwork Access Protocol (SNAP)                           | 26                             |
| 802.1Q/Ethernet version 2                                                   | 18                             |
| ATM Cell Relay                                                              | 4                              |
| ATM permanent virtual connection (PVC)                                      | 12                             |
| Cisco HDLC                                                                  | 4                              |
| Ethernet 802.3                                                              | 17                             |
| Ethernet circuit cross-connect (CCC) and virtual private LAN service (VPLS) | 4                              |
| Ethernet over ATM                                                           | 32                             |
| Ethernet SNAP                                                               | 22                             |
| Ethernet translational cross-connect (TCC)                                  | 18                             |
| Ethernet version 2                                                          | 14                             |
| Extended virtual local area network (VLAN) CCC and VPLS                     | 4                              |
| Extended VLAN TCC                                                           | 22                             |
| Frame Relay                                                                 | 4                              |
| PPP                                                                         | 4                              |
| VLAN CCC                                                                    | 4                              |
| VLAN VPLS                                                                   | 4                              |
| VLAN TCC                                                                    | 22                             |

### Media MTU Sizes by Interface Type for M5 and M7i Routers with CFEB, M10 and M10i Routers with CFEB, and M20 and M40 Routers

**Table 275: Media MTU Sizes by Interface Type for M5 and M7i Routers with CFEB, M10 and M10i Routers with CFEB, and M20 and M40 Routers**

| Interface Type                                | Default Media MTU (Bytes) | Maximum MTU (Bytes)                                                                                                                           | Default IP Protocol MTU (Bytes) |
|-----------------------------------------------|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| Adaptive Services (MTU size not configurable) | 9192                      | N/A                                                                                                                                           | N/A                             |
| ATM                                           | 4482                      | 9192                                                                                                                                          | 4470                            |
| E1/T1                                         | 1504                      | 9192                                                                                                                                          | 1500                            |
| E3/T3                                         | 4474                      | 9192                                                                                                                                          | 4470                            |
| Fast Ethernet                                 | 1514                      | 1533 (4-port)<br>1532 (8-port)<br>1532 (12-port)<br><br><i>NOTE:</i> The maximum MTU for two 100Base-TX Fast Ethernet port FIC is 9192 bytes. | 1500 (IPv4),<br>1497 (ISO)      |
| Gigabit Ethernet                              | 1514                      | 9192<br><br><i>NOTE:</i> The maximum MTU for one Gigabit Ethernet port FIC is 9192 bytes.                                                     | 1500 (IPv4),<br>1497 (ISO)      |
| Serial                                        | 1504                      | 9192                                                                                                                                          | 1500 (IPv4),<br>1497 (ISO)      |
| SONET/SDH                                     | 4474                      | 9192                                                                                                                                          | 4470                            |

### Media MTU Sizes by Interface Type for M40e Routers

**Table 276: Media MTU Sizes by Interface Type for M40e Routers**

| Interface Type                                | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes) |
|-----------------------------------------------|---------------------------|---------------------|---------------------------------|
| Adaptive Services (MTU size not configurable) | 9192                      | N/A                 | N/A                             |
| ATM                                           | 4482                      | 9192                | 4470                            |

**Table 276: Media MTU Sizes by Interface Type for M40e Routers** *(continued)*

| Interface Type   | Default Media MTU (Bytes) | Maximum MTU (Bytes)                                                                                                                                                                                                                                                                                       | Default IP Protocol MTU (Bytes) |
|------------------|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| E1/T1            | 1504                      | 4500                                                                                                                                                                                                                                                                                                      | 1500                            |
| E3/T3            | 4474                      | 4500<br>9192 (4-port)                                                                                                                                                                                                                                                                                     | 4470                            |
| E3/DS3 IQ        | 4474                      | 9192                                                                                                                                                                                                                                                                                                      | 4470                            |
| Fast Ethernet    | 1514                      | 1533                                                                                                                                                                                                                                                                                                      | 1500 (IPv4),<br>1497 (ISO)      |
| Gigabit Ethernet | 1514                      | 9192 (1- or 2-port)<br>9192 (4-port)                                                                                                                                                                                                                                                                      | 1500 (IPv4),<br>1497 (ISO)      |
| Serial           | 1504                      | 9192                                                                                                                                                                                                                                                                                                      | 1500 (IPv4),<br>1497 (ISO)      |
| SONET/SDH        | 4474                      | 4500 (1-port nonconcatenated)<br>9192 (4-port OC3)<br>9192 (4-port OC3c)<br>4500 (1-port OC12)<br>4500 (4-port OC12)<br>4500 (4-port OC12c)<br>4500 (1-port OC48)<br>9192 (2-port OC3)<br>9192 (2-port OC3c)<br>9192 (1-port OC12c)<br>9192 (1-port OC48c)<br>4500 (1-port OC192)<br>9192 (1-port OC192c) | 4470                            |

## Media MTU Sizes by Interface Type for M160 Routers

Table 277: Media MTU Sizes by Interface Type for M160 Routers

| Interface Type                                   | Default Media MTU (Bytes) | Maximum MTU (Bytes)                                                   | Default IP Protocol MTU (Bytes) |
|--------------------------------------------------|---------------------------|-----------------------------------------------------------------------|---------------------------------|
| Adaptive Services<br>(MTU size not configurable) | 9192                      | N/A                                                                   | N/A                             |
| ATM                                              | 4482                      | 9192                                                                  | 4470                            |
| E1/T1                                            | 1504                      | 4500                                                                  | 1500                            |
| E3/T3                                            | 4474                      | 4500                                                                  | 4470                            |
| E3/DS3 IQ                                        | 4474                      | 9192                                                                  | 4470                            |
| Fast Ethernet                                    | 1514                      | 1533                                                                  | 1500 (IPv4),<br>1497 (ISO)      |
| Gigabit Ethernet                                 | 1514                      | 9192 (1- or 2-port)<br>4500 (4-port)                                  | 1500 (IPv4),<br>1497 (ISO)      |
| Serial                                           | 1504                      | 9192                                                                  | 1500 (IPv4),<br>1497 (ISO)      |
| SONET/SDH                                        | 4474                      | 4500 (1-port nonconcatenated)<br>9192 (1- or 2-port)<br>4500 (4-port) | 4470                            |

## Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers

Table 278: Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers

| Interface Type      | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes) |
|---------------------|---------------------------|---------------------|---------------------------------|
| ATM2 IQ             | 4482                      | 9192                | 4470                            |
| Channelized DS3 IQ  | 4471                      | 4500                | 4470                            |
| Channelized E1 IQ   | 1504                      | 4500                | 1500                            |
| Channelized OC12 IQ | 4474                      | 9192                | 4470                            |

**Table 278: Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers (*continued*)**

| Interface Type             | Default Media MTU (Bytes) | Maximum MTU (Bytes)                         | Default IP Protocol MTU (Bytes) |
|----------------------------|---------------------------|---------------------------------------------|---------------------------------|
| Channelized STM1 IQ        | 4474                      | 9192                                        | 4470                            |
| DS3                        | 4471                      | 4500                                        | 4470                            |
| E1                         | 1504                      | 4500                                        | 1500                            |
| E3 IQ                      | 4471                      | 4500                                        | 4470                            |
| Fast Ethernet              | 1514                      | 1533 (4-port)<br>1532 (8-, 12- and 48-port) | 1500 (IPv4),<br>1497 (ISO)      |
| Gigabit Ethernet           | 1514                      | 9192                                        | 1500 (IPv4),<br>1497 (ISO)      |
| SONET/SDH                  | 4474                      | 9192                                        | 4470                            |
| T1                         | 1504                      | 4500                                        | 1500                            |
| CT3 IQ<br>(excluding M120) | 4474                      | 9192                                        | 4470                            |

#### Media MTU Sizes by Interface Type for MX Series Routers

**Table 279: Media MTU Sizes by Interface Type for MX Series Routers**

| Interface Type      | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes)            |
|---------------------|---------------------------|---------------------|--------------------------------------------|
| Gigabit Ethernet    | 1514                      | 9192                | 1500 (IPv4),<br>1488 (MPLS),<br>1497 (ISO) |
| 10-Gigabit Ethernet | 1514                      | 9192                | 1500 (IPv4),<br>1488 (MPLS),<br>1497 (ISO) |
| Multi-Rate Ethernet | 1514                      | 9192                | 1500 (IPv4),<br>1488 (MPLS),<br>1497 (ISO) |
| Tri-Rate Ethernet   | 1514                      | 9192                | 1500 (IPv4),<br>1488 (MPLS),<br>1497 (ISO) |

**Table 279: Media MTU Sizes by Interface Type for MX Series Routers (*continued*)**

|                                             |      |      |                                            |
|---------------------------------------------|------|------|--------------------------------------------|
| Channelized SONET/SDH OC3/STM1 (Multi-Rate) | 1514 | 9192 | 1500 (IPv4),<br>1488 (MPLS),<br>1497 (ISO) |
| DS3/E3 (Multi-Rate)                         | 1514 | 9192 | 1500 (IPv4),<br>1488 (MPLS),<br>1497 (ISO) |

**Media MTU Sizes by Interface Type for T320 Routers****Table 280: Media MTU Sizes by Interface Type for T320 Routers**

| Interface Type      | Default Media MTU (Bytes) | Maximum MTU (Bytes)                     | Default IP Protocol MTU (Bytes) |
|---------------------|---------------------------|-----------------------------------------|---------------------------------|
| ATM                 | 4482                      | 9192                                    | 4470                            |
| ATM2 IQ             | 4482                      | 9192                                    | 4470                            |
| Channelized OC12 IQ | 4474                      | 9192                                    | 4470                            |
| Channelized STM1 IQ | 4474                      | 9192                                    | 4470                            |
| DS3                 | 4471                      | 4500                                    | 4470                            |
| Fast Ethernet       | 1514                      | 1533 (4-port)<br>1532 (12- and 48-port) | 1500 (IPv4),<br>1497 (ISO)      |
| Gigabit Ethernet    | 1514                      | 9192                                    | 1500 (IPv4),<br>1497 (ISO)      |
| SONET/SDH           | 4474                      | 9192                                    | 4470                            |
| CT3 IQ              | 4474                      | 9192                                    | 4470                            |

**Media MTU Sizes by Interface Type for T640 Platforms****Table 281: Media MTU Sizes by Interface Type for T640 Platforms**

| Interface Type        | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes) |
|-----------------------|---------------------------|---------------------|---------------------------------|
| ATM2 IQ               | 4482                      | 9192                | 4470                            |
| 48-port Fast Ethernet | 1514                      | 1532                | 1500 (IPv4),<br>1497 (ISO)      |
| Gigabit Ethernet      | 1514                      | 9192                | 1500 (IPv4),<br>1497 (ISO)      |

**Table 281: Media MTU Sizes by Interface Type for T640 Platforms** *(continued)*

| Interface Type | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes) |
|----------------|---------------------------|---------------------|---------------------------------|
| SONET/SDH      | 4474                      | 9192                | 4470                            |
| CT3 IQ         | 4474                      | 9192                | 4470                            |

#### Media MTU Sizes by Interface Type for J2300 Platforms

**Table 282: Media MTU Sizes by Interface Type for J2300 Platforms**

| Interface Type         | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes) |
|------------------------|---------------------------|---------------------|---------------------------------|
| Fast Ethernet (10/100) | 1514                      | 9192                | 1500                            |
| G.SHDSL                | 4482                      | 9150                | 4470                            |
| ISDN BRI               | 1504                      | 4092                | 1500                            |
| Serial                 | 1504                      | 9150                | 1500                            |
| T1 or E1               | 1504                      | 9150                | 1500                            |

#### Media MTU Sizes by Interface Type for J4300 and J6300 Platforms

**Table 283: Media MTU Sizes by Interface Type for J4300 and J6300 Platforms**

| Interface Type                                        | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes) |
|-------------------------------------------------------|---------------------------|---------------------|---------------------------------|
| ADSL2+ PIM                                            | 4482                      | 9150                | 4470                            |
| Dual-port Fast Ethernet (10/100) PIM                  | 1514                      | 9192                | 1500                            |
| Dual-port Serial PIM                                  | 1504                      | 9150                | 1500                            |
| Dual-port T1 or E1 PIM                                | 1504                      | 9150                | 1500                            |
| Dual-port Channelized T1/E1 PIM (channelized to DS0s) | 1504                      | 4500                | 1500                            |

**Table 283: Media MTU Sizes by Interface Type for J4300 and J6300 Platforms (*continued*)**

| Interface Type                                           | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes) |
|----------------------------------------------------------|---------------------------|---------------------|---------------------------------|
| Dual-port Channelized T1/E1 PIM (clear channel T1 or E1) | 1504                      | 9150                | 1500                            |
| Fast Ethernet (10/100) built-in interface                | 1514                      | 9192                | 1500                            |
| G.SHDSL PIM                                              | 4482                      | 9150                | 4470                            |
| 4-port ISDN BRI PIM                                      | 1504                      | 4092                | 1500                            |
| T3 (DS3) or E3 PIM                                       | 4474                      | 9192                | 4470                            |

**Media MTU Sizes by Interface Type for J4350 and J6350 Platforms****Table 284: Media MTU Sizes by Interface Type for J4350 and J6350 Platforms**

| Interface Type                                           | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes) |
|----------------------------------------------------------|---------------------------|---------------------|---------------------------------|
| 4-port ISDN BRI PIM                                      | 1504                      | 4092                | 1500                            |
| ADSL2+ PIM                                               | 4482                      | 9150                | 4470                            |
| Dual-port Fast Ethernet (10/100) PIM                     | 1514                      | 9192                | 1500                            |
| Dual-port Serial PIM                                     | 1504                      | 9150                | 1500                            |
| Dual-port T1 or E1 PIM                                   | 1504                      | 9150                | 1500                            |
| Dual-port Channelized T1/E1 PIM (channelized to DS0s)    | 1504                      | 4500                | 1500                            |
| Dual-port Channelized T1/E1 PIM (clear channel T1 or E1) | 1504                      | 9150                | 1500                            |
| 4-port Fast Ethernet (10/100) ePIM                       | 1518                      | 1518                | 1500                            |



**Table 284: Media MTU Sizes by Interface Type for J4350 and J6350 Platforms** (*continued*)

| Interface Type                                                           | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes) |
|--------------------------------------------------------------------------|---------------------------|---------------------|---------------------------------|
| Gigabit Ethernet (10/100/1000) built-in interface                        | 1514                      | 9018                | 1500                            |
| Gigabit Ethernet (10/100/1000) Enhanced Physical Interface Module (ePIM) | 1514                      | 9018                | 1500                            |
| Gigabit Ethernet (10/100/1000) SFP ePIM                                  | 1514                      | 9018                | 1500                            |
| G.SHDSL PIM                                                              | 4482                      | 9150                | 4470                            |
| T3 (DS3) or E3 PIM                                                       | 4474                      | 9192                | 4470                            |



**NOTE:** On Gigabit Ethernet ePIMs in J4350 and J6350 Services Routers, you can configure a maximum transmission unit (MTU) size of only 9018 bytes even though the CLI indicates that you can configure an MTU of up to 9192 bytes. If you configure an MTU greater than 9018 bytes, the router does not accept the configuration and generates a system log error message similar to the following:

```
/kernel: ge-0/0/0: Illegal media change. MTU invalid: 9192. Max MTU supported on this PIC: 9018
```

On 4-port Fast Ethernet ePIMs in J4350 and J6350 Services Routers, you can configure a maximum transmission unit (MTU) size of only 1518 bytes even though the CLI indicates that you can configure an MTU of up to 9192 bytes. If you configure an MTU greater than 1518 bytes, the router does not accept the configuration and generates a system log error message similar to the following:

```
/kernel: fe-3/0/1: Illegal media change. MTU invalid: 9192. Max MTU supported on this PIC: 1518
```

## Media MTU Sizes by Interface Type for EX Series Switches and ACX Series Routers

Table 285: Media MTU Sizes by Interface Type for EX Series Switches and ACX Series Routers

| Interface Type      | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes) |
|---------------------|---------------------------|---------------------|---------------------------------|
| Gigabit Ethernet    | 1514                      | 9192                | 1500 (IPv4),<br>1497 (ISO)      |
| 10-Gigabit Ethernet | 1514                      | 9192                | 1500 (IPv4),<br>1497 (ISO)      |



**NOTE:** On ACX Series routers, you can configure the protocol MTU by including the `mtu` statement at the `[edit interfaces interface-name unit logical-unit-number family inet]` or `[edit interfaces interface-name unit logical-unit-number family inet6]` hierarchy level.

- If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.
- If you are configuring the protocol MTU for both `inet` and `inet6` families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for `inet` and `inet6` families that are configured on the same logical interface.

## Media MTU Sizes by Interface Type for PTX Series Packet Transport Routers

Table 286: Media MTU Sizes by Interface Type for PTX Series Packet Transport Routers

| Interface Type       | Default Media MTU (Bytes) | Maximum MTU (Bytes) | Default IP Protocol MTU (Bytes)            |
|----------------------|---------------------------|---------------------|--------------------------------------------|
| 10-Gigabit Ethernet  | 1514                      | 9500                | 1500 (IPv4),<br>1488 (MPLS),<br>1497 (ISO) |
| 40-Gigabit Ethernet  | 1514                      | 9500                | 1500 (IPv4),<br>1488 (MPLS),<br>1497 (ISO) |
| 100-Gigabit Ethernet | 1514                      | 9500                | 1500 (IPv4),<br>1488 (MPLS),<br>1497 (ISO) |

## Related Documentation

- [Configuring Interface Encapsulation on Physical Interfaces](#)
- [Setting the Protocol MTU on page 2651](#)

## Setting the Protocol MTU

When you initially configure an interface, the protocol maximum transmission unit (MTU) is calculated automatically. If you subsequently change the media MTU, the protocol MTU on existing address families automatically changes.

For a list of default protocol MTU values, see [“Configuring the Media MTU” on page 2638](#).

To modify the MTU for a particular protocol family, include the **mtu** statement:

**mtu** *bytes*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]**

If you increase the size of the protocol MTU, you must ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. For a list of encapsulation overhead values, see [Table 274 on page 2641](#). If you reduce the media MTU size, but there are already one or more address families configured and active on the interface, you must also reduce the protocol MTU size. (You configure the media MTU by including the **mtu** statement at the **[edit interfaces *interface-name*]** hierarchy level, as discussed in [“Configuring the Media MTU” on page 2638](#).)



**NOTE:** Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

The maximum number of data-link connection identifiers (DLCIs) is determined by the MTU on the interface. If you have keepalives enabled, the maximum number of DLCIs is 1000, with the MTU set to 5012.

The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the MTU. For example, the default protocol MTU for a Gigabit Ethernet interface is 1500 bytes, but the largest possible frame size is actually 1504 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

### Related Documentation

- [Configuring the Media MTU on page 2638](#)

## Interface Ranges



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Interface Ranges](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

Junos OS allows you to group a range of identical interfaces into an *interface range*. You first specify the group of identical interfaces in the interface range. Then you can apply a common configuration to the specified interface range, reducing the number of configuration statements required and saving time while producing a compact configuration.

- [Configuring Interface Ranges on page 2652](#)
- [Expanding Interface Range Member and Member Range Statements on page 2655](#)
- [Configuration Inheritance for Member Interfaces on page 2656](#)
- [Member Interfaces Inheriting Configuration from Configuration Groups on page 2657](#)
- [Interfaces Inheriting Common Configuration on page 2658](#)
- [Configuring Inheritance Range Priorities on page 2658](#)
- [Configuration Expansion Where Interface Range Is Used on page 2659](#)

---

### Configuring Interface Ranges

---

To configure an interface range, include the **interface-range** statement at the **[edit interfaces]** hierarchy level.

The **interface-range** statement accepts only physical networking interface names in its definition.

Interfaces can be grouped either as a range of interfaces or using a number range under the **interface-range** statement definition.

Interfaces in an **interface-range** definition can be added as part of a member range or as individual members or multiple members using a number range.

To specify a member range, use the **member-range** statement at the **[edit interfaces interface-range *name*]** hierarchy level.

To specify interfaces in lexical order, use the **member-range start-range to end-range** statement.

A range for a member statement must contain the following:

- **\***—All, specifies sequential interfaces from 0 through 47.



**CAUTION:** The wildcard **\*** in a member statement does not take into account the interface numbers supported by a specific interface type. Irrespective of the interface type, **\*** includes interface numbers ranging from 0 through 47 to the interface group. Therefore, use **\*** in a member statement with caution.

- **num**—Number; specifies one specific interface by its number.
- **[low-high]**—Numbers between low to high; specifies a range of sequential interfaces.
- **[num1, num2, num3]**—Numbers **num1**, **num2**, and **num3** specify multiple specific interfaces.

**Example: Specifying an  
Interface Range  
Member Range**

```
member-range ge-0/0/0 to ge-4/0/40;
```

To specify one or multiple members, use the **member** statement at the **[edit interfaces interface-range *name*]** hierarchy level.

To specify the list of interface range members individually or for multiple interfaces using regex, use the **member *list of interface names*** statement.

**Example: Specifying an  
Interface Range  
Member**

```
member ge-0/0/0;  
member ge-0/*/*  
member ge-0/[1-10]/0;  
member ge-0/[1,2,3]/3;
```

Regex or wildcards are not supported for interface-type prefixes. For example, prefixes **ge**, **fe**, and **xe** must be mentioned explicitly.

An **interface-range** definition can contain both **member** and **member-range** statements within it. There is no maximum limit on the number of **member** or **member-range** statements within an interface-range. However, at least one **member** or **member-range** statement must exist within an **interface-range** definition.

**Example: Interface  
Range Common  
Configuration**

Configuration common to an interface range can be added as a part of the **interface-range** definition, as follows:

```
[edit]  
interfaces {  
  + interface-range foo {  
    + member-range ge-1/0/0 to ge-4/0/40;  
    + member ge-0/1/1;  
    + member ge-5/[1-10]/*;  
    /*Common configuration is added as part of interface-range definition*/  
    mtu 256;  
    hold-time up 10;  
    ether-options {  
      flow-control;  
      speed {  
        100m;  
      }  
      802.3ad primary;  
    }  
  }  
}
```

An **interface-range** definition having just **member** or **member-range** statements and no common configurations statements is valid.

These defined interface ranges can be used in other configuration hierarchies, in places where an **interface** node exists.

**Example:  
Interface-Range foo  
Used Under the  
Protocols Hierarchy**

```
protocols {  
  dot1x {  
    authenticator {  
      interface foo{  
        retries 1;
```

```
    }  
  }  
}
```

**foo** should be an **interface-range** defined at the **[interfaces]** hierarchy level. In the above example, the **interface** node can accept both individual interfaces and interface ranges.



**TIP:** To view an interface range in expanded configuration, use the **(show | display inheritance)** command. For more information, see the *CLI User Guide*.

The defined interface ranges can be used at places where the **interface** node is used in the following configuration hierarchies:

- forwarding-options analyzer *name* input egress interface
- forwarding-options analyzer *name* input ingress interface
- poe interface
- protocols dot1x authenticator interface
- protocols igmp interface
- protocols isis interface
- protocols layer2-control bpdu-block interface
- protocols link-management peer *name* lmp-control-channel
- protocols link-management te-link *name* interface
- protocols lldp interface
- protocols lldp-med interface
- protocols mstp interface
- protocols oam ethernet link-fault-management interface
- protocols ospf area *area-id* interface
- protocols pim interface
- protocols router-advertisement interface
- protocols router-discovery interface
- protocols rsvp interface
- protocols sflow interfaces
- protocols vstp vlan *vlan-id* interface
- switch-options redundant-trunk-group group-name interface
- switch-options voip interface

### Expanding Interface Range Member and Member Range Statements

All **member** and **member-range** statements in an interface range definition are expanded to generate the final list of interface names for the specified interface range.

#### Example: Expanding Interface Range Member and Member Range Statements

```
[edit]
interfaces {
  interface-range range-1 {
    member-range ge-0/0/0 to ge-4/0/20;
    member ge-10/1/1;
    member ge-5/[0-5]/*;
    /*Common configuration is added part of the interface-range definition*/
    mtu 256;
    hold-time up 10;
    ether-options {
      flow-control;
      speed {
        100m;
      }
      802.3ad primary;
    }
  }
}
```

For the **member-range** statement, all possible interfaces between **start-range** and **end-range** are considered in expanding the members. For example, the following **member-range** statement:

**member-range ge-0/0/0 to ge-4/0/20**

expands to:

```
[ge-0/0/0, ge-0/0/1 ... ge-0/0/max_ports
ge-0/1/0 ge-0/1/1 ... ge-0/1/max_ports
ge-0/2/0 ge-0/2/1 ... ge-0/2/max_ports
.
.
ge-0/MAX_PICS/0 ... ge-0/max_pics/max_ports
ge-1/0/0 ge-1/0/1 ... ge-1/0/max_ports
.
ge-1/MAX_PICS/0 ... ge-1/max_pics/max_ports
.
.
ge-4/0/0 ge-4/0/1 ... ge-4/0/max_ports]
```

The following **member** statement:

**ge-5/[0-5]/\***

expands to:

```
ge-5/0/0 ... ge-5/0/max_ports
ge-5/1/0 ... ge-5/0/max_ports
.
.
ge-5/5/0 ... ge-5/5/max_ports
```

The following **member** statement:

```
ge-5/1/[2,3,6,10]
```

expands to:

```
ge-5/1/2
ge-5/1/3
ge-5/1/6
ge-5/1/10
```

### Configuration Inheritance for Member Interfaces

---

When the Junos OS expands the **member** and **member-range** statements present in an **interface-range**, it creates *interface objects* if they are not explicitly defined in the configuration. The common configuration is copied to all its member interfaces in the **interface-range**.

**Example:** Foreground interface configuration takes priority compared to configuration inherited by the interface through the **interface-range**.

#### Configuration Priorities

```
interfaces {
  interface-range range-1 {
    member-range ge-1/0/0/ to ge-10/0/47;
    mtu 256;
  }
  ge-1/0/1 {
    mtu 1024;
  }
}
```

In the preceding example, interface **ge-1/0/1** will have an MTU value of 1024.

This can be verified with output of the **show interfaces | display inheritance** command, as follows:

```
user@host: # show interfaces | display inheritance
## 'ge-1/0/0' was expanded from interface-range 'range-1'
##
ge-1/0/0 {
  ##
  ## '256' was expanded from interface-range 'range-1'
  ##
  mtu 256;
}
ge-1/0/1 {
  mtu 1024;
}
##
## 'ge-1/0/2' was expanded from interface-range 'range-1'
##
ge-1/0/2 {
  ##
  ## '256' was expanded from interface-range 'range-1'
  ##
  mtu 256;
}
.....
.....
##
## 'ge-10/0/47' was expanded from interface-range 'range-1'
##
```



```

ge-10/0/47 {
  ##
  ## '256' was expanded from interface-range 'range-1'
  ##
  mtu 256;
}

```

### Member Interfaces Inheriting Configuration from Configuration Groups

Interface range member interfaces inherit the config-groups configuration like any other foreground configuration. **interface-range** is similar to any other foreground configuration statement. The only difference is that the **interface-range** goes through a member interfaces expansion before Junos OS reads this configuration.

```

groups {
  global {
    interfaces {
      <*> {
        hold-time up 10;
      }
    }
  }
  apply-groups [global];
  interfaces {
    interface-range range-1 {
      member-range ge-1/0/0 to ge-10/0/47;
      mtu 256;
    }
  }
}

```

The **hold-time** configuration is applied to all members of **interface-range range-1**.

This can be verified with **show interfaces | display inheritance** as follows:

```

user@host# show interfaces | display inheritance
ge-1/0/0 {
  ##
  ## '256' was expanded from interface-range 'range-1'
  ##
  mtu 256;
  ##
  ## 'hold-time' was inherited from group 'global'
  ## '10' was inherited from group 'global'
  ##
  hold-time up 10;
}
ge-1/0/1 {
  ##
  ## '256' was expanded from interface-range 'range-1'
  ##
  mtu 256;
  ##
  ## 'hold-time' was inherited from group 'global'
  ## '10' was inherited from group 'global'
  ##
  hold-time up 10;
}
ge-10/0/47 {

```

```
##
## '256' was expanded from interface-range 'range-1'
##
mtu 256;
##
## 'hold-time' was inherited from group 'global'
## '10' was inherited from group 'global'
##
hold-time up 10;
}
```

---

### Interfaces Inheriting Common Configuration

If an interface is a member of several interface ranges, that interface will inherit the common configuration from all of those interface ranges.

```
[edit]
interfaces {
  interface-range range-1 {
    member-range ge-1/0/0 to ge-10/0/47;
    mtu 256;
  }
}
interfaces {
  interface-range range-1 {
    member-range ge-10/0/0 to ge-10/0/47;
    hold-time up 10;
  }
}
```

In this example, interfaces **ge-10/0/0** through **ge-10/0/47** will have both **hold-time** and **mtu**.

---

### Configuring Inheritance Range Priorities

The interface ranges are defined in the order of inheritance priority, with the first interface range configuration data taking priority over subsequent interface ranges.

```
[edit]
interfaces {
  interface-range int-grp-one {
    member-range ge-0/0/0 to ge-4/0/40;
    member ge-1/1/1;
    /*Common config is added part of the interface-range definition*/
    mtu 256;
    hold-time up 10;
  }
}
interfaces {
  interface-range int-grp-two {
    member-range ge-5/0/0 to ge-10/0/40;
    member ge-1/1/1;
    mtu 1024;
  }
}
```

Interface **ge-1/1/1** exists in both **interface-range int-grp-one** and **interface-range int-grp-two**. This interface inherits **mtu 256** from **interface-range int-grp-one** because it was defined first.

### Configuration Expansion Where Interface Range Is Used

In this example, **interface-range range-1** is used under the **protocols** hierarchy:

```
[edit]
interfaces {
  interface-range range-1 {
    member ge-10/1/1;
    member ge-5/5/1;
    mtu 256;
    hold-time up 10;
    ether-options {
      flow-control;
      speed {
        100m;
      }
      802.3ad primary;
    }
  }
}
protocols {
  dot1x {
    authenticator {
      interface range-1 {
        retries 1;
      }
    }
  }
}
```

The **interface** node present under **authenticator** is expanded into member interfaces of the **interface-range range-1** as follows:

```
protocols {
  dot1x {
    authenticator {
      interface ge-10/1/1 {
        retries 1;
      }
      interface ge-5/5/1 {
        retries 1;
      }
    }
  }
}
```

The **interface range-1** statement is expanded into two interfaces, **ge-10/1/1** and **ge-5/5/1**, and configuration **retries 1** is copied under those two interfaces.

This configuration can be verified using the **show protocols dot1x | display inheritance** command.

## Configuring Accounting for the Physical Interface

Juniper Networks routers and switches can collect various kinds of data about traffic passing through the router and switch. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the router or switch retains before discarding, and the number of bytes per file
- The polling period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the **[edit accounting-options]** hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the **interface-profile** statement at the **[edit accounting-options]** hierarchy level. You configure filter profiles by including the **filter-profile** statement at the **[edit accounting-options]** hierarchy level. For more information, see the *Network Management Administration Guide for Routing Devices*.

You apply filter profiles by including the **accounting-profile** statement at the **[edit firewall filter *filter-name*]** and **[edit firewall family *family* filter *filter-name*]** hierarchy levels. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

### Applying an Accounting Profile to the Physical Interface

---

To enable accounting on an interface, include the **accounting-profile** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
  accounting-profile name;
```

You can also reference profiles by logical unit; for more information, see “[Configuring Accounting for the Logical Interface](#)” on page 2661.

#### *Example: Applying an Accounting Profile to the Physical Interface*

Configure an accounting profile for an interface and apply it to a physical interface:

```
[edit]  
accounting-options {  
  file if_stats {  
    size 4m files 10 transfer-interval 15;  
    archive-sites {  
      "ftp://login:password@host/path";  
    }  
  }  
  interface-profile if_profile {  
    interval 15;  
    file if_stats {  
      fields {  
        input-bytes;  
        output-bytes;  
      }  
    }  
  }  
}
```

```

        input-packets;
        output-packets;
        input-errors;
        output-errors;
    }
}
}
[edit interfaces ge-1/0/1]
accounting-profile if_profile;

```

## Configuring Accounting for the Logical Interface

Juniper Networks routers or switches can collect various kinds of data about traffic passing through the router or switch. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the router or switch retains before discarding, and the number of bytes per file
- The period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the **[edit accounting-options]** hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the **interface-profile** statement at the **[edit accounting-options]** hierarchy level. You configure filter profiles by including the **filter-profile** statement at the **[edit accounting-options]** hierarchy level. For more information, see the *Network Management Administration Guide for Routing Devices*.

You apply filter profiles by including the **accounting-profile** statement at the **[edit firewall filter *filter-name*]** and **[edit firewall family *family* filter *filter-name*]** hierarchy levels. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

### Applying an Accounting Profile to the Logical Interface

To enable accounting on a logical interface, include the **accounting-profile** statement:

```
accounting-profile name;
```

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**

You can also reference profiles for the physical interface; for more information, see [“Configuring Accounting for the Physical Interface” on page 2660](#).

#### *Example: Applying an Accounting Profile to the Logical Interface*

Configure an accounting profile for an interface and apply it to a logical interface:

```

[edit]
accounting-options {

```

```
file if_stats {
  size 4m files 10 transfer-interval 15;
  archive-sites {
    "ftp://login:password@host/path";
  }
}
interface-profile if_profile {
  interval 15;
  file if_stats {
    fields {
      input-bytes;
      output-bytes;
      input-packets;
      output-packets;
      input-errors;
      output-errors;
    }
  }
}
[edit interfaces ge-1/0/1 unit 1]
accounting-profile if_profile;
```

To reference profiles by physical interface, see [“Applying an Accounting Profile to the Physical Interface” on page 2660](#). For information about configuring a firewall filter accounting profile, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

## Configuring Ethernet Loopback Capability

By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system. To place an interface in loopback mode, include the **loopback** statement:

**loopback;**



**NOTE:** If you configure a local loopback on a 1-port 10-Gigabit IQ2 and IQ2-E PIC using the loopback statement at the [edit interfaces *interface-name* *gigether-options*] hierarchy level, the transmit-path stops working, causing the remote end to detect a link down.

To return to the default—that is, to disable loopback mode—delete the **loopback** statement from the configuration:

```
[edit]
user@host# delete interfaces fe-fpc/pic/port fastether-options loopback
```

To explicitly disable loopback mode, include the **no-loopback** statement:

**no-loopback;**

You can include the **loopback** and **no-loopback** statements at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* ether-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]

#### Related Documentation

- [loopback on page 2803](#)
- *Ethernet Interfaces Overview*
- [EX Series Switches Interfaces Overview on page 2577](#)
- *Ethernet Interfaces Feature Guide for Routing Devices*

## Configuring Gratuitous ARP

Gratuitous Address Resolution Protocol (ARP) requests provide duplicate IP address detection. A gratuitous ARP request is a broadcast request for a router's own IP address. If a router or switch sends an ARP request for its own IP address and no ARP replies are received, the router- or switch-assigned IP address is not being used by other nodes. If a router or switch sends an ARP request for its own IP address and an ARP reply is received, the router- or switch-assigned IP address is already being used by another node.

By default, the router or switch responds to gratuitous ARP requests. On Ethernet interfaces, you can disable responses to gratuitous ARP requests. To disable responses to gratuitous ARP requests, include the **no-gratuitous-arp-request** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
no-gratuitous-arp-request;
```

To return to the default—that is, to respond to gratuitous ARP requests—delete the **no-gratuitous-arp-request** statement from the configuration:

```
[edit]
user@host# delete interfaces interface-name no-gratuitous-arp-request
```

Gratuitous ARP replies are reply packets sent to the broadcast MAC address with the target IP address set to be the same as the sender's IP address. When the router or switch receives a gratuitous ARP reply, the router or switch can insert an entry for that reply in the ARP cache.

By default, updating the ARP cache on gratuitous ARP replies is disabled on the router or switch. On Ethernet interfaces, you can enable handling of gratuitous ARP replies on a specific interface by including the **gratuitous-arp-reply** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
gratuitous-arp-reply;
```

To restore the default behavior, include the **no-gratuitous-arp-reply** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
no-gratuitous-arp-reply;
```

**Related  
Documentation**

- [gratuitous-arp-reply on page 2776](#)
- [no-gratuitous-arp-request on page 2817](#)
- [Ethernet Interfaces Overview](#)
- [EX Series Switches Interfaces Overview on page 2577](#)
- [Ethernet Interfaces Feature Guide for Routing Devices](#)

## Configuring Static ARP Table Entries

To configure static ARP table entries, include the **arp** statement:

```
arp ip-address (mac | multicast-mac) mac-address <publish>;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**

The IP address that you specify must be part of the subnet defined in the enclosing **address** statement.

To associate a multicast MAC address with a unicast IP address, include the **multicast-mac** statement.

Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn*; for example, 0011.2233.4455 or 00:11:22:33:44:55.

For unicast MAC addresses only, if you include the **publish** option, the router or switch replies to proxy ARP requests.



**NOTE:** By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. By including the **arp** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet policer]** hierarchy level, you can apply a specific ARP-packet policer to an interface. This feature is not available on EX Series switches.

When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the **unnumbered-address** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level.

---





**NOTE:** The Junos OS supports the IPv6 static neighbor discovery cache entries, similar to the static ARP entries in IPv4.

### Example: Configuring Static ARP Table Entries

Configure two static ARP table entries on the router or switch's management interface:

```
[edit interfaces]
fxp0 {
  unit 0 {
    family inet {
      address 10.10.0.11/24 {
        arp 10.10.0.99 mac 0001.0002.0003;
        arp 10.10.0.101 mac 00:11:22:33:44:55 publish;
      }
    }
  }
}
```

#### Related Documentation

- *Management Ethernet Interface Overview*
- [EX Series Switches Interfaces Overview on page 2577](#)
- *Applying Policers*
- *Configuring an Unnumbered Interface*
- *Ethernet Interfaces Feature Guide for Routing Devices*

### Disabling the Transmission of Redirect Messages on an Interface

By default, the interface sends protocol redirect messages. To disable the sending of these messages on an interface, include the **no-redirects** statement:

```
no-redirects;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family family]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family]`

To disable the sending of protocol redirect messages for the entire router or switch, include the **no-redirects** statement at the `[edit system]` hierarchy level.

#### Related Documentation

- [no-redirects on page 2817](#)

### Configuring Restricted and Unrestricted Proxy ARP

To configure restricted or unrestricted proxy ARP, include the **proxy-arp** statement:

```
proxy-arp (restricted |unrestricted);
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* ]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

To return to the default—that is, to disable restricted or unrestricted proxy ARP—delete the **proxy-arp** statement from the configuration:

```
[edit]
user@host# delete interfaces interface-name unit logical-unit-number proxy-arp
```

You can track the number of restricted or unrestricted proxy ARP requests processed by the router or switch by issuing the **show system statistics arp** operational mode command.



**NOTE:** When proxy ARP is enabled as default or unrestricted, the router or switch responds to any ARP request as long as the device has an active route to the target address of the ARP request. This gratuitous ARP behavior can result in an error when the receiving interface and target response interface are the same and the end device (for example, a client) performs a duplicate address check. To prevent this error, configure the router or switch interface with the **no-gratuitous-arp-reply** statement. See [“Configuring Gratuitous ARP” on page 2663](#) for information about how to disable responses to gratuitous ARP requests.

#### Related Documentation

- [proxy-arp on page 2422](#)
- *Restricted and Unrestricted Proxy ARP Overview*
- [Configuring Gratuitous ARP on page 2663](#)
- *Ethernet Interfaces Feature Guide for Routing Devices*

## Enabling or Disabling SNMP Notifications on Logical Interfaces

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. To explicitly enable these notifications on the logical interface, include the **traps** statement; to disable these notifications on the logical interface, include the **no-traps** statement:

```
(traps | no-traps);
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]



**NOTE:** Gigabit Ethernet interfaces on J Series routers do not support SNMP.

## Configuring Aggregated Ethernet Links (CLI Procedure)

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



**NOTE:** An interface with an already configured IP address cannot form part of the aggregation group.

To configure aggregated Ethernet interfaces, using the CLI:

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

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

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



**NOTE:** By default, only one link must be up for the bundle to be labeled *up*.

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options minimum-links number
```

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

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options link-speed speed
```

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

```
[edit interfaces]
user@switch# set xe-fpc/pic/port ether-options 802.3ad ae0
user@switch# set xe-fpc/pic/port ether-options 802.3ad ae0
```

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

```
[edit interfaces]
user@switch# set ae0 unit 0 family inet address address
```

For information about adding LACP to a LAG, see “Configuring Aggregated Ethernet LACP (CLI Procedure)” on page 2671.

### Related Documentation

- [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\)](#) on page 2668
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\)](#) on page 2671
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\)](#) on page 2672
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch](#)

- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- [Verifying the Status of a LAG Interface on page 2837](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2582](#)

## Configuring Aggregated Ethernet Interfaces (J-Web Procedure)

---



**NOTE:** This topic applies only to the J-Web Application package.

---

Use the link aggregation feature to aggregate one or more Ethernet interfaces to form a virtual link or link aggregation group (LAG) on an EX Series switch. The MAC client can treat this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability. You can use the J-Web interface to configure aggregated Ethernet interfaces, or a LAG, on the switch.

---



**NOTE:** Interfaces that are already configured with MTU, duplex, flow control, or logical interfaces are listed but are not available for aggregation.

---

To configure an aggregated Ethernet interface (also referred to as a LAG):

1. Select **Configure > Interfaces > Link Aggregation**.

The list of aggregated interfaces is displayed.

---



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

---

2. Click one of the following:

- **Add**—Creates an aggregated Ethernet interface, or LAG. Enter information as specified in [Table 287 on page 2669](#).
- **Edit**—Modifies a selected LAG.
  - **Aggregation**—Modifies settings for the selected LAG. Enter information as specified in [Table 287 on page 2669](#).
  - **VLAN**—Specifies VLAN options for the selected LAG. Enter information as specified in [Table 288 on page 2670](#).
  - **IP Option**—Specifies IP options for the selected LAG. Enter information as specified in [Table 289 on page 2670](#).
- **Delete**—Deletes the selected LAG.

- **Disable Port** or **Enable Port**—Disables or enables the administrative status on the selected interface.
- **Device Count**—Configures the number of aggregated logical devices available to the switch. Select the number and click **OK**.

Table 287: Aggregated Ethernet Interface Options

| Field                | Function                                                                                                                                                                                                                                                                                                                                                                                      | Your Action                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aggregated Interface | Specifies the name of the aggregated interface.                                                                                                                                                                                                                                                                                                                                               | None. The name is supplied by the software.                                                                                                                                                                                                                                                                                                                                                                       |
| LACP Mode            | <p>Specifies the mode in which LACP packets are exchanged between the interfaces. The modes are:</p> <ul style="list-style-type: none"> <li>• <b>None</b>—Indicates that no mode is applicable.</li> <li>• <b>Active</b>—Indicates that the interface initiates transmission of LACP packets</li> <li>• <b>Passive</b>—Indicates that the interface responds only to LACP packets.</li> </ul> | Select from the list.                                                                                                                                                                                                                                                                                                                                                                                             |
| Description          | Specifies a description for the LAG.                                                                                                                                                                                                                                                                                                                                                          | Enter a description.                                                                                                                                                                                                                                                                                                                                                                                              |
| Interface            | Specifies the interfaces in the LAG.                                                                                                                                                                                                                                                                                                                                                          | <p>To add interfaces to the LAG, select the interfaces and click <b>Add</b>. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list. Click <b>OK</b>.</p> <p>To remove an interface from the LAG, select the interface and click <b>Remove</b>.</p> <p><b>NOTE:</b> Only interfaces that are configured with the same speed can be selected together for a LAG.</p> |
| Enable Log           | Specifies whether to enable generation of log entries for the LAG.                                                                                                                                                                                                                                                                                                                            | Select the check box to enable log generation, or clear the check box to disable log generation.                                                                                                                                                                                                                                                                                                                  |

Table 288: VLAN Options

| Field     | Function                                                       | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port Mode | Specifies the mode of operation for the port: trunk or access. | <p>If you select <b>Trunk</b>, you can:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b> to add a VLAN member.</li> <li>2. Select the VLAN and click <b>OK</b>.</li> <li>3. (Optional) Associate a native VLAN ID with the port.</li> </ol> <p>If you select <b>Access</b>, you can:</p> <ol style="list-style-type: none"> <li>1. Select the VLAN member to be associated with the port.</li> <li>2. (Optional) Associate a VoIP VLAN with the interface. Only a VLAN with a VLAN ID can be associated as a VoIP VLAN.</li> </ol> <p>Click <b>OK</b>.</p> |

Table 289: IP Options

| Field        | Function                                        | Your Action                                                                                                                                                                                                                                                                                            |
|--------------|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPv4 Address | Specifies an IPv4 address for the selected LAG. | <ol style="list-style-type: none"> <li>1. Select the check box <b>IPv4 address</b>.</li> <li>2. Type an IP address—for example, <b>10.10.10.10</b>.</li> <li>3. Enter the subnet mask or address prefix. For example, 24 bits represents <b>255.255.255.0</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> |
| IPv6 Address | Specifies an IPv6 address for the selected LAG. | <ol style="list-style-type: none"> <li>1. Select the check box <b>IPv6 address</b>.</li> <li>2. Type an IP address—for example, <b>2001:ab8:85a3::8a2e:370:7334</b>.</li> <li>3. Enter the subnet mask or address prefix.</li> <li>4. Click <b>OK</b>.</li> </ol>                                      |

**Related Documentation**

- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch](#)
- [Verifying the Status of a LAG Interface on page 2837](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 2671](#)

- [Understanding Aggregated Ethernet Interfaces and LACP on page 2582](#)

## Configuring Aggregated Ethernet LACP (CLI Procedure)

For aggregated Ethernet interfaces on EX Series switches, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure aggregated Ethernet interfaces with or without LACP enabled.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the bundle without user intervention
- Link monitoring to check whether both ends of the bundle are connected to the correct group



**NOTE:** You can also configure LACP link protection on aggregated Ethernet interfaces. For information, see [“Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\)” on page 2672](#).

The Junos OS implementation of LACP provides link monitoring but not automatic addition and deletion of links.

Before you configure LACP, be sure you have:

- Configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs). See [“Configuring Aggregated Ethernet Links \(CLI Procedure\)” on page 2667](#)

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



**NOTE:** Do not add LACP to a LAG if the remote end of the LAG link is a security device, unless the security device supports LACP. Security devices often do not support LACP because they require a deterministic configuration.

To configure LACP:

1. Configure at least one side of the aggregated Ethernet link as active:

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options lacp active
```

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

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



**NOTE:** The LACP process exists in the system only if you configure the system in either active or passive LACP mode.

**Related Documentation**

- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\) on page 2672](#)
- [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\) on page 2668](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch](#)
- [Verifying the Status of a LAG Interface on page 2837](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2582](#)

## Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)

You can configure LACP link protection and system priority at the global level on the switch or for a specific aggregated Ethernet interface. When using LACP link protection to protect a single link in the aggregated ethernet bundle, you configure only two member links for an aggregated Ethernet interface: one active and one standby. LACP link protection ensures that only one link—the link with the higher priority—is used for traffic. The other link is forced to stay in a *waiting* state.

When using LACP link protection to protect multiple links in an aggregated ethernet bundle, you configure links into primary and backup subgroups. A link protection subgroup is a collection of ethernet links within the aggregated ethernet bundle. When you use link protection subgroups, you configure a primary subgroup and a backup subgroup. The configuration process includes assigning member links to each subgroup. When the configuration process is complete, the primary subgroup is used to forward traffic until a switchover event, such as a link failure, occurs and causes the backup subgroup to assume control of traffic that was travelling on the links in the primary subgroup within the bundle.

By default LACP link protection reverts to a higher-priority (lower-numbered) link when the higher-priority link becomes operational or when a higher-priority link is added to the aggregated Ethernet bundle. For priority purposes, LACP link protection treats subgroups like links. You can suppress link calculation by adding the **non-revertive** statement to the link protection configuration. In nonrevertive mode, when a link is active in sending and receiving LACP packets, adding a higher-priority link to the bundle does not change the status of the currently active link. It remains active.

If LACP link configuration is specified to be nonrevertive at the global **[edit chassis]** hierarchy level, you can specify the **revertive** statement in the LACP link protection configuration at the aggregated Ethernet interface level to override the nonrevertive setting for the interface. In revertive mode, adding a higher-priority link to the aggregated



Ethernet bundle results in LACP recalculating the priority and switching the status from the currently active link to the newly added, higher-priority link.



**NOTE:** When LACP link protection is enabled on both local and remote sides of the link, both sides must use the same mode (either revertive or nonrevertive).

Configuring LACP link configuration at the aggregated Ethernet level results in only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

Before you configure LACP link protection, be sure you have:

- Configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs). See [“Configuring Aggregated Ethernet Links \(CLI Procedure\)” on page 2667](#).
- Configured LACP for the interface. See [“Configuring Aggregated Ethernet LACP \(CLI Procedure\)” on page 2671](#).

You can configure LACP link protection for all aggregated Ethernet interfaces on the switch by enabling it at the global level on the switch or configure it for a specific aggregated Ethernet interface by enabling it on that interface.

- [Configuring LACP Link Protection for a Single Link at the Global Level on page 2674](#)
- [Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level on page 2674](#)
- [Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface on page 2675](#)

### Configuring LACP Link Protection for a Single Link at the Global Level

---

To configure LACP link protection for aggregated Ethernet interfaces at the global level:

1. Enable LACP link protection on the switch:

```
[edit chassis aggregated-devices ethernet lacp]
user@switch# set link-protection
```

2. (Optional) Configure the LACP link protection for the aggregated Ethernet interfaces to be in nonrevertive mode:



**NOTE:** LACP link protection is in revertive mode by default.

```
[edit chassis aggregated-devices ethernet lacp link-protection]
user@switch# set non-revertive
```

3. (Optional) To configure LACP system priority for the aggregated Ethernet interfaces:

```
[edit chassis aggregated-devices ethernet lacp]
user@switch# set system-priority
```

### Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level

---

To enable LACP link protection for a specific aggregated Ethernet interface:

1. Enable LACP link protection for the interface:

```
[edit interfaces aeX aggregated-ether-options lacp]
user@switch# set link-protection
```

2. (Optional) Configure the LACP link protection for the aggregated Ethernet interface to be in revertive or nonrevertive mode:

- To specify revertive mode:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
user@switch# set revertive
```

- To specify nonrevertive mode:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
user@switch# set non-revertive
```

3. (Optional) To configure LACP system priority for an aggregated Ethernet interface:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
user@switch# set system-priority
```

4. (Optional) To configure LACP port priority for an aggregated Ethernet interface:

```
[edit interfaces ge-fpc/pic/port ether-options 802.3ad lacp]
user@switch# set port-priority
```

## Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface

You can configure link protection subgroup bundles to provide link protection for multiple links in an aggregated ethernet bundle.

Link protection subgroups allow you to provide link protection to a collection of Ethernet links within a LAG bundle, instead of providing protection to a single link in the aggregated ethernet bundle only. You can, for instance, configure a primary subgroup with three member links and a backup subgroup with three different member links and use the backup subgroup to provide link protection for the primary subgroup.

To configure link protection using subgroups:

1. Configure the primary link protection subgroup in the aggregated ethernet interface:

```
[edit interfaces aeX aggregated-ether-options]
user@switch# set link-protection-sub-group group-name primary
```

For instance, to create a primary link protection subgroup named **subgroup-primary** for interface **ae0**:

```
[edit interfaces ae0 aggregated-ether-options]
user@switch# set link-protection-sub-group subgroup-primary primary
```

2. Configure the backup link protection subgroup in the aggregated ethernet interface:

```
[edit interfaces aeX aggregated-ether-options]
user@switch# set link-protection-sub-group group-name backup
```

For instance, to create a backup link protection subgroup named **subgroup-backup** for interface **ae0**:

```
[edit interfaces ae0 aggregated-ether-options]
user@switch# set link-protection-sub-group subgroup-backup backup
```



**NOTE:** You can create one primary and one backup link protection subgroup per aggregated ethernet interface.

3. Attach interfaces to the link protection subgroups:

```
[edit interfaces interface-name ether-options 802.3ad]
user@switch# set link-protection-sub-group group-name
```



**NOTE:** The primary and backup link protection subgroups must contain the same number of interfaces. For instance, if the primary link protection subgroup contains three interfaces, the backup link protection subgroup must also contain three interfaces.

For instance, to configure interfaces **ge-0/0/0** and **ge-0/0/1** into link protection subgroup **subgroup-primary** and interfaces **ge-0/0/2** and **ge-0/0/3** into link protection subgroup **subgroup-backup**:

```
[edit interfaces ge-0/0/0 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-primary
[edit interfaces ge-0/0/1 ether-options 802.3ad]
```

```
user@switch# set link-protection-sub-group subgroup-primary
[edit interfaces ge-0/0/2 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-backup
[edit interfaces ge-0/0/3 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-backup
```

4. (Optional) Configure the port priority for link protection:

```
[edit interfaces interface-name ether-options 802.3ad]
user@switch# set port-priority priority
```

The port priority is used to select the active link.

5. Enable link protection

To enable link protection at the LAG level:

```
[edit interfaces aeX aggregated-ether-options]
user@switch# set link-protection
```

To enable link protection at the LACP level:

```
[edit interfaces aeX aggregated-ether-options lacp]
user@switch# set link-protection
```

For instance, to enable link protection on **ae0** at the LAG level:

```
[edit interfaces ae0 aggregated-ether-options]
user@switch# set link-protection
```

For instance, to enable link protection on **ae0** at the LACP level:

```
[edit interfaces ae0 aggregated-ether-options lacp]
user@switch# set link-protection
```

**Related  
Documentation**

- [Understanding Aggregated Ethernet Interfaces and LACP on page 2582](#)

## Configuring Aggregated Ethernet Link Protection

You can configure link protection for aggregated Ethernet interfaces to provide QoS on the links during operation.

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.



**NOTE:** Link protection is not supported on MX80.

---

- [Configuring Link Protection for Aggregated Ethernet Interfaces on page 2677](#)
- [Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces on page 2677](#)

- [Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link on page 2677](#)
- [Disabling Link Protection for Aggregated Ethernet Interfaces on page 2677](#)

### Configuring Link Protection for Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces support link protection to ensure QoS on the interface.

To configure link protection:

1. Specify that you want to configure the options for an aggregated Ethernet interface.

```
user@host# edit interfaces aex aggregated-ether-options
```

2. Configure the link protection mode.

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

### Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces

To configure link protection, you must specify a primary and a secondary, or backup, link.

To configure a primary link and a backup link:

1. Configure the primary logical interface.

```
[edit interfaces interface-name]
user@host# set (fastether-options | gigether-options) 802.3ad aex primary
```

2. Configure the backup logical interface.

```
[edit interfaces interface-name]
user@host# set (fastether-options | gigether-options) 802.3ad aex backup
```

### Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.

To manually control when traffic should be diverted back to the primary link from the designated backup link, enter the following operational command:

```
user@host> request interface revert aex
```

### Disabling Link Protection for Aggregated Ethernet Interfaces

To disable link protection, issue the **delete interface revert aex** configuration command.

```
user@host# delete interfaces aex aggregated-ether-options link-protection
```

## Configuring Aggregated Ethernet Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. Generally, all interfaces that make up a bundle must have the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed that you specify in the **link-speed** parameter, an error message is logged. However, starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)



---

### NOTE:

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
  - Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.
- 

To configure member links of mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers, you need to configure the **mixed** option for the **[edit interfaces aex aggregated-ether-options link-speed]** statement.

To set the required link speed:

1. Specify that you want to configure the aggregated Ethernet options.

```
user@host# edit interfaces interface-name aggregated-ether-options
```

2. Configure the link speed.

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

**speed** can be in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet interfaces on the M120 router can have one of the following speeds:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **oc192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**—Links are 10 Mbps.
- **100m**—Links are 100 Mbps.
- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.
- **50g**—Links are 50 Gbps.

Aggregated Ethernet links on T Series routers can be configured to operate at one of the following speeds:

- **100g**—Links are 100 Gbps.
- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **40g**—Links are 40 Gbps.
- **50g**—Links are 50 Gbps.
- **80g**—Links are 80 Gbps.
- **8g**—Links are 8 Gbps.
- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

**Related  
Documentation**

- *aggregated-ether-options*
- *Configuring Mixed Aggregated Ethernet Links*
- *Ethernet Interfaces Feature Guide for Routing Devices*

## Configuring Aggregated Ethernet Minimum Links

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled **up**. By default, only one link must be up for the bundle to be labeled **up**.

To configure the minimum number of links:

1. Specify that you want to configure the aggregated Ethernet options.

```
user@host# edit interfaces interface-name aggregated-ether-options
```

2. Configure the minimum number of links.

```
[edit interfaces interface-name aggregated-ether-options]
```

```
user@host# set minimum-links number
```

On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, and EX 9200 switches, the valid range for **minimum-links *number*** is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for **minimum-links *number*** is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On EX8200 switches, the range of valid values for **minimum-links *number*** is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

If the number of links configured in an aggregated Ethernet interface is less than the minimum link value configured under the **aggregated-ether-options** statement, the configuration commit fails and an error message is displayed.

**Related  
Documentation**

- *aggregated-ether-options*
- *minimum-links*
- *Ethernet Interfaces Feature Guide for Routing Devices*



## Configuring Multichassis Link Aggregation

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two MC-LAG peers (for example, EX9200 switches). An MC-LAG provides redundancy and load balancing between the two MC-LAG peers, multihoming support, and a loop-free Layer 2 network without running the Spanning Tree Protocol (STP).

On one end of an MC-LAG, there is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to have an MC-LAG configured. On the other side of MC-LAG, there are two MC-LAG peers. Each of the MC-LAG peers has one or more physical links connected to a single client device.

The MC-LAG peers use Interchassis Control Protocol (ICCP) to exchange control information and coordinate with each other to ensure that data traffic is forwarded properly.



**NOTE:** An interface with an already configured IP address cannot form part of the aggregated Ethernet interface or multichassis aggregated Ethernet interface group.

For information about MC-LAG functional behaviors, see the “MC-LAG Configuration Guidelines and Functional Behavior” section in the topic [“Understanding Multichassis Link Aggregation” on page 2599](#). For a list of best practice MC-LAG configuration guidelines, see the “MC-LAG Configuration Guidelines and Functional Behavior” section in the topic [“Understanding Multichassis Link Aggregation” on page 2599](#).

Perform the following steps on each switch that hosts an MC-LAG:

1. Specify the same multichassis aggregated Ethernet identification number for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

```
[edit interfaces]
user@switch# set aex aggregated-ether-options mc-ae mc-ae-id number
```

For example:

```
[edit interfaces]
user@switch# set ael aggregated-ether-options mc-ae mc-ae-id 3
```

2. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

```
[edit interfaces]
user@switch# set aex aggregated-ether-options mc-ae chassis-id number
```

For example:

```
[edit interfaces]
user@switch# set ael aggregated-ether-options mc-ae chassis-id 0
```

3. Specify the mode of the MC-LAG the aggregated Ethernet interface belongs to.

```
[edit interfaces]
user@switch# set aex aggregated-ether-options mc-ae mode mode
```

For example:

```
[edit interfaces]
user@switch# set ael aggregated-ether-options mc-ae mode active-active
```

4. Specify whether the aggregated Ethernet interface participating in the MC-LAG is primary or secondary. Primary is **active**, and secondary is **standby**.



**NOTE:** You must configure status control on both switches that host the MC-LAG. If one switch is in active mode, the other must be in standby mode.

```
[edit interfaces]
user@switch# set aex aggregated-ether-options mc-ae status-control (active | standby)
```

For example:

```
[edit interfaces]
user@switch# set ael aggregated-ether-options mc-ae status-control active
```



**NOTE:** You can configure the `prefer-status-control-active` statement with the `mc-ae status-control standby` configuration to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure. Use this configuration only if you can ensure that ICCP will not go down unless the router or switch is down. You must also configure the `hold-time down` value (at the `[edit interfaces interface-name]` hierarchy level) for the ICL with the `mc-ae status-control standby` configuration to be higher than the ICCP BFD timeout. This configuration prevents data traffic loss by ensuring that when the router or switch with the `mc-ae status-control active` configuration goes down, the router or switch with the `mc-ae status-control standby` configuration does not go into standby mode.

To make the `prefer-status-control-active` configuration work with the `mc-ae status-control standby` configuration when an ICL logical interface is configured on an aggregated Ethernet interface, you must either configure the `lACP periodic interval` statement at the `[edit interfaces interface-name aggregated-ether-options]` hierarchy level as `slow` or configure the `detection-time threshold` statement at the `[edit protocols iccp peer liveness-detection]` hierarchy level as less than 3 seconds.

On EX9200 switches, the `prefer-status-control-active` statement was added in Junos OS Release 13.2R1.

5. Specify the same LACP system ID on each switch.

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lACP system-id mac-address
```

For example:

```
[edit interfaces]
user@switch# set ael aggregated-ether-options lACP system-id 00:01:02:03:04:05
```

6. Specify the same LACP administration key on each switch.

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lACP admin-key number
```

For example:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lACP admin-key 3
```

7. Configure ICCP by performing the following steps on each switch that hosts the MC-LAG:

- a. Configure the local IP address to be used by the switches that host the MC-LAG.

```
[edit protocols]
user@switch# set iccp local-ip-addr local-ip-address
```

For example:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.1
```

- b. (Optional) Configure the IP address of the switch and the time during which an ICCP connection must be established between the switches that host the MC-LAG.

Configured session establishment hold time results in faster ICCP connection establishment. The recommended value is 50 seconds.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address session-establishment-hold-time seconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 session-establishment-hold-time 50
```

- c. (Optional) We recommend that you configure the backup liveness detection feature to implement faster failover of data traffic during an MC-LAG peer reboot. Configure the **backup-liveness-detection** statement on the management interface (fxp0) only.



**NOTE:** On EX9200 switches, the **backup-liveness-detection** statement was added in Junos OS Release 13.2R1.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address backup-liveness-detection backup-peer-ip ip-address
```

For example:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.232
```

- d. Configure the minimum interval at which the switch must receive a reply from the other switch with which it has established a Bidirectional Forwarding Detection (BFD) session.



**NOTE:** Configuring the minimum receive interval is required to enable BFD. We recommend a minimum receive interval value of 6 seconds.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address liveness-detection minimum-receive-interval
seconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60
```

- e. Configure the minimum transmit interval during which a switch must receive a reply from a switch with which it has established a BFD session.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address liveness-detection transmit-interval
minimum-interval seconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval
60
```

8. Configure a multichassis protection link between the switches.

```
[edit]
user@switch# set multi-chassis multi-chassis-protection peer-ip-address interface
interface-name
```

For example:

```
[edit protocols]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

9. Disable Rapid Spanning Tree Protocol (RSTP) on the ICL interfaces on both switches.



**NOTE:** STP is not supported on the ICL or MC-LAG interfaces.

```
[edit]
user@switch# set protocols rstp interface interface-name disable
```

For example:

```
[edit]
user@switch# set protocols rstp interface ae0.0 disable
```

#### Related Documentation

- [Understanding Multichassis Link Aggregation on page 2599](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using VRRP on EX9200 Switches](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on EX9200 Switches](#)

## Configuring Energy Efficient Ethernet on Interfaces (CLI Procedure)

Energy Efficient Ethernet (EEE), an Institute of Electrical and Electronics Engineers (IEEE) 802.3az standard, reduces the power consumption of physical layer devices (PHYs) during periods of low link utilization. EEE saves energy by putting part of the transmission circuit into low power mode when a link is idle.



**NOTE:** Configure EEE only on EEE-capable Base-T copper Ethernet ports. If you configure EEE on unsupported ports, the console displays the message: “EEE not supported”.

This topic describes:

- [Enabling EEE on an EEE-Capable Base-T Copper Ethernet Port on page 2685](#)
- [Disabling EEE on a Base-T Copper Ethernet Port on page 2685](#)

### [Enabling EEE on an EEE-Capable Base-T Copper Ethernet Port](#)

To enable EEE on an EEE-capable Base-T copper Ethernet interface:

[edit]

```
user@switch# set interfaces interface-name ether-options ieee-802-3az-eee
```

You can view the EEE status by using the `show interfaces interface-name detail` command.

### [Disabling EEE on a Base-T Copper Ethernet Port](#)

To disable EEE on a Base-T copper Ethernet interface:

[edit]

```
user@switch# delete interfaces interface-name ether-options ieee-802-3az-eee
```

By default, EEE is disabled on EEE-capable ports.

#### Related Documentation

- [Verifying That EEE Is Saving Energy on Configured Ports on page 2837](#)
- [Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces on page 2590](#)

## Configuring Local Link Bias (CLI Procedure)

Local link bias is used to conserve bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis or VCF on a different member link in the LAG bundle.

You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG out of a local link. You should not enable local link bias if you want egress traffic load-balanced as it exits the Virtual Chassis or VCF.

To enable local link bias on a LAG bundle:

```
[edit]
user@switch# set interface aex aggregated-ether-options local-bias
where aex is the name of the aggregated Ethernet link bundle.
```

For instance, to enable local link bias on aggregated Ethernet interface ae0:

```
[edit]
user@switch# set interface ae0 aggregated-ether-options local-bias
```

### Related Documentation

- [Understanding Local Link Bias on page 2590](#)

## Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

Juniper Networks EX Series and QFX Series switches use a hashing algorithm to determine how to forward traffic over a Link Aggregation group (LAG) bundle or to the next-hop device when equal-cost multipath (ECMP) is enabled.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. You can configure some of the fields that are used by the hashing algorithm.

Configuring the fields used by the hashing algorithm is useful in scenarios where most of the traffic entering the bundle is similar and the traffic needs to be managed in the LAG bundle. For instance, if the only difference in the IP packets for all incoming traffic is the source and destination IP address, you can tune the hashing algorithm to make hashing decisions more efficiently by configuring the algorithm to make hashing decisions using only those fields.

- [Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing on page 2687](#)
- [Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing on page 2687](#)
- [Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing on page 2687](#)

### Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing

To configure the hashing algorithm to use fields in the Layer 2 header for hashing:

1. Configure the hash mode to Layer 2 header:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set hash-mode layer2-header
```

The default hash mode is Layer 2 payload. Therefore, this step must be performed if you have not previously configured the hash mode.

2. Configure the fields in the Layer 2 header that the hashing algorithm uses for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set layer2 {no-destination-mac-address | no-ether-type |
no-source-mac-address | vlan-id}
```

By default, the hashing algorithm uses the values in the destination MAC address, Ethertype, and source MAC address fields in the header to hash traffic on the LAG. You can configure the hashing algorithm to not use the values in these fields by configuring **no-destination-mac-address**, **no-ether-type**, or **no-source-mac-address**.

You can also configure the hashing algorithm to include the VLAN ID field in the header by configuring the **vlan-id** option.

If you want the hashing algorithm to not use the Ethertype field for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set layer2 no-ether-type
```

### Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing

To configure the hashing algorithm to use fields in the IP payload for hashing:

1. Configure the hash mode to Layer 2 payload:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set hash-mode layer2-payload
```

The IP payload is not checked by the hashing algorithm unless the hash mode is set to Layer 2 payload. The default hash mode is Layer 2 payload.

2. Configure the fields in the IP payload that the hashing algorithm uses for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet {no-ipv4-destination-address | no-ipv4-source-address |
no-l4-destination-port | no-l4-source-port | no-protocol | vlan-id}
```

For instance, if you want the hashing algorithm to ignore the Layer 4 destination port, Layer 4 source port, and protocol fields and instead hash traffic based only on the IPv4 source and destination addresses:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet no-l4-destination-port no-l4-source-port no-protocol
```

### Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing

To configure the hashing algorithm to use fields in the IPv6 payload for hashing:

1. Configure the hash mode to Layer 2 payload:

```
[edit forwarding-options enhanced-hash-key]
```

```
user@switch# set hash-mode layer2-payload
```

The IPv6 payload is not checked by the hashing algorithm unless the hash mode is set to Layer 2 payload. The default hash mode is Layer 2 payload.

2. Configure the fields in the IPv6 payload that the hashing algorithm uses for hashing:

```
[edit forwarding-options enhanced-hash-key]
```

```
user@switch# set inet6 {no-ipv6-destination-address | no-ipv6-source-address |  
no-l4-destination-port | no-l4-source-port | no-next-header | vlan-id}
```

For instance, if you want the hashing algorithm to ignore the Layer 4 destination port, Layer 4 source port, and the Next Header fields and instead hash traffic based only on the IPv6 source and IPv6 destination address fields only:

```
[edit forwarding-options enhanced-hash-key]
```

```
user@switch# set inet6 no-l4-destination-port no-l4-source-port no-next-header
```

- Related Documentation**
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2585](#)
  - [Understanding Aggregated Ethernet Interfaces and LACP on page 2582](#)

## Configuring Tagged Aggregated Ethernet Interfaces

To specify aggregated Ethernet interfaces, include the **vlan-tagging** statement at the **[edit interfaces aex]** hierarchy level:

```
[edit interfaces aex]  
vlan-tagging;
```

You must also include the **vlan-id** statement:

```
vlan-id number;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

For more information about the **vlan-tagging** and **vlan-id** statements, see [“802.1Q VLANs Overview” on page 2614](#).

- Related Documentation**
- [vlan-id](#)
  - [vlan-tagging on page 2834](#)



## Configuring a Layer 3 Subinterface (CLI Procedure)

EX Series switches use Layer 3 subinterfaces to divide a physical interface into multiple logical interfaces, each corresponding to a VLAN. The switch uses the Layer 3 subinterfaces to route traffic between subnets.

To configure Layer 3 subinterfaces, you enable VLAN tagging and partition one or more physical ports into multiple logical interfaces, each corresponding to a VLAN ID.

Before you begin, make sure you set up your VLANs. See *Configuring VLANs for EX Series Switches (CLI Procedure)* or “[Configuring VLANs for EX Series Switches \(J-Web Procedure\)](#)” on page 2334.

To configure Layer 3 subinterfaces:

1. Enable VLAN tagging:

```
[edit interfaces interface-name]
user@switch# set vlan-tagging
```

2. Bind each VLAN ID to a logical interface:

```
[edit interfaces interface-name]
user@switch# set unit logical-unit-number vlan-id (VLAN Tagging and Layer 3 Subinterfaces)
vlan-id-number
```

### Related Documentation

- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- [Verifying That Layer 3 Subinterfaces Are Working on page 2840](#)
- [Understanding Layer 3 Subinterfaces on page 2591](#)

## Configuring Unicast RPF (CLI Procedure)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Enabling unicast RPF on the switch interfaces filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. When a packet comes into an interface, if that interface is not the best return path to the source, the switch discards the packet. If the incoming interface is the best return path to the source, the switch forwards the packet.



**NOTE:** On EX3200, EX4200, and EX4300 switches, you can enable unicast RPF only globally—that is, on all switch interfaces. You cannot enable unicast RPF on a per-interface basis.

Before you begin:

- On an EX8200, EX6200, or QFX Series switch, ensure that the selected switch interface is symmetrically routed before you enable unicast RPF. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. Do not enable unicast RPF on asymmetrically routed interfaces. An

asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.

- On an EX3200, EX4200, or EX4300 switch, ensure that *all* switch interfaces are symmetrically routed before you enable unicast RPF on an interface. When you enable unicast RPF on any interface, it is enabled globally on all switch interfaces. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.

To enable unicast RPF, configure it explicitly on a selected customer-edge interface:

[edit interfaces]

user@switch# **set ge-1/0/10 unit 0 family inet *rpf-check***



**BEST PRACTICE:** On EX3200, EX4200, and EX4300 switches, unicast RPF is enabled globally on *all* switch interfaces, regardless of whether you configure it explicitly on only one interface or only on some interfaces.

On EX3200, EX4200, and EX4300 switches, we recommend that you enable unicast RPF explicitly on either all interfaces or only one interface. To avoid possible confusion, do not enable it on only some interfaces:

- Enabling unicast RPF explicitly on only one interface makes it easier if you choose to disable it in the future because you must explicitly disable unicast RPF on every interface on which you explicitly enabled it. If you explicitly enable unicast RPF on two interfaces and you disable it on only one interface, unicast RPF is still implicitly enabled globally on the switch. The drawback of this approach is that the switch displays the flag that indicates that unicast RPF is enabled only on interfaces on which unicast RPF is explicitly enabled, so even though unicast RPF is enabled on all interfaces, this status is not displayed.
- Enabling unicast RPF explicitly on all interfaces makes it easier to know whether unicast RPF is enabled on the switch because every interface shows the correct status. (Only interfaces on which you explicitly enable unicast RPF display the flag that indicates that unicast RPF is enabled.) The drawback of this approach is that if you want to disable unicast RPF, you must explicitly disable it on every interface. If unicast RPF is enabled on any interface, it is implicitly enabled on all interfaces.

---

**Related  
Documentation**

- *Example: Configuring Unicast RPF on an EX Series Switch*
- [Verifying Unicast RPF Status on page 2841](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 2691](#)
- [Troubleshooting Unicast RPF on page 2933](#)
- [Understanding Unicast RPF on page 2592](#)

## Disabling Unicast RPF (CLI Procedure)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Unicast RPF filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. If the network configuration changes so that an interface that has unicast RPF enabled becomes a trusted interface or becomes asymmetrically routed (the interface that receives a packet is not the best return path to the packet's source), disable unicast RPF.

To disable unicast RPF on an EX3200, EX4200, or EX4300 switch, you must delete it from every interface on which you explicitly configured it. If you do not disable unicast RPF on every interface on which you explicitly enabled it, it remains implicitly enabled on all interfaces. If you attempt to delete unicast RPF from an interface on which it was not explicitly enabled, the **warning: statement not found** message appears. If you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces of the EX3200, EX4200, or EX4300 switch.

On EX8200, EX6200, and QFX Series switches, the switch does not apply unicast RPF to an interface unless you explicitly enable that interface for unicast RPF.

To disable unicast RPF, delete its configuration from the interface:

[edit interfaces]

user@switch# **delete** ge-1/0/10 unit 0 family inet **rpf-check**



**NOTE:** On EX3200, EX4200, and EX4300 switches, if you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces.

### Related Documentation

- [Example: Configuring Unicast RPF on an EX Series Switch](#)
- [Verifying Unicast RPF Status on page 2841](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 2689](#)
- [Understanding Unicast RPF on page 2592](#)

## Configuring IP Directed Broadcast (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring IP Directed Broadcast (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

You can use IP directed broadcast on an EX Series switch to facilitate remote network management by sending broadcast packets to hosts on a specified subnet without broadcasting to the entire network. IP directed broadcast packets are broadcast on only the target subnet. The rest of the network treats IP directed broadcast packets as unicast packets and forwards them accordingly.

Before you begin to configure IP directed broadcast:

- Ensure that the subnet on which you want broadcast packets using IP direct broadcast is not directly connected to the Internet.
- Configure an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI) for the subnet that will be enabled for IP direct broadcast. See [“Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\)” on page 2340](#), [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#) , or [“Configuring VLANs for EX Series Switches \(J-Web Procedure\)” on page 2334](#).



**NOTE:** We recommend that you do not enable IP directed broadcast on subnets that have a direct connection to the Internet because of increased exposure to denial-of-service (DoS) attacks.

To enable IP directed broadcast for a specified subnet:

1. Add the target subnet’s logical interfaces to the VLAN:  

```
[edit interfaces]
user@switch# set ge-0/0/0.0 family ethernet-switching vlan members v1
user@switch# set ge-0/0/1.0 family ethernet-switching vlan members v1
```
2. Configure the Layer 3 interface on the VLAN that is the target of the IP directed broadcast packets:  

```
[edit interfaces]
user@switch# set irb.1 family inet address 10.1.2.1/24
```
3. Associate a Layer 3 interface with the VLAN:  

```
[edit vlans]
user@switch# set vl l3-interface (VLANs) irb.1
```
4. Enable the Layer 3 interface for the VLAN to receive IP directed broadcasts:  

```
[edit interfaces]
user@switch# set irb.1 family inet targeted-broadcast
```

**Related  
Documentation**

- [Example: Configuring IP Directed Broadcast on an EX Series Switch](#)
- [Understanding IP Directed Broadcast for EX Series Switches on page 2596](#)

## Tracing Operations of an Individual Router or Switch Interface

To trace the operations of individual router or switch interfaces, include the **traceoptions** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
traceoptions {
  flag flag;
```

```
}
```

You can specify the following interface tracing flags:

- **all**—Trace all interface operations.
- **event**—Trace all interface events.
- **ipc**—Trace all interface interprocess communication (IPC) messages.
- **media**—Trace all interface media changes.

The interfaces **traceoptions** statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system **syslog** files.

#### Related Documentation

- *Tracing Operations of the Interface Process*
- *Tracing Interface Operations Overview*

## Tracing Operations of the Interface Process

To trace the operations of the router or switch interface process, dcd, include the **traceoptions** statement at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
traceoptions {
  file <filename> <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag <flag> <disable>;
  no-remote-trace;
}
```

By default, interface process operations are placed in the file named dcd and three 1-MB files of tracing information are maintained.

You can specify the following flags in the **interfaces traceoptions** statement:

- **change-events**—Log changes that produce configuration events.
- **config-states**—Log the configuration state machine changes.
- **kernel**—Log configuration IPC messages to kernel.
- **kernel-detail**—Log details of configuration messages to kernel.

For general information about tracing, see the tracing and logging information in the *Junos OS Administration Library for Routing Devices*.

#### Related Documentation

- *Tracing Interface Operations Overview*
- [Tracing Operations of an Individual Router or Switch Interface on page 2692](#)

## Configuration Statements

---

- [\[edit chassis\] Configuration Statement Hierarchy on EX Series Switches on page 2696](#)
- [\[edit forwarding-options\] Configuration Statement Hierarchy on EX Series Switches on page 2698](#)
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2699](#)
- [\[edit interfaces ae\] Configuration Statement Hierarchy on EX Series Switches on page 2700](#)
- [\[edit interfaces et\] Configuration Statement Hierarchy on EX Series Switches on page 2705](#)
- [\[edit interfaces ge\] Configuration Statement Hierarchy on EX Series Switches on page 2711](#)
- [\[edit interfaces interface-range\] Configuration Statement Hierarchy on EX Series Switches on page 2716](#)
- [\[edit interfaces irb\] Configuration Statement Hierarchy on EX Series Switches on page 2725](#)
- [\[edit interfaces lo\] Configuration Statement Hierarchy on EX Series Switches on page 2728](#)
- [\[edit interfaces me\] Configuration Statement Hierarchy on EX Series Switches on page 2731](#)
- [\[edit interfaces vme\] Configuration Statement Hierarchy on EX Series Switches on page 2735](#)
- [\[edit interfaces xe\] Configuration Statement Hierarchy on EX Series Switches on page 2738](#)
- [\[edit protocols lacp\] Configuration Statement Hierarchy on EX Series Switches on page 2743](#)
- [802.3ad on page 2745](#)
- [accounting-profile on page 2746](#)
- [address on page 2747](#)
- [aggregated-devices on page 2749](#)
- [aggregated-ether-options on page 2750](#)
- [arp \(Interfaces\) on page 2752](#)
- [auto-negotiation on page 2753](#)
- [backup-liveness-detection on page 2754](#)
- [backup-peer-ip on page 2755](#)
- [bandwidth \(Interfaces\) on page 2756](#)
- [broadcast on page 2757](#)
- [chassis on page 2758](#)
- [description \(Interfaces\) on page 2760](#)
- [device-count on page 2761](#)
- [disable \(Interface\) on page 2762](#)

- [enhanced-hash-key](#) on page 2764
- [ether-options](#) on page 2766
- [ethernet \(Aggregated Devices\)](#) on page 2767
- [eui-64](#) on page 2767
- [family](#) on page 2768
- [filter](#) on page 2774
- [flow-control](#) on page 2775
- [force-up](#) on page 2776
- [gratuitous-arp-reply](#) on page 2776
- [hash-mode](#) on page 2777
- [hold-time \(Physical Interface\)](#) on page 2779
- [iccp](#) on page 2781
- [ieee-802-3az-eee](#) on page 2782
- [inet \(enhanced-hash-key\)](#) on page 2783
- [inet6 \(enhanced-hash-key\)](#) on page 2785
- [interface \(Multichassis Protection\)](#) on page 2786
- [interface-mode](#) on page 2787
- [interface-range](#) on page 2789
- [lACP \(Aggregated Ethernet\)](#) on page 2791
- [lACP \(802.3ad\)](#) on page 2793
- [layer2 \(enhanced-hash-key\)](#) on page 2794
- [link-mode](#) on page 2796
- [link-protection](#) on page 2798
- [link-speed \(Aggregated Ethernet\)](#) on page 2800
- [liveness-detection](#) on page 2801
- [local-bias](#) on page 2802
- [local-ip-addr \(ICCP\)](#) on page 2802
- [loopback \(Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet\)](#) on page 2803
- [mc-ae](#) on page 2804
- [mc-ae-id](#) on page 2806
- [member \(Interface Ranges\)](#) on page 2807
- [member-range](#) on page 2808
- [members](#) on page 2809
- [minimum-interval \(Liveness Detection\)](#) on page 2811
- [minimum-receive-interval \(Liveness Detection\)](#) on page 2811
- [mtu](#) on page 2812
- [multi-chassis](#) on page 2814

- [multi-chassis-protection](#) on page 2815
- [native-vlan-id](#) on page 2816
- [no-gratuitous-arp-request](#) on page 2817
- [no-redirects](#) on page 2817
- [peer \(ICCP\)](#) on page 2818
- [periodic](#) on page 2819
- [preferred](#) on page 2820
- [primary \(Address on Interface\)](#) on page 2821
- [proxy-arp](#) on page 2822
- [rpf-check](#) on page 2823
- [session-establishment-hold-time](#) on page 2824
- [speed \(Ethernet\)](#) on page 2825
- [traceoptions \(Individual Interfaces\)](#) on page 2826
- [traceoptions \(Interface Process\)](#) on page 2828
- [transmit-interval \(Liveness Detection\)](#) on page 2829
- [traps](#) on page 2830
- [unit](#) on page 2831
- [vlan \(802.1Q Tagging\)](#) on page 2832
- [vlan-id \(VLAN Tagging and Layer 3 Subinterfaces\)](#) on page 2833
- [vlan-tagging](#) on page 2834

## **[edit chassis] Configuration Statement Hierarchy on EX Series Switches**

This topic lists supported and unsupported configuration statements in the **[edit chassis]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit chassis\] Hierarchy Level](#) on page 2696

---

### **Supported Statements in the [edit chassis] Hierarchy Level**

The following hierarchy shows the **[edit chassis]** configuration statements supported on EX Series switches:

```
chassis {  
  aggregated-devices {  
    ethernet {
```



```

    device-count number;
    lACP {
        link-protection non-revertive;
        system-priority system-priority-number
    }
}
alarm {
    ethernet {
        link-down (ignore | red | yellow);
    }
    management-ethernet {
        link-down (ignore | red | yellow);
    }
}
container-devices {
    device-count device-count-number;
}
disk-partition {
    /config {
        level (full | high) {
            free-space (free-space-threshold-value | mb | percent);
        }
    }
    /var {
        level (full | high) {
            free-space (free-space-threshold-value | mb | percent);
        }
    }
}
}
fpc slot-number {
    pic pic-number {
        no-multi-rate;
        q-pic-large-buffer (large-scale | small-scale);
    }
}
}
maximum-ecmp maximum-ecmp-routes;
lcd-menu {
    fpc slot-number {
        menu-item menu-name);
        disable;
    }
}
pseudowire-service {
    device-count device-count-number;
}
psu {
    redundancy {
        n-plus-n;
    }
}
redundancy {
    graceful-switchover;
}
slow-pfe-alarm;
}

```

**Related Documentation**

- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\) on page 33](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 5116](#)
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 2512](#)
- [Configuring the Power Priority of Line Cards \(CLI Procedure\)](#)
- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade \(CLI Procedure\)](#)

## **[edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches**

This topic lists supported and unsupported configuration subhierarchies in the **[edit forwarding-options]** hierarchy level on EX Series switches.

- *Supported* subhierarchies are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* subhierarchies are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).
- [Supported Subhierarchies in the \[edit forwarding-options\] Hierarchy Level on page 2698](#)
- [Unsupported Subhierarchies in the \[edit forwarding-options\] Hierarchy Level on page 2698](#)

### **Supported Subhierarchies in the [edit forwarding-options] Hierarchy Level**

The following list shows the **[edit forwarding-options]** subhierarchies supported on EX Series switches:

Each of the following topics lists the statements at a subhierarchy of the **[edit forwarding-options]** hierarchy.

- [\[edit forwarding-options analyzer\] Configuration Statement Hierarchy on page 4158](#)
- [\[edit forwarding-options dhcp-relay\] Configuration Statement Hierarchy for EX Series Switches on page 1550](#)
- [\[edit forwarding-options enhanced-hash-key\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit forwarding-options port-mirroring\] Configuration Statement Hierarchy on page 4157](#)
- [\[edit forwarding-options storm-control-profiles\] Configuration Statement Hierarchy for EX Series Switches on page 2211](#)

### **Unsupported Subhierarchies in the [edit forwarding-options] Hierarchy Level**

All subhierarchies in the **[edit forwarding-options]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 290: Unsupported [edit forwarding-options] Subhierarchies on EX Series Switches

| Subhierarchy | Hierarchy Level           |
|--------------|---------------------------|
| accounting   | [edit forwarding-options] |
| helpers      | [edit forwarding-options] |
| sampling     | [edit forwarding-options] |

**Related  
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

## [edit interfaces] Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the [edit interfaces] hierarchy:

- [\[edit interfaces ae\] Configuration Statement Hierarchy on EX Series Switches on page 2700](#)
- [\[edit interfaces et\] Configuration Statement Hierarchy on EX Series Switches on page 2705](#)
- [\[edit interfaces ge\] Configuration Statement Hierarchy on EX Series Switches on page 2711](#)
- [\[edit interfaces interface-range\] Configuration Statement Hierarchy on EX Series Switches on page 2716](#)
- [\[edit interfaces irb\] Configuration Statement Hierarchy on EX Series Switches on page 2725](#)
- [\[edit interfaces lo\] Configuration Statement Hierarchy on EX Series Switches on page 360](#)
- [\[edit interfaces me\] Configuration Statement Hierarchy on EX Series Switches on page 363](#)
- [\[edit interfaces vme\] Configuration Statement Hierarchy on EX Series Switches on page 370](#)
- [\[edit interfaces xe\] Configuration Statement Hierarchy on EX Series Switches on page 2738](#)

**Related  
Documentation**

- [EX Series Switches Interfaces Overview on page 2577](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 2689](#)
- [Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\) on page 2340](#)
- [Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis \(CLI Procedure\)](#)

- [Junos OS Interfaces Fundamentals Configuration Guide](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

## [edit interfaces ae] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces ae]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces ae\] Hierarchy Level on page 2700](#)
- [Unsupported Statements in the \[edit interfaces ae\] Hierarchy Level on page 2704](#)

### Supported Statements in the [edit interfaces ae] Hierarchy Level

---

The following hierarchy shows the **[edit interfaces ae]** configuration statements supported on EX Series switches.

```
interfaces {
  aeX {
    accounting-profile name;
    aggregated-ether-options {
      ethernet-switch-profile {
        tag-protocol-id identifiers;
      }
      (flow-control | no-flow-control);
      lcp {
        (active | passive);
        link-protection {
          disable;
          (revertive | non-revertive);
        }
        periodic interval;
        system-priority number;
      }
      (link-protection | no-link-protection);
      link-speed speed;
      local-bias;
      (loopback | no-loopback);
      mc-ae {
        chassis-id chassis-id;
        events {
          iccp-peer-down {
            force-icl-down;
            prefer-status-control-active;
          }
        }
      }
    }
  }
}
```

```

    }
  }
  init-delay-time seconds;
  mc-ae-id mc-ae-id;
  mode (active-active | active-standby);
  redundancy-group group-id;
  revert-time revert-time;
  status-control (active | standby);
  switchover-mode (non-revertive | revertive);
}
minimum-links number;
rebalance-periodic;
}
description text;
disable;
encapsulation type;
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
mtu bytes;
native-vlan-id
no-gratuitous-arp-request;
traceoptions {
  flag flag;
}
(traps | no-traps);
unit logical-unit-number {
  accounting-profile name;
  arp-resp (restricted | unrestricted);
  bandwidth rate;
  description text;
  disable;
  encapsulation type;
  family ccc {
    filter {
      group group-number;
      input filter-name;
      input-list [filter-names];
      output filter-name;
      output-list [filter-names];
    }
    policer {
      input policer-name;
      output policer-name;
    }
  }
}
family ethernet-switching {
  filter {
    input filter-name;
    output filter-name;
  }
  interface-mode (access | trunk);
  recovery-timeout seconds;
  storm-control profile-name;
  vlan {
    members [members];
  }
}

```

```
}
family inet {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
  address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    preferred;
    primary;
    vrrp-group group-number {
      (accept-data | no-accept-data);
      advertise-interval seconds;
      advertisements-threshold number;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      (preempt | no-preempt) {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
      }
      virtual-address [addresses];
      vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
      }
    }
  }
}
filter {
  filter {
    input filter-name;
    output filter-name;
  }
  mtu bytes;
  no-neighbor-learn;
  no-redirects;
  primary;
  rpf-check {
    fail-filter filter-name;
    mode {
      loose;
    }
  }
}
family inet6 {
  accounting {
    destination-class-usage;
```

```

source-class-usage {
    input;
    output;
}
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-name;
            active-interface interface-name;
        }
    }
}
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time time;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
}

```

```

family iso {
    address interface-address;
    mtu bytes;
}
input-vlan-map action;
output-vlan-map action;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id vlan-id;
vlan-id-list [vlan-id vlan-id-vlan-id];
vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
vlan-tagging;
}

```

### Unsupported Statements in the [edit interfaces ae] Hierarchy Level

All statements in the [edit interfaces ae] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 291: Unsupported [edit interfaces ae] Configuration Statements for EX Series Switches

| Statement            | Hierarchy                                           |
|----------------------|-----------------------------------------------------|
| stacked-vlan-tagging | [edit interfaces ae]                                |
| admin-key            | [edit interfaces ae aggregated-ether-options lacp]  |
| system-id            | [edit interfaces ae aggregated-ether-options lacp]  |
| layer2-policer       | [edit interfaces ae unit]                           |
| native-inner-vlan-id | [edit interfaces ae unit]                           |
| swap-by-poppush      | [edit interfaces ae unit]                           |
| vlan-id-range        | [edit interfaces ae unit]                           |
| vlan-tags            | [edit interfaces ae unit]                           |
| mpls                 | [edit interfaces ae unit family]                    |
| vpls                 | [edit interfaces ae unit family]                    |
| bridge-domain-type   | [edit interfaces ae unit family ethernet-switching] |
| inner-vlan-id-list   | [edit interfaces ae unit family ethernet-switching] |
| vlan-rewrite         | [edit interfaces ae unit family ethernet-switching] |
| policer              | [edit interfaces ae unit family inet]               |



Table 291: Unsupported [edit interfaces ae] Configuration Statements for EX Series Switches (*continued*)

| Statement           | Hierarchy                                                                 |
|---------------------|---------------------------------------------------------------------------|
| sampling            | [edit interfaces ae unit family inet]                                     |
| service             | [edit interfaces ae unit family inet]                                     |
| simple-filter       | [edit interfaces ae unit family inet]                                     |
| targeted-broadcast  | [edit interfaces ae unit family inet]                                     |
| unnumbered-address  | [edit interfaces ae unit family inet]                                     |
| bandwidth-threshold | [edit interfaces ae unit family inet address vrrp-group track interface]  |
| service             | [edit interfaces ae unit family inet6]                                    |
| bandwidth-threshold | [edit interfaces ae unit family inet6 address vrrp-group track interface] |
| group               | [edit interfaces ae unit family inet6 filter]                             |
| pop                 | [edit interfaces ae unit input-vlan-map]                                  |
| push                | [edit interfaces ae unit output-vlan-map]                                 |

**Related Documentation** • [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

## [edit interfaces et] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces et]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces et\] Hierarchy Level on page 2706](#)
- [Unsupported Statements in the \[edit interfaces et\] Hierarchy Level on page 2709](#)

### Supported Statements in the [edit interfaces et] Hierarchy Level

The following hierarchy shows the [edit interfaces et] configuration statements supported on EX Series switches.

```
interfaces {
  et-fpc/pic/port {
    accounting-profile name;
    description text;
    disable;
    encapsulation type;
    ether-options {
      802.3ad {
        aex;
        (backup | primary);
        lacp {
          force-up;
          port-priority number;
        }
      }
      ethernet-switch-profile {
        tag-protocol-id [tpids];
      }
      (flow-control | no-flow-control);
      (loopback | no-loopback);
      no-auto-mdix;
    }
    flexible-vlan-tagging;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    native-vlan-id
    no-gratuitous-arp-request;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      bandwidth rate;
      description text;
      disable;
      encapsulation type;
      family ccc;
      filter {
        group group-number;
        input filter-name;
        input-list [filter-names];
        output filter-name;
        output-list [filter-names];
      }
      policer {
        input policer-name;
        output policer-name;
      }
    }
  }
}
```

```

}
family ethernet-switching {
  filter {
    input filter-name;
    output filter-name;
  }
  interface-mode (access | trunk);
  recovery-timeout seconds;
  storm-control profile-name;
  vlan {
    members (vlan-name | [-vlan-names] | all);
  }
}
family inet {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
}
address ipv4-address {
  arp ip-address (mac | multicast-mac) mac-address <publish>;
  broadcast address;
  preferred;
  primary;
  vrrp-group group-number {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    advertisements-threshold number;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    (preempt | no-preempt) {
      hold-time seconds;
    }
    priority number;
    track {
      interface interface-name {
        priority-cost number;
      }
      priority-hold-time seconds;
      route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [addresses];
    vrrp-inherit-from {
      active-group group-number;
      active-interface interface-name;
    }
  }
}
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;

```

```
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-name;
            active-interface interface-name;
        }
    }
}
(dad-disable | no-dad-disable);
filter {

    input filter-name;

    output filter-name;
```

```

}
mtu bytes;
nd6-stale-time time;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
family iso {
    address interface-address;
    mtu bytes;
}
input-vlan-map action;
output-vlan-map action;
proxy-arp (restricted | unrestricted);
swap-by-poppush;
(traps | no-traps);
vlan-id vlan-id-number;
vlan-id-list [vlan-id vlan-id-vlan-id];
}
vlan-tagging;
}
}

```

### Unsupported Statements in the [edit interfaces et] Hierarchy Level

All statements in the [edit interfaces et] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 292: Unsupported [edit interfaces et] Configuration Statements for EX Series Switches**

| Statement                 | Hierarchy                          |
|---------------------------|------------------------------------|
| passive-monitor-mode      | [edit interfaces et]               |
| stacked-vlan-tagging      | [edit interfaces et]               |
| asynchronous-notification | [edit interfaces et ether-options] |
| ignore-l3-incompletes     | [edit interfaces et ether-options] |
| mpls                      | [edit interfaces et ether-options] |
| source-address-filter     | [edit interfaces et ether-options] |

Table 292: Unsupported [edit interfaces et] Configuration Statements for EX Series Switches (*continued*)

| Statement            | Hierarchy                                                                 |
|----------------------|---------------------------------------------------------------------------|
| source-filtering     | [edit interfaces et ether-options]                                        |
| no-source-filtering  | [edit interfaces et ether-options]                                        |
| accept-source-mac    | [edit interfaces et unit]                                                 |
| layer2-policer       | [edit interfaces et unit]                                                 |
| native-inner-vlan-id | [edit interfaces et unit]                                                 |
| vlan-id-range        | [edit interfaces et unit]                                                 |
| vlan-tags            | [edit interfaces et unit]                                                 |
| mpls                 | [edit interfaces et unit family]                                          |
| tcc                  | [edit interfaces et unit family]                                          |
| vpls                 | [edit interfaces et unit family]                                          |
| bridge-domain-type   | [edit interfaces et unit family ethernet-switching]                       |
| inner-vlan-id-list   | [edit interfaces et unit family ethernet-switching]                       |
| vlan-rewrite         | [edit interfaces et unit family ethernet-switching]                       |
| policer              | [edit interfaces et unit family inet]                                     |
| sampling             | [edit interfaces et unit family inet]                                     |
| service              | [edit interfaces et unit family inet]                                     |
| targeted-broadcast   | [edit interfaces et unit family inet]                                     |
| unnumbered-address   | [edit interfaces et unit family inet]                                     |
| bandwidth-threshold  | [edit interfaces et unit family inet address vrrp-group track interface]  |
| service              | [edit interfaces et unit family inet6]                                    |
| bandwidth-threshold  | [edit interfaces et unit family inet6 address vrrp-group track interface] |
| group                | [edit interfaces et unit family inet6 filter]                             |
| pop                  | [edit interfaces et unit input-vlan-map]                                  |

Table 292: Unsupported [edit interfaces et] Configuration Statements for EX Series Switches (*continued*)

| Statement | Hierarchy                                 |
|-----------|-------------------------------------------|
| push      | [edit interfaces et unit output-vlan-map] |

**Related Documentation**

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

## [edit interfaces ge] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces ge]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces ge\] Hierarchy Level on page 2711](#)
- [Unsupported Statements in the \[edit interfaces ge\] Hierarchy Level on page 2715](#)

### Supported Statements in the [edit interfaces ge] Hierarchy Level

The following hierarchy shows the **[edit interfaces ge]** configuration statements supported on EX Series switches.

```

interfaces {
  ge-fpc/pic/port {
    accounting-profile name;
    description text;
    disable;
    encapsulation type;
    ether-options {
      802.3ad {
        aex;
        (backup | primary);
        lacp {
          force-up;
          port-priority number;
        }
      }
    }
    (auto-negotiation | no-auto-negotiation);
    ethernet-switch-profile {
      tag-protocol-id [tpids];
    }
    (flow-control | no-flow-control);
  }
}

```

```
ieee-802-3az-eee;  
(loopback | no-loopback);  
no-auto-mdix;  
}  
flexible-vlan-tagging;  
(gratuitous-arp-reply | no-gratuitous-arp-reply);  
hold-time up milliseconds down milliseconds;  
link-mode {  
    full-duplex;  
}  
mtu bytes;  
native-vlan-id  
no-gratuitous-arp-request;  
speed speed;  
traceoptions {  
    flag flag;  
}  
(traps | no-traps);  
unit logical-unit-number {  
    accounting-profile name;  
    arp-resp (restricted | unrestricted);  
    bandwidth rate;  
    description text;  
    disable;  
    encapsulation type;  
    family ccc;  
        filter {  
            group group-number;  
            input filter-name;  
            input-list [filter-names];  
            output filter-name;  
            output-list [filter-names];  
        }  
        policer {  
            input policer-name;  
            output policer-name;  
        }  
    }  
}  
family ethernet-switching {  
    filter {  
        input filter-name;  
        output filter-name;  
    }  
    interface-mode (access | trunk);  
    recovery-timeout seconds;  
    storm-control profile-name;  
    vlan {  
        members (vlan-name | [-vlan-names] | all);  
    }  
}  
family inet {  
    accounting {  
        destination-class-usage;  
        source-class-usage {  
            input;  
            output;  
        }  
    }  
}
```



```

    }
}
address ipv4-address {
  arp ip-address (mac | multicast-mac) mac-address <publish>;
  broadcast address;
  preferred;
  primary;
  vrrp-group group-number {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    advertisements-threshold number;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    (preempt | no-preempt) {
      hold-time seconds;
    }
    priority number;
    track {
      interface interface-name {
        priority-cost number;
      }
      priority-hold-time seconds;
      route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [addresses];
    vrrp-inherit-from {
      active-group group-number;
      active-interface interface-name;
    }
  }
}
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
  fail-filter filter-name;
  mode {
    loose;
  }
}
}
family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
}
address address {

```

```
eui-64;
ndp ip-address (mac | multicast-mac) mac-address <publish>;
preferred;
primary;
vrrp-inet6-group group-id {
    accept-data | no-accept-data;
    advertisements-threshold number;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    inet6-advertise-interval milliseconds;
    preempt | no-preempt {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
    }
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time time;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
family iso {
    address interface-address;
    mtu bytes;
}
input-vlan-map action;
interface-shared-with {
```

```

        psd-name;
    }
    output-vlan-map action;
    proxy-arp (restricted | unrestricted);
    swap-by-poppush;
    (traps | no-traps);
    vlan-id vlan-id-number;
    vlan-id-list [vlan-id vlan-id-vlan-id];
}
vlan-tagging;
}
}

```

### Unsupported Statements in the [edit interfaces ge] Hierarchy Level

All statements in the [edit interfaces ge] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 293: Unsupported [edit interfaces ge] Configuration Statements for EX Series Switches**

| Statement                 | Hierarchy                          |
|---------------------------|------------------------------------|
| passive-monitor-mode      | [edit interfaces ge]               |
| stacked-vlan-tagging      | [edit interfaces ge]               |
| asynchronous-notification | [edit interfaces ge ether-options] |
| ignore-l3-incompletes     | [edit interfaces ge ether-options] |
| mpls                      | [edit interfaces ge ether-options] |
| source-address-filter     | [edit interfaces ge ether-options] |
| source-filtering          | [edit interfaces ge ether-options] |
| no-source-filtering       | [edit interfaces ge ether-options] |
| accept-source-mac         | [edit interfaces ge unit]          |
| layer2-policer            | [edit interfaces ge unit]          |
| native-inner-vlan-id      | [edit interfaces ge unit]          |
| vlan-id-range             | [edit interfaces ge unit]          |
| vlan-tags                 | [edit interfaces ge unit]          |
| mpls                      | [edit interfaces ge unit family]   |
| tcc                       | [edit interfaces ge unit family]   |

Table 293: Unsupported [edit interfaces ge] Configuration Statements for EX Series Switches (*continued*)

| Statement           | Hierarchy                                                                 |
|---------------------|---------------------------------------------------------------------------|
| vpls                | [edit interfaces ge unit family]                                          |
| bridge-domain-type  | [edit interfaces ge unit family ethernet-switching]                       |
| inner-vlan-id-list  | [edit interfaces ge unit family ethernet-switching]                       |
| vlan-rewrite        | [edit interfaces ge unit family ethernet-switching]                       |
| policer             | [edit interfaces ge unit family inet]                                     |
| sampling            | [edit interfaces ge unit family inet]                                     |
| service             | [edit interfaces ge unit family inet]                                     |
| simple-filter       | [edit interfaces ge unit family inet]                                     |
| targeted-broadcast  | [edit interfaces ge unit family inet]                                     |
| unnumbered-address  | [edit interfaces ge unit family inet]                                     |
| bandwidth-threshold | [edit interfaces ge unit family inet address vrrp-group track interface]  |
| service             | [edit interfaces ge unit family inet6]                                    |
| bandwidth-threshold | [edit interfaces ge unit family inet6 address vrrp-group track interface] |
| group               | [edit interfaces ge unit family inet6 filter]                             |
| pop                 | [edit interfaces ge unit input-vlan-map]                                  |
| push                | [edit interfaces ge unit output-vlan-map]                                 |

**Related Documentation** • [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

### [edit interfaces interface-range] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit interfaces interface-range] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.

- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces interface-range\] Hierarchy Level on page 2717](#)
- [Unsupported Statements in the \[edit interfaces interface-range\] Hierarchy Level on page 2720](#)

### Supported Statements in the [edit interfaces interface-range] Hierarchy Level

The following hierarchy shows the **[edit interfaces interface-range]** configuration statements supported on EX Series switches.

```

interfaces {
  interface-range name {
    accounting-profile name;
    aggregated-ether-options {
      ethernet-switch-profile {
        tag-protocol-id identifier;
      }
      (flow-control | no-flow-control);
      lACP {
        (active | passive);
        admin-key key;
        periodic interval;
        system-id mac-address;
      }
      (link-protection | no-link-protection);
      link-speed speed;
      (loopback | no-loopback);
      minimum-links number;
      rebalance-periodic;
      source-address-filter filter;
      source-filtering | no-source-filtering;
    }
    description text;
    disable;
    ether-options {
      802.3ad {
        aex;
        (backup | primary);
        lACP {
          force-up;
        }
      }
      (auto-negotiation | no-auto-negotiation);
      (flow-control | no-flow-control);
      (loopback | no-loopback);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    link-mode mode;
    member interface-name;
  }
}

```

```
member-range starting-interface name to ending-interface name;  
mtu bytes;  
native-vlan-id  
no-gratuitous-arp-request;  
speed speed;  
traceoptions {  
    flag flag;  
}  
(traps | no-traps);  
unit logical-unit-number {  
    accept-source-mac {  
        mac-address mac-address {  
            policer {  
                input policer-name;  
                output policer-name;  
            }  
        }  
    }  
}  
accounting-profile name;  
arp-resp;  
bandwidth rate;  
description text;  
disable;  
family ccc;  
family ethernet-switching {  
    filter {  
        input filter-name;  
        output filter-name;  
    }  
    interface-mode (access | trunk);  
    recovery-timeout seconds;  
    storm-control profile-name;  
    vlan {  
        members [ members];  
    }  
}  
family inet {  
    accounting {  
        destination-class-usage;  
        source-class-usage;  
    }  
    address ipv4-address {  
        arp ip-address (mac | multicast-mac) mac-address <publish>;  
        broadcast address;  
        destination-class-usage;  
        destination-profile;  
        master-only;  
        preferred;  
        primary;  
        vrrp-group group-number {  
            (accept-data | no-accept-data);  
            advertise-interval seconds;  
            authentication-key key;  
            authentication-type authentication;  
            fast-interval milliseconds;  
            (preempt | no-preempt) {
```

```

        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [ addresses ];
    virtual-link-local-address address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check;
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage;
    }
    address address {
        eui-64;
        ndp ip-address (mac | multicast-mac) mac-address <publish>;
        preferred;
        primary;
        vrrp-inet6-group group-id {
            accept-data | no-accept-data;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            inet6-advertise-interval milliseconds;
            preempt | no-preempt {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    priority-cost number;
                }
                priority-hold-time seconds;
                route ( address | routing-instance routing-instance-name );
            }
            virtual-inet6-address [ addresses ];

```

```

    virtual-link-local-address ipv6-address;
  }
  vrrp-inherit-from {
    active-group group-name;
    active-interface interface-name;
  }
}
(dad-disable | no-dad-disable);
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
  input policer-name;
  output policer-name;
}
rpf-check;
}
family iso {
  address interface-address;
  mtu bytes;
}
minimum-links;
mtu;
proxy-arp (restricted | unrestricted);
swap-by-poppush;
(traps | no-traps);
vlan-id vlan-id-number;
}
vlan-tagging;
}

```

### Unsupported Statements in the [edit interfaces interface-range] Hierarchy Level

All statements in the [edit interfaces interface-range] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 294: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches**

| Statement                                                                                                | Hierarchy                         |
|----------------------------------------------------------------------------------------------------------|-----------------------------------|
| <i>NOTE:</i> Variables, such as <i>interface-range</i> , are not shown in the statements or hierarchies. |                                   |
| cesopsn-options                                                                                          | [edit interfaces interface-range] |
| container-options                                                                                        | [edit interfaces interface-range] |
| framing                                                                                                  | [edit interfaces interface-range] |
| lmi                                                                                                      | [edit interfaces interface-range] |



Table 294: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches (*continued*)

| Statement                 | Hierarchy                                       |
|---------------------------|-------------------------------------------------|
| logical-tunnel-options    | [edit interfaces interface-range]               |
| lsq-failure-options       | [edit interfaces interface-range]               |
| multiservice-options      | [edit interfaces interface-range]               |
| passive-monitor-mode      | [edit interfaces interface-range]               |
| ppp-options               | [edit interfaces interface-range]               |
| receive-bucket            | [edit interfaces interface-range]               |
| satop-options             | [edit interfaces interface-range]               |
| serial-options            | [edit interfaces interface-range]               |
| stacked-vlan-tagging      | [edit interfaces interface-range]               |
| transmit-bucket           | [edit interfaces interface-range]               |
| vdsl-options              | [edit interfaces interface-range]               |
| asynchronous-notification | [edit interfaces interface-range ether-options] |
| ethernet-switch-profile   | [edit interfaces interface-range ether-options] |
| ieee-802-3az-eee          | [edit interfaces interface-range ether-options] |
| ignore-l3-incompletes     | [edit interfaces interface-range ether-options] |
| mpls                      | [edit interfaces interface-range ether-options] |
| no-source-filtering       | [edit interfaces interface-range ether-options] |
| source-address-filter     | [edit interfaces interface-range ether-options] |
| source-filtering          | [edit interfaces interface-range ether-options] |
| accept-source-mac         | [edit interfaces interface-range unit]          |
| allow-any-vci             | [edit interfaces interface-range unit]          |
| atm-l2circuit-mode        | [edit interfaces interface-range unit]          |
| atm-scheduler-map         | [edit interfaces interface-range unit]          |

Table 294: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches (*continued*)

| Statement                   | Hierarchy                              |
|-----------------------------|----------------------------------------|
| cell-bundle-size            | [edit interfaces interface-range unit] |
| clear-don-fragment-bit      | [edit interfaces interface-range unit] |
| compression-device          | [edit interfaces interface-range unit] |
| copy-tos-to-outer-ip-header | [edit interfaces interface-range unit] |
| diable-mlppp-inner-ppp-pfc  | [edit interfaces interface-range unit] |
| dlci                        | [edit interfaces interface-range unit] |
| drop-timeout                | [edit interfaces interface-range unit] |
| epd-threshold               | [edit interfaces interface-range unit] |
| fragment-threshold          | [edit interfaces interface-range unit] |
| input-vlan-map              | [edit interfaces interface-range unit] |
| interface-shared-with       | [edit interfaces interface-range unit] |
| interleave-fragments        | [edit interfaces interface-range unit] |
| inverse-arp                 | [edit interfaces interface-range unit] |
| layer2-policer              | [edit interfaces interface-range unit] |
| link-layer-overhead         | [edit interfaces interface-range unit] |
| load-balancing-options      | [edit interfaces interface-range unit] |
| mrru                        | [edit interfaces interface-range unit] |
| multicast-dlci              | [edit interfaces interface-range unit] |
| multicast-vci               | [edit interfaces interface-range unit] |
| multilink-max-classes       | [edit interfaces interface-range unit] |
| multipoint                  | [edit interfaces interface-range unit] |
| native-inner-vlan-id        | [edit interfaces interface-range unit] |
| oam-liveness                | [edit interfaces interface-range unit] |

Table 294: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches (*continued*)

| Statement       | Hierarchy                                     |
|-----------------|-----------------------------------------------|
| oam-period      | [edit interfaces interface-range unit]        |
| output-vlan-map | [edit interfaces interface-range unit]        |
| peer-unit       | [edit interfaces interface-range unit]        |
| plp-to-clp      | [edit interfaces interface-range unit]        |
| point-to-point  | [edit interfaces interface-range unit]        |
| ppp-options     | [edit interfaces interface-range unit]        |
| receive-lap     | [edit interfaces interface-range unit]        |
| service-domain  | [edit interfaces interface-range unit]        |
| shaping         | [edit interfaces interface-range unit]        |
| short-sequence  | [edit interfaces interface-range unit]        |
| transmit-lsp    | [edit interfaces interface-range unit]        |
| transmit-weight | [edit interfaces interface-range unit]        |
| trunk-bandwidth | [edit interfaces interface-range unit]        |
| trunk-id        | [edit interfaces interface-range unit]        |
| tunnel          | [edit interfaces interface-range unit]        |
| vci             | [edit interfaces interface-range unit]        |
| vci-range       | [edit interfaces interface-range unit]        |
| vlan-id-list    | [edit interfaces interface-range unit]        |
| vpi             | [edit interfaces interface-range unit]        |
| mlfr-end-to-end | [edit interfaces interface-range unit family] |
| mlfr-uni-nni    | [edit interfaces interface-range unit family] |
| mlppp           | [edit interfaces interface-range unit family] |
| mpls            | [edit interfaces interface-range unit family] |

Table 294: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches (*continued*)

| Statement               | Hierarchy                                                                              |
|-------------------------|----------------------------------------------------------------------------------------|
| tcc                     | [edit interfaces interface-range unit family]                                          |
| vpls                    | [edit interfaces interface-range unit family]                                          |
| bridge-domain-type      | [edit interfaces interface-range unit family ethernet-switching]                       |
| inner-vlan-id-list      | [edit interfaces interface-range unit family ethernet-switching]                       |
| vlan-rewrite            | [edit interfaces interface-range unit family ethernet-switching]                       |
| ipsec-sa                | [edit interfaces interface-range unit family inet]                                     |
| multicast-only          | [edit interfaces interface-range unit family inet]                                     |
| negotiate-address       | [edit interfaces interface-range unit family inet]                                     |
| next-hop-tunnel         | [edit interfaces interface-range unit family inet]                                     |
| policer                 | [edit interfaces interface-range unit family inet]                                     |
| receive-options-packets | [edit interfaces interface-range unit family inet]                                     |
| receive-ttl-exceeded    | [edit interfaces interface-range unit family inet]                                     |
| sampling                | [edit interfaces interface-range unit family inet]                                     |
| service                 | [edit interfaces interface-range unit family inet]                                     |
| simple-filter           | [edit interfaces interface-range unit family inet]                                     |
| targeted-broadcast      | [edit interfaces interface-range unit family inet]                                     |
| unnumbered-address      | [edit interfaces interface-range unit family inet]                                     |
| bandwidth-threshold     | [edit interfaces interface-range unit family inet address vrrp-group track interface]  |
| service                 | [edit interfaces interface-range unit family inet6]                                    |
| group                   | [edit interfaces interface-range unit family inet6 filter]                             |
| bandwidth-threshold     | [edit interfaces interface-range unit family inet6 address vrrp-group track interface] |

**Related Documentation** • [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

## [edit interfaces irb] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces irb]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces irb\] Hierarchy Level on page 2725](#)
- [Unsupported Statements in the \[edit interfaces irb\] Hierarchy Level on page 2728](#)

### Supported Statements in the [edit interfaces irb] Hierarchy Level

The following hierarchy shows the **[edit interfaces irb]** configuration statements supported on EX Series switches.

```

interfaces {
  irb {
    accounting-profile name;
    description text;
    disable;

    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      bandwidth rate;
      description text;
      disable;
      family inet {
        accounting {
          destination-class-usage;
          source-class-usage {
            input;
            output;
          }
        }
      }
      address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
      }
    }
  }
}

```

```
preferred;
primary;
vrrp-group group-number {
  (accept-data | no-accept-data);
  advertise-interval seconds;
  advertisements-threshold number;
  authentication-key key;
  authentication-type authentication;
  fast-interval milliseconds;
  (preempt | no-preempt) {
    hold-time seconds;
  }
  priority number;
  track {
    interface interface-name {
      priority-cost number;
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
  }
  virtual-address [addresses];
  vrrp-inherit-from {
    active-group group-number;
    active-interface interface-name;
  }
}
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
  fail-filter filter-name;
  mode {
    loose;
  }
}
}
family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
}
address address {
  eui-64;
  ndp ip-address (mac | multicast-mac) mac-address <publish>;
  preferred;
  primary;
  vrrp-inet6-group group-id {
```

```

accept-data | no-accept-data;
advertisements-threshold number;
authentication-key key;
authentication-type authentication;
fast-interval milliseconds;
inet6-advertise-interval milliseconds;
preempt | no-preempt {
    hold-time seconds;
}
priority number;
track {
    interface interface-name {
        priority-cost number;
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
}
virtual-inet6-address [addresses];
virtual-link-local-address ipv6-address;
vrrp-inherit-from {
    active-group group-number;
    active-interface interface-name;
}
}
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
family iso {
    address interface-address;
    mtu bytes;
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}

```

}

### Unsupported Statements in the [edit interfaces irb] Hierarchy Level

All statements in the **[edit interfaces irb]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 295: Unsupported [edit interfaces irb] Configuration Statements for EX Series Switches**

| Statement           | Hierarchy                                                                        |
|---------------------|----------------------------------------------------------------------------------|
| encapsulation       | [edit interfaces irb]                                                            |
| layer2-policer      | [edit interfaces irb unit]                                                       |
| ccc                 | [edit interfaces irb unit family]                                                |
| mpls                | [edit interfaces irb unit family]                                                |
| tcc                 | [edit interfaces irb unit family]                                                |
| vpls                | [edit interfaces irb unit family]                                                |
| policer             | [edit interfaces irb unit family inet]                                           |
| sampling            | [edit interfaces irb unit family inet]                                           |
| service             | [edit interfaces irb unit family inet]                                           |
| targeted-broadcast  | [edit interfaces irb unit family inet]                                           |
| unnumbered-address  | [edit interfaces irb unit family inet]                                           |
| bandwidth-threshold | [edit interfaces irb unit family inet address vrrp-group track interface]        |
| service             | [edit interfaces irb unit family inet6]                                          |
| bandwidth-threshold | [edit interfaces irb unit family inet6 address vrrp-inet6-group track interface] |
| group               | [edit interfaces irb unit family inet6 filter]                                   |

**Related Documentation**

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

### [edit interfaces lo] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces lo]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.



- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces lo\] Hierarchy Level on page 2729](#)
- [Unsupported Statements in the \[edit interfaces lo\] Hierarchy Level on page 2731](#)

### Supported Statements in the [edit interfaces lo] Hierarchy Level

The following hierarchy shows the [edit interfaces lo] configuration statements supported on EX Series switches.

```

interfaces {
  lo0 {
    accounting-profile name;
    description text;
    disable;
    hold-time down milliseconds up milliseconds ;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family ccc;
      family inet {
        address ipv4-address {
          preferred;
          primary;
          vrrp-group group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            (preempt | no-preempt) {
              hold-time seconds;
            }
          }
          priority number;
          track {
            interface interface-name {
              bandwidth-threshold bandwidth;
              priority-cost number;
            }
          }
          priority-hold-time seconds;
          route ip-address/mask routing-instance instance-name priority-cost cost;
        }
      }
    }
  }
}

```

```
virtual-address [ addresses ];
virtual-link-local-address address;
vrrp-inherit-from {
    active-group group-number;
    active-interface interface-name;
}
}
}
dhcp {
    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinite);
    retransmission-attempt number;
    retransmission-interval sections;
    server-address ip-address;
    update-server
    vendor-id
}
filter {
    input filter-name;
    output filter-name;
}
no-neighbor-learn;
no-redirects;
primary;
}
family inet6 {
    address address {
        preferred;
        primary;
        vrrp-inet6-group group-id {
            accept-data | no-accept-data;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            inet6-advertise-interval milliseconds;
            preempt | no-preempt {
                hold-time seconds;
            }
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ( address | routing-instance routing-instance-name );
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-name;
            active-interface interface-name;
        }
    }
}
(dad-disable | no-dad-disable);
filter {
```

```

        group group-name;
        input filter-name;
        output filter-name;
    }
    no-neighbor-learn;
    policer {
        input policer-name;
        output policer-name;
    }
}
family iso {
    address interface-address;
}
family mpls;
(traps | no-traps);
}
}

```

### Unsupported Statements in the [edit interfaces lo] Hierarchy Level

All statements in the [edit interfaces lo] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 296: Unsupported [edit interfaces lo] Configuration Statements for EX Series Switches**

| Statement          | Hierarchy                             |
|--------------------|---------------------------------------|
| layer2-policer     | [edit interfaces lo unit]             |
| any                | [edit interfaces lo unit family]      |
| tcc                | [edit interfaces lo unit family]      |
| policer            | [edit interfaces lo unit family inet] |
| unnumbered-address | [edit interfaces lo unit family inet] |

- Related Documentation**
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)
  - [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

### [edit interfaces me] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit interfaces me] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.

- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces me\] Hierarchy Level on page 2732](#)
- [Unsupported Statements in the \[edit interfaces me\] Hierarchy Level on page 2734](#)

---

### Supported Statements in the [edit interfaces me] Hierarchy Level

---

The following hierarchy shows the [edit interfaces me] configuration statements supported on EX Series switches.

```
interfaces {
  me0 {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    no-gratuitous-arp-request;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family ethernet-switching {
        filter {
          input filter-name;
          output filter-name;
        }
        native-vlan-id vlan-id-number;
        port-mode (access | trunk);
        vlan {
          members [ members ];
        }
      }
    }
    family inet {
      accounting {
        destination-class-usage;
        source-class-usage {
          input;
          output;
        }
      }
      address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        master-only;
        preferred;
      }
    }
  }
}
```

```

    primary;
}
dhcp {
    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinite);
    retransmission-attempt number;
    retransmission-interval seconds;
    server-address ip-address;
    update-server
    vendor-id
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
primary;
rpf-check;
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
}
(dad-disable | no-dad-disable);
filter {
    group group-name;
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
swap-by-poppush;

```

```

        (traps | no-traps);
        vlan-id vlan-id-number;
    }
    vlan-tagging;
}

```

### Unsupported Statements in the [edit interfaces me] Hierarchy Level

All statements in the [edit interfaces me] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 297: Unsupported [edit interfaces me] Configuration Statements for EX Series Switches**

| Statement            | Hierarchy                                      |
|----------------------|------------------------------------------------|
| encapsulation        | [edit interfaces me]                           |
| link-mode            | [edit interfaces me]                           |
| encapsulation        | [edit interfaces me unit]                      |
| layer2-policer       | [edit interfaces me unit]                      |
| native-inner-vlan-id | [edit interfaces me unit]                      |
| vlan-id-list         | [edit interfaces me unit]                      |
| vlan-id-range        | [edit interfaces me unit]                      |
| ccc                  | [edit interfaces me unit family]               |
| tcc                  | [edit interfaces me unit family]               |
| vpls                 | [edit interfaces me unit family]               |
| no-redirects         | [edit interfaces me unit family inet]          |
| policer              | [edit interfaces me unit family inet]          |
| sampling             | [edit interfaces me unit family inet]          |
| service              | [edit interfaces me unit family inet]          |
| unnumbered-address   | [edit interfaces me unit family inet]          |
| vrrp-group           | [edit interfaces me unit family inet address]  |
| service              | [edit interfaces me unit family inet6]         |
| vrrp-inet6-group     | [edit interfaces me unit family inet6 address] |

- Related Documentation**
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)
  - [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

## [\[edit interfaces vme\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit interfaces vme]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces vme\] Hierarchy Level on page 2735](#)
- [Unsupported Statements in the \[edit interfaces vme\] Hierarchy Level on page 2738](#)

### [Supported Statements in the \[edit interfaces vme\] Hierarchy Level](#)

The following hierarchy shows the **[edit interfaces vme]** configuration statements supported on EX Series switches.

```

interfaces {
  vme {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family inet {
        accounting {
          destination-class-usage;
          source-class-usage {
            input;
            output;
          }
        }
      }
    }
  }
}

```

```
address ipv4-address {
  arp ip-address (mac | multicast-mac) mac-address <publish>;
  broadcast address;
  master-only;
  preferred;
  primary;
  vrrp-group group-number {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    (preempt | no-preempt) {
      hold-time seconds;
    }
    priority number;
    track {
      interface interface-name {
        bandwidth-threshold bandwidth;
        priority-cost number;
      }
      priority-hold-time seconds;
      route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [ addresses ];
    virtual-link-local-address address;
    vrrp-inherit-from {
      active-group group-number;
      active-interface interface-name;
    }
  }
}

dhcp {
  client-identifier (ascii client-id | hexadecimal client-id);
  lease-time (seconds | infinite);
  retransmission-attempt number;
  retransmission-interval sections;
  server-address ip-address;
  update-server
  vendor-id
}

filter {
  input filter-name;
  output filter-name;
}

mtu bytes;
no-neighbor-learn;
primary;
rpf-check;
}

family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
}
```



```

    }
  }
  address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
      accept-data | no-accept-data;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds;
      preempt | no-preempt {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          bandwidth-threshold bandwidth priority-cost number;
          priority-cost number;
        }
        priority-hold-time seconds;
        route ( address | routing-instance routing-instance-name );
      }
      virtual-inet6-address [addresses];
      virtual-link-local-address ipv6-address;
      vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
      }
    }
  }
  (dad-disable | no-dad-disable);
  filter {
    group group-name;
    input filter-name;
    output filter-name;
  }
  mtu bytes;
  no-neighbor-learn;
  policer {
    input policer-name;
    output policer-name;
  }
  rpf-check;
}
family iso {
  address interface-address;
  mtu bytes;
}
family mpls {
  mtu bytes;
}
(traps | no-traps);
vlan-id vlan-id-number;

```

```
    }  
    vlan-tagging;  
  }  
}
```

### Unsupported Statements in the [edit interfaces vme] Hierarchy Level

All statements in the [edit interfaces vme] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 342](#)
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

## [edit interfaces xe] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit interfaces xe] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces xe\] Hierarchy Level on page 2738](#)
- [Unsupported Statements in the \[edit interfaces xe\] Hierarchy Level on page 2742](#)

### Supported Statements in the [edit interfaces xe] Hierarchy Level

The following hierarchy shows the [edit interfaces xe] configuration statements supported on EX Series switches.

```
interfaces {  
  xe-fpc/pic/port {  
    accounting-profile name;  
    description text;  
    disable;  
    encapsulation type;  
    ether-options {  
      802.3ad {  
        (backup | primary);  
        lacp {  
          force-up;  
          port-priority number;  
        }  
      }  
    }  
    (auto-negotiation | no-auto-negotiation);  
    ethernet-switch-profile {
```

```

    tag-protocol-id [tpids];
  }
  (flow-control | no-flow-control);
  (loopback | no-loopback);
}
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
link mode {
    full-duplex;
}
mtu bytes;
native-vlan-id
no-gratuitous-arp-request;
traceoptions {
    flag flag;
}
(traps | no-traps);
unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    encapsulation type;
    family ccc {
        filter {
            input filter-name;
            output filter-name;
        }
        policer{
            input policer-name;
            output policer-name;
        }
    }
}
family ethernet-switching {
    filter {
        input filter-name;
        output filter-name;
    }
    interface-mode (access | trunk) ;
    recovery-timeout seconds;
    storm-control profile-name;
    vlan {
        members (vlan-name | [vlan-names] | all;
    }
}
family inet {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;

```

```
    broadcast address;
    preferred;
    primary;
    vrrp-group group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        advertisements-thresholds number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [addresses];
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
```

```

vrp-inet6-group group-id {
    accept-data | no-accept-data;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    inet6-advertise-interval milliseconds;
    preempt | no-preempt {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            priority-cost number;
        }
        priority-hold-time seconds;
        route (address | routing-instance routing-instance-name);
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
    }
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
family iso {
    address interface-address;
    mtu bytes;
}
input-vlan-map action;
interface-shared-with psdnumerical-index;
output-vlan-map action;
proxy-arp (restricted | unrestricted);
swap-by-poppush;
(traps | no-traps);
vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
vlan-id-list [vlan-id vlan-id-vlan-id];

```

```

    }
    vlan-tagging;
  }
}

```

### Unsupported Statements in the [edit interfaces xe] Hierarchy Level

All statements in the [edit interfaces xe] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 298: Unsupported [edit interfaces xe] Configuration Statements for EX Series Switches**

| Statement                 | Hierarchy                                           |
|---------------------------|-----------------------------------------------------|
| clocking                  | [edit interfaces xe]                                |
| framing                   | [edit interfaces xe]                                |
| passive-monitor-mode      | [edit interfaces xe]                                |
| stacked-vlan-tagging      | [edit interfaces xe]                                |
| asynchronous-notification | [edit interfaces xe ether-options]                  |
| ignore-l3-incompletes     | [edit interfaces xe ether-options]                  |
| mpls                      | [edit interfaces xe ether-options]                  |
| source-address-filter     | [edit interfaces xe ether-options]                  |
| source-filtering          | [edit interfaces xe ether-options]                  |
| no-source-filtering       | [edit interfaces xe ether-options]                  |
| accept-source-mac         | [edit interfaces xe unit]                           |
| layer2-policer            | [edit interfaces xe unit]                           |
| native-inner-vlan-id      | [edit interfaces xe unit]                           |
| vlan-id-range             | [edit interfaces xe unit]                           |
| vlan-tags                 | [edit interfaces xe unit]                           |
| mpls                      | [edit interfaces xe unit family]                    |
| tcc                       | [edit interfaces xe unit family]                    |
| vpls                      | [edit interfaces xe unit family]                    |
| bridge-domain-type        | [edit interfaces xe unit family ethernet-switching] |

Table 298: Unsupported [edit interfaces xe] Configuration Statements for EX Series Switches (*continued*)

| Statement           | Hierarchy                                                                 |
|---------------------|---------------------------------------------------------------------------|
| inner-vlan-id-list  | [edit interfaces xe unit family ethernet-switching]                       |
| vlan-rewrite        | [edit interfaces xe unit family ethernet-switching]                       |
| policer             | [edit interfaces xe unit family inet]                                     |
| sampling            | [edit interfaces xe unit family inet]                                     |
| service             | [edit interfaces xe unit family inet]                                     |
| simple-filter       | [edit interfaces xe unit family inet]                                     |
| targeted-broadcast  | [edit interfaces xe unit family inet]                                     |
| unnumbered-address  | [edit interfaces xe unit family inet]                                     |
| bandwidth-threshold | [edit interfaces xe unit family inet address vrrp-group track interface]  |
| service             | [edit interfaces xe unit family inet6]                                    |
| bandwidth-threshold | [edit interfaces xe unit family inet6 address vrrp-group track interface] |
| group               | [edit interfaces xe unit family inet6 filter]                             |
| pop                 | [edit interfaces xe unit input-vlan-map]                                  |
| push                | [edit interfaces xe unit output-vlan-map]                                 |

**Related Documentation** • [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 2366](#)

## [edit protocols lACP] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit protocols lACP] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols lacp\] Hierarchy Level on page 2744](#)
- [Unsupported Statements in the \[edit protocols lacp\] Hierarchy Level on page 2744](#)

---

### Supported Statements in the [edit protocols lacp] Hierarchy Level

The following hierarchy shows the **[edit protocols lacp]** configuration statements supported on EX Series switches:

```
protocols {
  lacp {
    ppm {
      centralized
    }
    traceoptions {
      file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
      flag <flag>;
      no-remote-trace;
    }
  }
}
```

---

### Unsupported Statements in the [edit protocols lacp] Hierarchy Level

All statements in the **[edit protocols lacp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)



## 802.3ad

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> 802.3ad {     aex;     (backup   primary);     lacp {         force-up;         port-priority     } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ether-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Configure membership in a link aggregation group (LAG).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <ul style="list-style-type: none"> <li>• <b>aex</b>—Name of the LAG.</li> <li>• <b>backup</b>—Designate the interface as the backup interface for link-protection mode.</li> <li>• <b>primary</b>—Designate the interface as the primary interface for link-protection mode.</li> </ul> <p>The remaining statements are described separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li> <li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li> <li>• <i>Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches</i></li> <li>• <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure) on page 2667</a></li> <li>• <a href="#">Configuring Aggregated Ethernet LACP (CLI Procedure) on page 2671</a></li> <li>• <a href="#">Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 2672</a></li> </ul> |

## accounting-profile

---

|                                 |                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>accounting-profile <i>name</i>;</code>                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ],<br>[edit interfaces interface-range <i>name</i> ]                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                              |
| <b>Description</b>              | Enable collection of accounting data for the specified physical or logical interface or interface range.                                                                                                                               |
| <b>Options</b>                  | <i>name</i> —Name of the accounting profile.                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Applying an Accounting Profile to the Physical Interface on page 2660</a></li><li>• <a href="#">Applying an Accounting Profile to the Logical Interface on page 2661</a></li></ul> |

## address

```

Syntax  address address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        destination address;
        destination-profile name;
        eui-64;
        master-only;
        multipoint-destination address dlcid dlcid-identifier;
        multipoint-destination address {
            epd-threshold cells;
            inverse-arp;
            oam-liveness {
                up-count cells;
                down-count cells;
            }
            oam-period (disable | seconds);
            shaping {
                (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
                 length);
                queue-length number;
            }
            vci vpi-identifier.vci-identifier;
        }
        primary;
        preferred;
        (vrrp-group | vrrp-inet6-group) group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-type authentication;
            authentication-key key;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority-number number;
            track {
                priority-cost seconds;
                priority-hold-time interface-name {
                    interface priority;
                    bandwidth-threshold bits-per-second {
                        priority;
                    }
                }
            }
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
    }
}

```

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family *family*],  
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*  
 family *family*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the interface address.

**Options** *address*—Address of the interface.

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.



**NOTE:** If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration, the remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
}
```

For more information on configuring the same address on multiple interfaces, see [“Configuring the Interface Address” on page 2632](#).

---

The remaining statements are explained separately.



**NOTE:** The edit logical-systems hierarchy is not available on QFabric systems.

---

**Required Privilege** interface—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring the Protocol Family*
  - *Junos OS Administration Library for Routing Devices*
  - [family on page 1647](#)
  - *negotiate-address*
  - *unnumbered-address (Ethernet)*

## aggregated-devices

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>aggregated-devices {   ethernet (Aggregated Devices) {     device-count <i>number</i>;     lacp   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Configure properties for aggregated devices on the switch.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>                  | Aggregated devices are disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li> <li>• <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure) on page 2667</a></li> <li>• <a href="#">Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 2672</a></li> <li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2582</a></li> <li>• <i>Junos OS Ethernet Interfaces Configuration Guide</i></li> </ul> |

## aggregated-ether-options

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> aggregated-ether-options {   ethernet-switch-profile {     tag-protocol-id;   }   (flow-control   no-flow-control);   lacp {     (active   passive);     admin-key <i>key</i>;     periodic <i>interval</i>;     system-id <i>mac-address</i>;   }   (link-protection   no-link-protection);   link-speed <i>speed</i>;   local-bias;   logical-interface-fpc-redundancy;   (loopback   no-loopback);   mc-ae {     chassis-id <i>chassis-id</i>;     events {       iccp-peer-down {         force-icl-down;         prefer-status-control-active;       }     }     init-delay-time <i>seconds</i>;     mc-ae-id <i>mc-ae-id</i>;     mode (active-active   active-standby);     redundancy-group <i>group-id</i>;     revert-time <i>revert-time</i>;     status-control (active   standby);     switchover-mode (non-revertive   revertive);   }   minimum-links <i>number</i>;   system-priority } </pre> |
| <b>Hierarchy Level</b>          | [edit interfaces aex]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3R2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure the aggregated Ethernet properties of a specific aggregated Ethernet interface.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 2671](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\) on page 2672](#)
- [Configuring Q-in-Q Tunneling \(CLI Procedure\) on page 2351](#)
- *Junos OS Ethernet Interfaces Configuration Guide*

## arp (Interfaces)

---

|                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                       | <code>arp ip-address (mac   multicast-mac) mac-address publish;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>                                                                                                                                                              | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                                                                                                                                                          | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>                                                                                                                                                                  | For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries, mapping IP addresses to MAC addresses.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                                                                                                                                                      | <p><b>ip-address</b>—IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing <b>address</b> statement.</p> <p><b>mac mac-address</b>—MAC address to map to the IP address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, <b>0011.2233.4455</b> or <b>00:11:22:33:44:55</b>.</p> <p><b>multicast-mac mac-address</b>—Multicast MAC address to map to the IP address. Specify the multicast MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, <b>0011.2233.4455</b> or <b>00:11:22:33:44:55</b>.</p> <p><b>publish</b>—(Optional) Have the router or switch reply to ARP requests for the specified IP address. If you omit this option, the router or switch uses the entry to reach the destination but does not reply to ARP requests.</p> |
| <div> <b>NOTE:</b> The edit logical-systems hierarchy is not available on QFabric systems.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b>                                                                                                                                                     | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>                                                                                                                                                        | <ul style="list-style-type: none"><li>• <a href="#">Configuring Static ARP Table Entries on page 2664</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |



## auto-negotiation

|                            |                                                                                                                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | (auto-negotiation   no-auto-negotiation) <remote-fault (local-interface-online   local-interface-offline)>;                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> ether-options],<br>[edit interfaces <i>interface-name</i> gigether-options],<br>[edit interfaces <i>ge-pim</i> /0/0 switch-options switch-port <i>port-number</i> ]                                                                                     |
| <b>Release Information</b> | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 8.4 for J Series Services Routers.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. |
| <b>Description</b>         | For Gigabit Ethernet interfaces on M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers explicitly enable autonegotiation and remote fault. For EX Series switches and J Series Services Routers, explicitly enable autonegotiation only.                                  |

- **auto-negotiation**—Enables autonegotiation. This is the default.
- **no-auto-negotiation**—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure the link mode and speed.

When you configure Tri-Rate Ethernet copper interfaces to operate at 1 Gbps, autonegotiation must be enabled.



**NOTE:** On EX Series switches, an interface configuration that disables autonegotiation and manually sets the link speed to 1 Gbps is accepted when you commit the configuration; however, if the interface you are configuring is a Tri-Rate Ethernet copper interface, the configuration is ignored as invalid and autonegotiation is enabled by default.

To correct the invalid configuration and disable autonegotiation:

1. Delete the **no-auto-negotiation** statement and commit the configuration.
2. Set the link speed to 10 or 100 Mbps, set **no-auto-negotiation**, and commit the configuration.

On J Series Services Routers with universal Physical Interface Modules (uPIMs) and on EX Series switches, if the link speed and duplex mode are also configured, the interfaces use the values configured as the desired values in the negotiation. If autonegotiation is disabled, the link speed and link mode must be configured.



**NOTE:** On T4000 routers, the **auto-negotiation** command is ignored for interfaces other than Gigabit Ethernet.

|                                 |                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b>remote-fault (local-interface-online   local-interface-offline)</b> —(Optional) For M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers only, manually configure remote fault on an interface.<br><b>Default:</b> local-interface-online                                                                                                 |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Gigabit Ethernet Autonegotiation Overview</i></li><li>• <i>Configuring Gigabit Ethernet Interfaces on J Series Services Routers</i></li><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li></ul> |

---

## backup-liveness-detection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>backup-liveness-detection {</b><br><b>backup-peer-ip</b> <i>ip4-address</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <b>iccp</b> peer]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2R1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Backup liveness detection determines the peer status (whether it is up or down) by exchanging keep alive messages (UDP-based packets) over the management link between the two Interchassis Control Protocol (ICCP) peers. When an ICCP connection is operationally down, the status of the peers hosting a multichassis link aggregation group (MC-LAG) is detected by sending liveness detection requests to each other. Peers must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, the liveness detection check fails, and a failure action is implemented. Backup liveness detection must be configured on both peers hosting the MC-LAG. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

---

## backup-peer-ip

---

|                                 |                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | backup-peer-ip <i>ip4-address</i> ;                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer backup-liveness-detection</a> ]                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2R1 for EX Series switches. |
| <b>Description</b>              | Specify the IP address of the peer being used as a backup peer in the Bidirectional Forwarding Detection (BFD) configuration.                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                          |

## bandwidth (Interfaces)

---

|                            |                                                                                                                                                                                                  |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>bandwidth rate;</code>                                                                                                                                                                     |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                        |
| <b>Description</b>         | Configure an informational-only bandwidth value for an interface. This statement is valid for all logical interface types except multilink and aggregated interfaces.                            |



**NOTE:** We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the `bandwidth` statement affects how the interface cost is calculated for a dynamic routing protocol, such as OSPF. By default, the interface cost for a dynamic routing protocol is calculated using the following formula:

$$\text{cost} = \text{reference-bandwidth} / \text{bandwidth},$$

where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the `bandwidth` statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <b>rate</b> —Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation <b>c</b> ; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps.<br><b>Range:</b> Not limited. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Interface Bandwidth on page 2637</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## broadcast

|                            |                                                                                                                                                                                                                                                                                                 |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>broadcast address;</code>                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> <b>address</b> <i>address</i> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                  |
| <b>Description</b>         | Set the broadcast address on the network or subnet. On a subnet you cannot specify a host address of 0, nor can you specify a broadcast address.                                                                                                                                                |
| <b>Default</b>             | The default broadcast address has a host portion of all ones.                                                                                                                                                                                                                                   |
| <b>Options</b>             | <b>address</b> —Broadcast address. The address must have a host portion of either all ones or all zeros. You cannot specify the addresses 0.0.0.0 or 255.255.255.255.                                                                                                                           |



**NOTE:** The edit logical-systems hierarchy is not available on QFabric systems.

|                                 |                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Interface Address on page 2632</a></li> </ul>      |

## chassis

```
Syntax  chassis {
        aggregated-devices {
            ethernet (Aggregated Devices) {
                device-count number;
            }
        }
        auto-image-upgrade;
        fpc slot {
            pic pic-number {
                sfppplus {
                    pic-mode mode;
                }
            }
            power-budget-priority priority;
        }
        lcd-menu {
            fpc slot-number {
                menu-item (menu-name | menu-option) {
                    disable;
                }
            }
        }
        nssu {
            upgrade-group group-name {
                fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
                member (NSSU Upgrade Groups) member-id {
                    fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
                }
            }
        }
        psu {
            redundancy {
                n-plus-n (Power Management);
            }
        }
        redundancy {
            graceful-switchover;
        }
    }
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure chassis-specific properties for the switch.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

**Related  
Documentation**

- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- [\*Upgrading Software by Using Automatic Software Download\*](#)
- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\) on page 33](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 5116](#)
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 2512](#)
- [\*Configuring the Power Priority of Line Cards \(CLI Procedure\)\*](#)
- [\*Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade \(CLI Procedure\)\*](#)

## description (Interfaces)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description text;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <code>[edit interfaces <i>interface-name</i>],</code><br><code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the <b>show interfaces</b> commands, and is also exposed in the <b>ifAlias</b> Management Information Base (MIB) object. It has no effect on the operation of the interface on the router or switch.</p> <p>The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.</p>                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>text</b> —Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <b>interface</b> —To view this statement in the configuration.<br><b>interface-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Interface Description</i></li><li>• <a href="#">Adding a Logical Unit Description to the Configuration on page 2629</a></li><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li><li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li><li>• <a href="#">Using DHCP Relay Agent Option 82 Information on page 1464</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li><li>• <a href="#">Example: Connecting Access Switches to a Distribution Switch on page 2291</a></li></ul> |



## device-count

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>device-count <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis aggregated-devices ethernet (Aggregated Devices)</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Range updated in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure the number of aggregated Ethernet logical devices available to the switch.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b><i>number</i></b>—Maximum number of aggregated Ethernet logical interfaces on the switch.</p> <p><b>Range:</b> 1 through 32 for EX2200, EX3200, and standalone EX3300 switches and for EX3300 Virtual Chassis</p> <p><b>Range:</b> 1 through 64 for standalone EX4200, standalone EX4500, and EX6200 switches and for EX4200 and EX4500 Virtual Chassis</p> <p><b>Range:</b> 1 through 239 for EX8200 Virtual Chassis</p> <p><b>Range:</b> 1 through 255 for standalone EX8200 switches</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li> <li>• <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure) on page 2667</a></li> <li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li> </ul>                                                                                                 |

## disable (Interface)

---

|                            |                                                                                                                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | disable;                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.                           |
| <b>Description</b>         | Disable a physical or a logical interface, effectively unconfiguring it.                                                                                                                                                                      |



---

### CAUTION:

- Dynamic subscribers and logical interfaces use physical interfaces for connection to the network. The Junos OS allows you to set the interface to disable and commit the change while dynamic subscribers and logical interfaces are still active. This action results in the loss of all subscriber connections on the interface. Use care when disabling interfaces.
  - If aggregated SONET links are configured between a T1600 router and a T4000 router, interface traffic is disrupted when you disable the physical interface configured on the T1600 router. If you want to remove the interface, we recommend that you deactivate the interface instead of disabling it.
- 



---

### NOTE:

- When you use the disable statement at the [edit interfaces] hierarchy level, depending on the PIC type, the interface might or might not turn off the laser. Older PIC transceivers do not support turning off the laser, but newer Gigabit Ethernet (GE) PICs with SFP and XFP transceivers and ATM MIC with SFP do support it and the laser will be turned off when the interface is disabled. If the ATM MIC with SFP is part of an APS group, then the laser will not be turned off when you use the disable statement at the [edit interfaces] hierarchy level..
  - When you disable or deactivate an interface, then all the references made to the deactivated interface must be removed from the routing instance.
- 



---

**WARNING:** Do not stare into the laser beam or view it directly with optical instruments even if the interface has been disabled.

---

**Required Privilege** interface—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

**Related Documentation**

- [Disabling a Physical Interface on page 2630](#)
- [Disabling a Logical Interface on page 2631](#)

## enhanced-hash-key

---

**Syntax**    enhanced-hash-key {  
              ecmp-resilient-hash;  
              fabric-load-balance {  
                  flowlet {  
                      inactivity-interval *interval*;  
                  }  
              per-packet;  
          }  
              hash-mode {  
                  layer2-header;  
                  layer2-payload;  
              }  
              inet {  
                  no-ipv4-destination-address;  
                  no-ipv4-source-address;  
                  no-l4-destination-port;  
                  no-l4-source-port;  
                  no-protocol;  
                  vlan-id;  
              }  
              inet6 {  
                  no-ipv6-destination-address;  
                  no-ipv6-source-address;  
                  no-l4-destination-port;  
                  no-l4-source-port;  
                  no-next-header;  
                  vlan-id;  
              }  
              layer2 {  
                  no-destination-mac-address;  
                  no-ether-type;  
                  no-source-mac-address;  
                  vlan-id;  
              }  
          }  
      }

**Hierarchy Level**    [edit forwarding-options]

**Release Information**    Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.  
                              Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.  
                              The **fabric-load-balance** statement introduced in Junos OS Release 14.1X53-D10.

**Description**    Configure the hashing key used to hash link aggregation group (LAG) and equal-cost multipath (ECMP) traffic, or enable adaptive load balancing (ALB) in a Virtual Chassis Fabric (VCF).

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The remaining statements are explained separately.

|                           |                                                               |
|---------------------------|---------------------------------------------------------------|
| <b>Required Privilege</b> | interface—To view this statement in the configuration.        |
| <b>Level</b>              | interface-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                             |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2686</a></li><li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2585</a></li></ul> |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## ether-options

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ether-options {<br/>    802.3ad {<br/>        aex;<br/>        (backup   primary);<br/>        lacp {<br/>            force-up;<br/>            port-priority<br/>        }<br/>    }<br/>    (auto-negotiation   no-auto-negotiation);<br/>    ethernet-switch-profile {<br/>        tag-protocol-id;<br/>    }<br/>    (flow-control   no-flow-control);<br/>    ieee-802-3az-eee;<br/>    link-mode <i>mode</i>;<br/>    (loopback   no-loopback);<br/>    speed (<i>speed</i>   auto-negotiation);<br/>}</pre>                                                                          |
| <b>Hierarchy Level</b>          | <pre>[edit interfaces <i>interface-name</i>],<br/>[edit interfaces interface-range <i>range</i>]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3R2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure Ethernet properties for a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 2619</a></li><li>• <a href="#">Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 2672</a></li><li>• <a href="#">Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</a></li><li>• <i>Junos OS Ethernet Interfaces Configuration Guide</i></li></ul> |

## ethernet (Aggregated Devices)

|                                 |                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ethernet {   device-count <i>number</i>;   lacp {     link-protection {       non-revertive;     }     system-priority;   } }</pre>                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis aggregated-devices</a> ]                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Configure properties for Ethernet aggregated devices on the switch.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure) on page 2667</a></li> <li>• <a href="#">Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 2672</a></li> <li>• <a href="#">Junos OS Ethernet Interfaces Configuration Guide</a></li> </ul> |

## eui-64

|                                 |                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | eui-64;                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>number</i> family inet6 address <i>address</i> ]                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for the QFX Series.</p> |
| <b>Description</b>              | For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.                                                                                 |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Interface Address on page 2632</a></li> </ul>                                                                                            |

## family

---

**Syntax** [family ccc on page 1647](#)  
[family ethernet-switching on page 1647](#)  
[family inet on page 1647](#)  
[family inet6 on page 1648](#)  
[family iso on page 1648](#)

**family ccc** family ccc;  
filter {  
    group *group-number*;  
    input *filter-name*;  
    input-list [*filter-names*];  
    output *filter-name*;  
    output-list [*filter-names*];  
}  
policer {  
    input *policer-name*;  
    output *policer-name*;  
}  
}

**family ethernet-switching** family ethernet-switching {  
    filter {  
        input *filter-name*;  
        output *filter-name*;  
    }  
    [interface-mode](#) (access | trunk);  
    [recovery-timeout](#) *seconds*;  
    storm-control *profile-name*;  
    vlan {  
        members (*vlan-name* | [*-vlan-names*] | all);  
    }  
}

**family inet** family inet {  
    accounting {  
        destination-class-usage;  
        source-class-usage {  
            input;  
            output;  
        }  
    }  
    [address](#) *ipv4-address* {  
        [arp](#) *ip-address* (mac | multicast-mac) *mac-address* <publish>;  
        [broadcast](#) *address*;  
        [preferred](#);  
        [primary](#);  
        vrrp-group *group-number* {  
            (accept-data | no-accept-data);  
            advertise-interval *seconds*;  
            advertisements-threshold *number*;  
            authentication-key *key*;  
            authentication-type *authentication*;



```
fast-interval milliseconds;  
(preempt | no-preempt) {  
    hold-time seconds;  
}  
priority number;  
track {  
    interface interface-name {  
        priority-cost number;  
    }  
    priority-hold-time seconds;  
    route ip-address/mask routing-instance instance-name priority-cost cost;  
}  
virtual-address [addresses];  
vrrp-inherit-from {  
    active-group group-number;  
    active-interface interface-name;  
}  
}  
}  
filter {  
    input filter-name;  
    output filter-name;  
}  
mtu bytes;  
no-neighbor-learn;  
no-redirects;  
primary;  
rpf-check {  
    fail-filter filter-name;  
    mode {  
        loose;  
    }  
}  
}
```

```
family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
  address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
      accept-data | no-accept-data;
      advertisements-threshold number;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds;
      preempt | no-preempt {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
      }
      virtual-inet6-address [addresses];
      virtual-link-local-address ipv6-address;
      vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
      }
    }
  }
  (dad-disable | no-dad-disable);
  filter {
    input filter-name;
    output filter-name;
  }
  mtu bytes;
  nd6-stale-time time;
  no-neighbor-learn;
  no-redirects;
  policer {
    input policer-name;
    output policer-name;
  }
  rpf-check {
    fail-filter filter-name;
    mode {
      loose;
    }
  }
}
```

```

    }
  }
}

```

```

family iso {
  address interface-address;
  mtu bytes;
}

```

|                     |                                                                                                                                                                                                                                                                                                                      |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hierarchy Level     | [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ],<br>[edit interfaces interface-range <i>name</i> unit <i>logical-unit-number</i> ]                                                                                                                                                   |
| Release Information | Statement introduced in Junos OS Release 9.0 for EX Series switches, including options <b>ethernet-switching</b> , <b>inet</b> , and <b>iso</b> .<br>Option <b>inet6</b> introduced in Junos OS Release 9.3 for EX Series switches.<br>Options <b>ccc</b> introduced in Junos OS Release 9.5 for EX Series switches. |
| Description         | Configure protocol family information for the logical interface on the switch.<br><br>You must configure a logical interface to be able to use the physical device.                                                                                                                                                  |
| Default             | Interfaces on EX4300 switches are set to <b>family ethernet-switching</b> by the default factory configuration. Before you can change the family setting for an interface to another family type, you must delete this default setting or any user-configured family setting.                                        |

**Options** See [Table 145 on page 1649](#) for protocol families available on the switch interfaces. Different protocol families support different subsets of the interface types on the switch. Interface types on the switch are:

- Aggregated Ethernet (**ae0**)
- 40-Gigabit Ethernet (**et**)
- Gigabit Ethernet (**ge**)
- Interface-range configuration (**interface-range**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Integrated Routing and Bridging (IRB) interfaces (IRB) (**irb**)
- Virtual management Ethernet (**vme**)
- 10-Gigabit Ethernet (**xe**)

If you are using an interface range, the supported protocol families are the ones supported by the interface types that compose the range.

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

**Table 299: Protocol Families and Supported Interface Types**

| Family             | Description                                | Supported Interface Types |    |    |     |     |     |     |    |
|--------------------|--------------------------------------------|---------------------------|----|----|-----|-----|-----|-----|----|
|                    |                                            | ae0                       | et | ge | irb | lo0 | me0 | vme | xe |
| ccc                | Circuit cross-connect protocol family      | ✓                         | ✓  | ✓  |     |     |     |     | ✓  |
| ethernet-switching | Ethernet switching protocol family         | ✓                         | ✓  | ✓  |     |     |     |     | ✓  |
| inet               | IPv4 protocol family                       | ✓                         | ✓  | ✓  | ✓   | ✓   | ✓   | ✓   | ✓  |
| inet6              | IPv6 protocol family                       | ✓                         | ✓  | ✓  | ✓   | ✓   | ✓   | ✓   | ✓  |
| iso                | Junos OS protocol family for IS-IS traffic | ✓                         | ✓  | ✓  | ✓   | ✓   | ✓   | ✓   | ✓  |

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring a DHCP Server on Switches (CLI Procedure)*
  - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)
  - [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
  - [Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\) on page 2340](#)

## filter

|                            |                                                                                                                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>filter {   group <i>filter-group-number</i>;   input <i>filter-name</i>;   input-list [ <i>filter-names</i> ];   output <i>filter-name</i>;   output-list [ <i>filter-names</i> ]; }</pre>                                                    |
| <b>Hierarchy Level</b>     | <pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</pre>   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                      |
| <b>Description</b>         | <p>Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure them under the <b>family ethernet-switching</b>, <b>inet</b>, <b>inet6</b>, <b>mpls</b>, or <b>vpls</b> only.</p> |




**NOTE:** On QFX3500 and QFX3600 switches running Enhanced Layer 2 Software, VPLS is not supported.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <p><b>group <i>filter-group-number</i></b>—Define an interface to be part of a filter group.<br/> <b>Range:</b> 1 through 255</p> <p><b>input <i>filter-name</i></b>—Name of one filter to evaluate when packets are received on the interface.</p> <p><b>output <i>filter-name</i></b>—Name of one filter to evaluate when packets are transmitted on the interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                  |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Applying a Filter to an Interface</i></li> <li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> <li>• <i>Junos OS Administration Library for Routing Devices</i></li> <li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li> </ul> |

- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
- [Configuring Firewall Filters and Policers for VPLS](#)
- [family on page 1647](#)
- [family on page 2768](#)

## flow-control

|                                                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                     | (flow-control   no-flow-control);                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                            | [edit interfaces <i>interface-name</i> aggregated-ether-options],<br>[edit interfaces <i>interface-name</i> ether-options],<br>[edit interfaces <i>interface-name</i> fastether-options],<br>[edit interfaces <i>interface-name</i> gigether-options],<br>[edit interfaces <i>interface-name</i> multiservice-options],<br>[edit interfaces interface-range <i>name</i> aggregated-ether-options],<br>[edit interfaces interface-range <i>name</i> ether-options] |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                        | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 in EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.                                                                                                                                                                                                                                                |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                | For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router or switch to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.                                                                                             |
| <div>  <p><b>NOTE:</b> On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                    | Flow control is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                   | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                      | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Flow Control on page 2632</a></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure)</a></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li> </ul>                                                                                                                                                                      |

## force-up

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | force-up;                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ether-options <a href="#">802.3ad lacp</a> ]                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Set the state of the interface as UP when the peer has limited LACP capability.                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 2619</a></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2582</a></li><li>• <i>Junos OS Ethernet Interfaces Configuration Guide</i></li></ul> |

## gratuitous-arp-reply

---

|                                 |                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (gratuitous-arp-reply   no-gratuitous-arp-reply);                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ]                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 in EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. |
| <b>Description</b>              | For Ethernet interfaces, enable updating of the Address Resolution Protocol (ARP) cache for gratuitous ARPs.                                                                                                       |
| <b>Default</b>                  | Updating of the ARP cache is disabled on all Ethernet interfaces.                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Gratuitous ARP on page 2663</a></li><li>• <a href="#">no-gratuitous-arp-request on page 2817</a></li></ul>                                         |



## hash-mode

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>hash-mode {     layer2-header;     layer2-payload; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit forwarding-options <a href="#">enhanced-hash-key</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Select the mode for the hashing algorithm.</p> <p>The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.</p> <p>For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.</p> <p>The hash mode that is set using this statement determines which fields are inspected by the hashing algorithm. You must set the hash mode to <b>layer2-payload</b> if you want the hashing algorithm to inspect fields in the Layer 2 payload when making hashing decisions. You must set the hash mode to <b>layer2-header</b> if you want the hashing algorithm to inspect fields in the Layer 2 header when making hashing decisions.</p> <p>If the hash mode is set to <b>layer2-payload</b>, you can set the fields used by the hashing algorithm to hash IPv4 traffic using the <b>set forwarding-options enhanced-hash-key inet</b> statement. You can set the fields used by the hashing algorithm to hash IPv6 traffic using the <b>set forwarding-options enhanced-hash-key inet6</b> statement.</p> <p>If the hash mode is set to <b>layer2-header</b>, you can set the fields that the hashing algorithm inspects in the Layer 2 header using the <b>set forwarding-options enhanced-hash-key layer2</b> statement.</p> |
| <b>Default</b>                  | layer2-payload                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>layer-2-payload</b>—Set the hashing algorithm to use fields in the Layer 2 payload to make hashing decisions.</p> <p><b>layer-2-header</b>—Set the hashing algorithm to use fields in the Layer 2 header to make hashing decisions.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Related  
Documentation**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2686](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2585](#)
- [enhanced-hash-key on page 2764](#)
- [inet on page 2783](#)
- [inet6 on page 2785](#)
- [layer2 on page 2794](#)

## hold-time (Physical Interface)

|                            |                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>hold-time up <i>milliseconds</i> down <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>     | <code>[edit interfaces <i>interface-name</i>],</code><br><code>[edit interfaces interface-range <i>interface-range-name</i>]</code>                                                                                                                                                                                                                              |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 10.4R5 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.                                                                          |
| <b>Description</b>         | Specify the <b>hold-time</b> value to use to damp shorter interface transitions milliseconds. When an interface goes from up to down, it is not advertised to the rest of the system as being down until it has remained down for the hold-time period. Similarly, an interface is not advertised as being up until it has remained up for the hold-time period. |



### NOTE:

- We recommend that you configure the hold-time value after determining an appropriate value by performing repeated tests in the actual hardware environment. This is because the appropriate value for hold-time depends on the hardware (XFP, SFP, SR, ER, or LR) used in the networking environment.
- The hold-time option is not available for controller interfaces.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | Interface transitions are not damped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>down <i>milliseconds</i></b>—Hold time to use when an interface transitions from up to down. Junos OS advertises the transition within 100 milliseconds of the time value you specify.</p> <p><b>Range:</b> 0 through 4,294,967,295</p> <p><b>Default:</b> 0 (interface transitions are not damped)</p> <p><b>up <i>milliseconds</i></b>—Hold time to use when an interface transitions from down to up. Junos OS advertises the transition within 100 milliseconds of the time value you specify.</p> <p><b>Range:</b> 0 through 4,294,967,295</p> <p><b>Default:</b> 0 (interface transitions are not damped)</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>advertise-interval</i></li> <li>• <i>interfaces (for EX Series switches)</i></li> <li>• <i>Physical Interface Damping Overview</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                 |

- *Damping Shorter Physical Interface Transitions*
- *Damping Longer Physical Interface Transitions*

## iccp

```
Syntax  iccp {
    authentication-key string;
    local-ip-addr local-ip-addr;
    peer ip-address {
        authentication-key string;
        backup-liveness-detection {
            backup-peer-ip ip-address;
        }
        liveness-detection {
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
        local-ip-addr ipv4-address;
        session-establishment-hold-time seconds;
    }
    session-establishment-hold-time seconds;
    traceoptions {
        file <filename> <files number> <match regular-expression> <microsecond-stamp>
        <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 10.0 for MX Series routers.  
Statement introduced in Junos OS Release 12.2 for the QFX Series.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** Configure Interchassis Control Protocol (ICCP) between the multichassis link aggregation group (MC-LAG) peers. ICCP replicates forwarding information, validates configurations, and propagates the operational state of the MC-LAG members.



**NOTE:** Backup liveness detection is not supported on MX Series routers.

The remaining statement are explained separately.

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

## ieee-802-3az-eee

---

**Syntax** ieee-802-3az-eee;

**Hierarchy Level** [edit interfaces *interface-name* ether-options]

**Release Information** Statement introduced in Junos OS Release 12.2 for EX Series switches.

**Description** Configure Energy Efficient Ethernet (EEE) on an EEE-capable Base-T copper interface.

**Default** EEE is disabled on EEE-capable interfaces.

**Required Privilege** system—To view this statement in the configuration.  
**Level** system-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Energy Efficient Ethernet on Interfaces \(CLI Procedure\) on page 2684](#)

## inet (enhanced-hash-key)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>inet {     no-ipv4-destination-address;     no-ipv4-source-address;     no-l4-destination-port;     no-l4-source-port;     no-protocol;     vlan-id; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>     | [edit forwarding-options <a href="#">enhanced-hash-key</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>         | <p>Select the payload fields in IPv4 traffic used by the hashing algorithm to make hashing decisions.</p> <p>When IPv4 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using the <b>inet</b> statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or how to forward traffic to the next hop device when ECMP is enabled.</p> <p>The hashing algorithm, when used to hash LAG bundle traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>The hashing algorithm only inspects the IPv4 fields in the payload to make hashing decisions when the hash mode is set to <b>layer2-payload</b>. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the <b>set forwarding-options enhanced-hash-key hash-mode layer2-payload</b> statement.</p> |
| <b>Default</b>             | <p>The following fields are used by the hashing algorithm to make hashing decisions for IPv4 traffic:</p> <ul style="list-style-type: none"> <li>• IP destination address</li> <li>• IP source address</li> <li>• Layer 4 destination port</li> <li>• Layer 4 source port</li> <li>• Protocol</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>             | <p><b>no-ipv4-destination-address</b>—Exclude the IPv4 destination address field from the hashing algorithm.</p> <p><b>no-ipv4-source-address</b>—Exclude the IPv4 source address field from the hashing algorithm.</p> <p><b>no-l4-destination-port</b>—Exclude the Layer 4 destination port field from the hashing algorithm.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

**no-l4-source-port**—Exclude the Layer 4 source port field from the hashing algorithm.

**no-protocol**—Exclude the protocol field from the hashing algorithm.

**vlan-id**—Include the VLAN ID field in the hashing algorithm.

|                           |                                                               |
|---------------------------|---------------------------------------------------------------|
| <b>Required Privilege</b> | interface—To view this statement in the configuration.        |
| <b>Level</b>              | interface-control—To add this statement to the configuration. |

- |                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2686</a></li><li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2585</a></li><li>• <a href="#">enhanced-hash-key on page 2764</a></li><li>• <a href="#">hash-mode on page 2777</a></li><li>• <a href="#">inet6 on page 2785</a></li></ul> |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|



## inet6 (enhanced-hash-key)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>inet6 {     no-ipv6-destination-address;     no-ipv6-source-address;     no-l4-destination-port;     no-l4-source-port;     no-next-header;     vlan-id; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | [edit forwarding-options <a href="#">enhanced-hash-key</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>         | <p>Select the payload fields in an IPv6 packet used by the hashing algorithm to make hashing decisions.</p> <p>When IPv6 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or to forward traffic to the next hop device when ECMP is enabled.</p> <p>The hashing algorithm, when used to hash LAG traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>The hashing algorithm only inspects the IPv6 fields in the payload to make hashing decisions when the hash mode is set to Layer 2 payload. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the <b>set forwarding-options enhanced-hash-key hash-mode layer2-payload</b> statement.</p> |
| <b>Default</b>             | <p>The data in the following fields are used by the hashing algorithm to make hashing decisions for IPv6 traffic:</p> <ul style="list-style-type: none"> <li>• IP destination address</li> <li>• IP source address</li> <li>• Layer 4 destination port</li> <li>• Layer 4 source port</li> <li>• Next header</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>             | <p><b>no-ipv6-destination-address</b>—Exclude the IPv6 destination address field from the hashing algorithm.</p> <p><b>no-ipv6-source-address</b>—Exclude the IPv6 source address field from the hashing algorithm.</p> <p><b>no-l4-destination-port</b>—Exclude the Layer 4 destination port field from the hashing algorithm.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

**no-l4-source-port**—Exclude the Layer 4 source port field from the hashing algorithm.

**no-next-header**—Exclude the Next Header field from the hashing algorithm.

**vlan-id**—Include the VLAN ID field in the hashing algorithm.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2686](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2585](#)
- [enhanced-hash-key on page 2764](#)
- [hash-mode on page 2777](#)
- [inet on page 2783](#)

---

## interface (Multichassis Protection)

---

**Syntax** interface *interface-name*;



**Hierarchy Level** [edit [multi-chassis multi-chassis-protection](#) peer]

**Release Information** Statement introduced in Junos OS Release 9.6 for MX Series routers.  
Statement introduced in Junos OS Release 12.2 for the QFX Series.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** Specify the name of the interface that is being used as an interchassis link-protection link (ICL-PL). The two switches hosting a multichassis link aggregation group (MC-LAG) use this link to pass Interchassis Control Protocol (ICCP) and data traffic.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

## interface-mode

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | interface-mode (access   trunk);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Hierarchy Level          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Release Information      | Statement introduced in Junos OS Release 9.2.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Description              |  <p><b>NOTE:</b> This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see <i>port-mode</i>. For ELS details, see <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a>.</p> <p>(QFX Series 3500 and 3600 standalone switches)—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the <b>trunk</b> option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the <b>vlan-id</b> or <b>vlan-id-list</b> statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the <b>access</b> option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the <b>vlan-id</b> statement.</p>  <p><b>NOTE:</b> On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure <b>interface-mode</b> and <b>irb</b> for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see <i>Configuring a Trunk Interface on a Bridge Network</i>.</p> |
| Options                  | <p><b>access</b>—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the <b>vlan-id</b> statement.</p> <p><b>trunk</b>—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the <b>vlan-id</b> or <b>vlan-id-list</b> statement.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Related Documentation    | <ul style="list-style-type: none"> <li>• <i>Configuring a Logical Interface for Access Mode</i></li> <li>• <i>Configuring a Logical Interface for Trunk Mode</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

- [Example: Connecting Access Switches to a Distribution Switch on page 2291](#)

## interface-range

```
Syntax  interface-range name {
        accounting-profile name;
        description text;
        disable;
        ether-options {
            802.3ad {
                aex;
                (backup | primary);
                lacp {
                    force-up;
                }
            }
            (auto-negotiation | no-auto-negotiation);
            (flow-control | no-flow-control);
            ieee-802-3az-eee;
            link-mode mode;
            (loopback | no-loopback);
            speed (auto-negotiation | speed);
        }
        (gratuitous-arp-reply | no-gratuitous-arp-reply);
        hold-time up milliseconds down milliseconds;
        member interface-name;
        member-range starting-interface name to ending-interface name;
        mtu bytes;
        no-gratuitous-arp-request;
        traceoptions {
            flag flag;
        }
        (traps | no-traps);
        unit logical-unit-number {
            accounting-profile name;
            bandwidth rate;
            description text;
            disable;
            family family-name {...}
            proxy-arp (restricted | unrestricted);
            (traps | no-traps);
            vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
        }
        vlan-tagging;
    }
```

**Hierarchy Level** [edit interfaces]

**Release Information** Statement introduced in Junos OS Release 10.0 for EX Series switches.

**Description** Group interfaces that share a common configuration profile.



**NOTE:** You can specify interface ranges only for Gigabit and 10-Gigabit Ethernet interfaces.

**Options**    *name*—Name of the interface range.




**NOTE:** You can use regular expressions and wildcards to specify the interfaces in the member configuration. Do not use wildcards for interface types.

The remaining statements are explained separately.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                         interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
  - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)
  - [Understanding Interface Ranges on EX Series Switches](#)
  - [Understanding Interface Ranges on EX Series Switches on page 2598](#)
  - [EX Series Switches Interfaces Overview on page 2577](#)
  - [Junos OS Interfaces Fundamentals Configuration Guide](#)

## lACP (Aggregated Ethernet)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>lACP {   (active   passive);   admin-key key;   accept-data;   fast-failover;   link-protection {     disable;     (revertive   non-revertive);   }   periodic interval;   system-id mac-address;   system-priority priority; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>     | <p>[edit interfaces aeX aggregated-ether-options]</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces aeX aggregated-ether-options]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>fast-failover</b> option introduced in Junos OS Release 12.2.</p> <p>Support for logical systems introduced in Junos OS Release 14.1.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>         | <p>Configure the Link Aggregation Control Protocol (LACP) for aggregated Ethernet interfaces only.</p> <p>When you configure the <b>accept-data</b> statement at the [edit interfaces aeX aggregated-ether-options lACP] hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> When you configure the <b>accept-data</b> statement at the [edit interfaces aeX aggregated-ether-options lACP] hierarchy level, this behavior occurs:</p> <ul style="list-style-type: none"> <li>• By default, the <b>accept-data</b> statement is not configured when LACP is enabled.</li> <li>• You can configure the <b>accept-data</b> statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.</li> <li>• When LACP is down and a member link receives packets, the router or switch does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the <b>accept-data</b> statement is configured.</li> </ul> </div> |
| <b>Default</b>             | If you do not specify LACP as either <b>active</b> or <b>passive</b> , LACP remains passive.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>             | <b>active</b> —Initiate transmission of LACP packets.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

**admin-key *number***—Specify an administrative key for the router or switch.



**NOTE:** You must also configure multichassis link aggregation (MC-LAG) when you configure the **admin-key**.

**fast-failover**—Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.

**passive**—Respond to LACP packets.

The remaining statements are explained separately.

|                                 |                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration. |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Configuring LACP for Aggregated Ethernet Interfaces</i></li><li>• <a href="#">Configuring Aggregated Ethernet LACP (CLI Procedure) on page 2671</a></li><li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li></ul> |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|



## lacp (802.3ad)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>lacp {     force-up;     port-priority }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit interfaces <i>interface-name</i> ether-options <a href="#">802.3ad</a>]</p> <p>[edit interfaces aeX aggregated-ether-options]</p> <p>[edit chassis aggregated-devices ethernet]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Support for LACP link protection introduced in Junos OS Release 11.4 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Configure the Link Aggregation Control Protocol (LACP) parameters for aggregated Ethernet interfaces on the global level (for all the aggregated Ethernet interfaces on the switch) or for a specific aggregated Ethernet interface.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li> <li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li> <li>• <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure) on page 2667</a></li> <li>• <a href="#">Configuring Aggregated Ethernet LACP (CLI Procedure) on page 2671</a></li> <li>• <a href="#">Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 2672</a></li> <li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2582</a></li> <li>• <i>Junos OS Ethernet Interfaces Configuration Guide</i></li> </ul> |

## layer2 (enhanced-hash-key)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>layer2 {<br/>    no-destination-mac-address;<br/>    no-ether-type;<br/>    no-source-mac-address;<br/>    vlan-id;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | [edit forwarding-options <b>enhanced-hash-key</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>         | <p>Select the fields in the Layer 2 header that are used by the hashing algorithm to make hashing decisions.</p> <p>When traffic enters a link aggregation group (LAG) bundle, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links. The hashing algorithm always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>When traffic is exiting a device that has enabled ECMP, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to forward traffic to the next hop device.</p> <p>The hashing algorithm only inspects the fields in the Layer 2 header when the hash mode is set to Layer 2 header. You can set the hash mode to Layer 2 header using the <b>set forwarding-options enhanced-hash-key hash-mode layer2-header</b> statement.</p> |
| <b>Default</b>             | <p>The hash mode of the hashing algorithm is set to Layer 2 payload, by default. When the hash mode is set to Layer 2 payload, the hashing algorithm does not use fields in the Layer 2 header to make hashing decisions.</p> <p>The following fields are used by the hashing algorithm when the hash mode of the hashing algorithm is set to Layer 2 header, by default:</p> <ul style="list-style-type: none"><li>• Destination MAC address</li><li>• Ethertype</li><li>• Source MAC address</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>             | <p><b>no-destination-mac-address</b>—Exclude the destination MAC address field from the hashing algorithm.</p> <p><b>no-ether-type</b>—Exclude the Ethertype field from the hashing algorithm.</p> <p><b>no-source-mac-address</b>—Exclude the source MAC address field from the hashing algorithm.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

**vlan-id**—Include the VLAN ID field in the hashing algorithm.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2686](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2585](#)
- [enhanced-hash-key on page 2764](#)
- [hash-mode on page 2777](#)

## link-mode

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>link-mode mode;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> ],<br>[edit interfaces <i>interface-name</i> ether-options],<br>[edit interfaces <i>ge-pim</i> /0/0 switch-options switch-port <i>port-number</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | Set the device's link connection characteristic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>             | <i>mode</i> —Link characteristics: <ul style="list-style-type: none"><li>• <b>automatic</b>—Link mode is negotiated. This is the default for EX Series switches.</li><li>• <b>full-duplex</b>—Connection is full duplex.</li><li>• <b>half-duplex</b>—Connection is half duplex.</li></ul> <p><b>Default:</b> Fast Ethernet interfaces, except the J Series ePIM Fast Ethernet interfaces, can operate in either full-duplex or half-duplex mode. The router's management Ethernet interface, <b>fxp0</b> or <b>em0</b>, the built-in Fast Ethernet interfaces on the FIC (M7i router), and the Gigabit Ethernet ports on J Series Services Routers with uPIMs installed and configured for access switching mode autonegotiate whether to operate in full-duplex or half-duplex mode. Unless otherwise noted here, all other interfaces operate only in full-duplex mode.</p> |



**NOTE:** On J Series ePIM Fast Ethernet interfaces, if you specify half-duplex (or if full-duplex mode is not autonegotiated), the following message is written to the system log: "Half-duplex mode not supported on this PIC, forcing full-duplex mode."



**NOTE:**

- On EX4300 switches, the interfaces operate in full-duplex mode only.
- On EX Series switches, if no-auto-negotiation is specified in [edit interfaces *interface-name* ether-options], you can select only full-duplex or half-duplex. If auto-negotiation is specified, you can select any mode.



**NOTE:** Member links of an aggregated Ethernet bundle must not be explicitly configured with a link mode. You must remove any such link-mode configuration before committing the aggregated Ethernet configuration.

|                                 |                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Link Characteristics on Ethernet Interfaces</i></li><li>• <i>Understanding Management Ethernet Interfaces</i></li><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li></ul> |

## link-protection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>link-protection {<br/>    disable;<br/>    (revertive  non-revertive);<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <pre>[edit interfaces aex aggregated-ether-options]<br/>[edit interfaces aex aggregated-ether-options lacp]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for <b>disable</b>, <b>revertive</b>, and <b>non-revertive</b> statements added in Junos OS Release 9.3.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the <b>primary</b> and <b>backup</b> statements at the <b>[edit interfaces ge-fpc/pic/port gigheter-options 802.3ad aex]</b> hierarchy level or the <b>[edit interfaces fe-fpc/pic/port fastether-options 802.3ad aex]</b> hierarchy level.</p> <p>On the switch, you can configure either Junos OS link protection for aggregated Ethernet interfaces or the LACP standards link protection for aggregated Ethernet interfaces.</p> <p>For Junos OS link protection, specify <b>link-protection</b> at the following hierarchy levels:</p> <ul style="list-style-type: none"><li>• <b>[edit interfaces ge-fpc/pic/port ether-options 802.3ad aex]</b></li><li>• <b>[edit interfaces xe-fpc/pic/port ether-options 802.3ad aex]</b></li></ul> <p>For LACP standards link protection, specify <b>link-protection</b> at the following hierarchy levels:</p> <ul style="list-style-type: none"><li>• For global LACP link protection, specify at <b>[edit chassis aggregated-devices ethernet lacp]</b></li><li>• For a specific aggregated Ethernet interface, specify at <b>[edit interfaces aeX aggregated-ether-options lacp]</b></li></ul> <p>To disable link protection, use the <b>delete interface ae aggregate-ether-options link-protection</b> statement at the <b>[edit interfaces aex aggregated-ether-options]</b> hierarchy level or the <b>[edit interfaces aex aggregated-ether-options lacp]</b> hierarchy level.</p> |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Aggregated Ethernet Link Protection on page 2676</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\)](#)  
on page 2672

## link-speed (Aggregated Ethernet)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | link-speed <i>speed</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | [edit interfaces aex aggregated-ether-options],<br>[edit interfaces interface-range <i>name</i> aggregated-ether-options],<br>[edit interfaces interface-range <i>name</i> aggregated-sonet-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>         | For aggregated Ethernet interfaces only, set the required link speed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b><i>speed</i></b>—For aggregated Ethernet links, you can specify <b><i>speed</i></b> in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p>Aggregated Ethernet links on the M120 router can have one of the following speeds:</p> <ul style="list-style-type: none"><li>• <b>100m</b>—Links are 100 Mbps.</li><li>• <b>10g</b>—Links are 10 Gbps.</li><li>• <b>1g</b>—Links are 1 Gbps.</li><li>• <b>oc192</b>—Links are OC192 or STM64c.</li></ul> <p>Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:</p> <ul style="list-style-type: none"><li>• <b>10m</b>—Links are 10 Mbps.</li><li>• <b>100m</b>—Links are 100 Mbps.</li><li>• <b>1g</b>—Links are 1 Gbps.</li><li>• <b>10g</b>—Links are 10 Gbps.</li></ul> <p>Aggregated Ethernet links on T Series routers can be configured to operate at one of the following speeds:</p> <ul style="list-style-type: none"><li>• <b>100g</b>—Links are 100 Gbps.</li><li>• <b>100m</b>—Links are 100 Mbps.</li><li>• <b>10g</b>—Links are 10 Gbps.</li><li>• <b>1g</b>—Links are 1 Gbps.</li><li>• <b>40g</b>—Links are 40 Gbps.</li><li>• <b>50g</b>—Links are 50 Gbps.</li><li>• <b>80g</b>—Links are 80 Gbps.</li><li>• <b>8g</b>—Links are 8 Gbps.</li></ul> |



- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Aggregated Ethernet Interfaces Overview*
- [Configuring Aggregated Ethernet Link Speed on page 2678](#)
- *Configuring Mixed Aggregated Ethernet Links*
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

## liveness-detection

**Syntax**

```
liveness-detection {
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}
```

**Hierarchy Level** [edit protocols [iccp peer](#)]

**Release Information** Statement introduced in Junos OS Release 10.0 for MX Series routers.  
Statement introduced in Junos OS Release 12.2 for the QFX Series.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** Enable Bidirectional Forwarding Detection (BFD). BFD enables rapid detection of communication failures between peers.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

## local-bias

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>local-bias;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit interfaces aex aggregated-ether-options]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Enable local link bias for all links in the aggregated Ethernet interface.</p> <p>Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic.</p> <p>You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced as it exits the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Local Link Bias (CLI Procedure) on page 2686</a></li><li>• <a href="#">Understanding Local Link Bias on page 2590</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## local-ip-addr (ICCP)

---

|                                 |                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>local-ip-addr <i>local-ip-address</i>;</code>                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <code>[edit protocols <a href="#">iccp</a>],</code><br><code>[edit protocols <a href="#">iccp</a> <a href="#">peer</a> <i>peer-IP-address</i>]</code>                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches. |
| <b>Description</b>              | Specify the local IP address of the interchassis link (ICL) interface that Interchassis Control Protocol (ICCP) uses to communicate to the peers that host a multichassis link aggregation group (MC-LAG).           |
| <b>Options</b>                  | <i>local-ip-address</i> —Default local IP address to be used by all peers.                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                  |

## loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)

|                            |                                                                                                                                                                                                                                                                                                                       |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | (loopback   no-loopback);                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> aggregated-ether-options],<br>[edit interfaces <i>interface-name</i> ether-options],<br>[edit interfaces <i>interface-name</i> fastether-options],<br>[edit interfaces <i>interface-name</i> gigether-options],<br>[edit interfaces interface-range <i>name</i> ether-options] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.                                                                                                   |
| <b>Description</b>         | For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.                                                                                                                                                                                        |



### NOTE:

- By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system.
- IPv6 Neighbor Discovery Protocol (NDP) addresses are not supported on Gigabit Ethernet interfaces when loopback mode is enabled on the interface. That is, if the loopback statement is configured at the [edit interfaces *ge-fpc/pic/port* gigether-options] hierarchy level, an NDP address cannot be configured at the [edit interfaces *ge-fpc/pic/port* unit *logical-unit-number* family inet6 address] hierarchy level.

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet Loopback Capability on page 2662</a></li> </ul> |

## mc-ae

**Syntax**

```
mc-ae {
  chassis-id chassis-id;
  events {
    iccp-peer-down {
      force-icl-down;
      prefer-status-control-active;
    }
  }
  init-delay-time seconds;
  mc-ae-id mc-ae-id;
  mode (active-active | active-standby);
  redundancy-group group-id;
  revert-time revert-time;
  status-control (active | standby);
  switchover-mode (non-revertive | revertive);
}
```

**Hierarchy Level** [edit interfaces aeX aggregated-ether-options]  
[edit logical-systems *logical-system-name* interfaces aeX aggregated-ether-options]

**Release Information** Statement introduced in Junos OS Release 9.6.  
**events** statement introduced in Junos OS Release 11.4R4.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.  
**prefer-status-control-active** statement introduced in Junos OS Release 13.2R1 for EX Series switches.  
**init-delay-time *seconds*** statement introduced in Junos OS Release 13.2R3 for EX Series switches.  
**switchover-mode** and **revert-time** statements introduced in Junos OS Release 13.3.  
Support for logical systems introduced in Junos OS Release 14.1.

**Description** Enable multichassis link aggregation groups (MC-LAG), which enables one device to form a logical LAG interface with two or more other devices.

**Options** **chassis-id**—Specify the chassis ID for Link Aggregation Control Protocol (LACP) to calculate the port number of MC-LAG physical member links.

**Values:** 0 or 1

**events**—Specify an action if a specific MC-LAG event occurs.

**iccp-peer-down**—Specify an action if the ICCP peer of this node goes down.

**force-icl-down**—If the node's ICCP peer goes down, bring down the interchassis-link logical interface.

**prefer-status-control-active**—Specify that the node configured as **status-control active** become the active node if the peer of this node goes down.



**NOTE:** The **prefer-status-control-active** statement can be configured with the **status-control standby** configuration to prevent

the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure. Use this configuration only if you can ensure that ICCP will not go down unless the router or switch is down. You must also configure the **hold-time down** value (at the **[edit interfaces *interface-name*]** hierarchy level) for the interchassis link with the **status-control standby** configuration to be higher than the ICCP BFD timeout. This configuration prevents data traffic loss by ensuring that when the router or switch with the **status-control active** configuration goes down, the router or switch with the **status-control standby** configuration does not go into standby mode.

To make the **prefer-status-control-active** configuration work with the **status-control standby** configuration when an interchassis-link logical interface is configured on aggregate Ethernet interface, you must either configure the **lacp periodic *interval*** statement at the **[edit interface *interface-name* aggregated-ether-options]** hierarchy level as **slow** or configure the **detection-time threshold** statement at the **[edit protocols iccp peer liveness-detection]** hierarchy level as less than 3 seconds.

**mc-ae-id**—Specify the identification number of the MC-LAG device. The two MC-LAG network devices that manage a given MC-LAG must have the same identification number.

**Range:** 1 through 65,535

**mode (active-active | active-standby)**—Specify whether the MC-LAG is in active-active or active-standby mode.



**NOTE:** You can configure IPv4 (**inet**) and IPv6 (**inet6**) addresses on **mc-ae** interfaces when the **active-standby** mode is configured.

**redundancy-group**—Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions.

**Range:** 1 through 4,294,967,294

**revert-time**—Wait interval (in minutes) before the switchover to the preferred node is performed when the **switchover-mode** is configured as **revertive**.

**Range:** 1 through 10

**status-control (active | standby)**—Specify whether the chassis becomes active or remains in standby when an interchassis link failure occurs.

**switchover-mode (non-revertive | revertive)**—Specify whether Junos OS should trigger a link switchover to the preferred node when the active node is available.



**NOTE:** For revertive mode to automatically switch over to the preferred node, the **status-control** statement should be configured as **active**.

**init-delay-time seconds**—To minimize traffic loss, specify the number of seconds by which to delay bringing the multichassis aggregated Ethernet (mc-ae) interface back to the up state when you reboot an MC-LAG peer.

The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview</i></li> <li>• <i>Configuring Multichassis Link Aggregation</i></li> <li>• <a href="#">Configuring Multichassis Link Aggregation on page 2681</a></li> <li>• <i>Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers</i></li> <li>• <i>Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain</i></li> <li>• <i>Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces</i></li> </ul> |

## mc-ae-id

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mc-ae-id mc-ae-id;</code>                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit interfaces aggregated-ether-options mc-ae]                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.                                                                                                                                        |
| <b>Description</b>              | Specify the multichassis aggregated Ethernet (MC-AE) identification number of the MC-AE that a given aggregated Ethernet interface belongs to. The two peers that host a given multichassis link aggregation group (MC-LAG) must have the same multichassis aggregated Ethernet ID. |
| <b>Options</b>                  | <b>Range:</b> 1 through 65535.                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                             |

## member (Interface Ranges)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>member <i>interface-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit interfaces <a href="#">interface-range</a> <i>interface-range-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Specify the name of the member interface belonging to an interface range on the EX Series switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure)</a></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li> <li>• <a href="#">Understanding Interface Ranges on EX Series Switches</a></li> <li>• <a href="#">Understanding Interface Ranges on EX Series Switches on page 2598</a></li> <li>• <a href="#">EX Series Switches Interfaces Overview on page 2577</a></li> <li>• <a href="#">Junos OS Interfaces Fundamentals Configuration Guide</a></li> </ul> |

## member-range

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>member-range <i>starting-interface-name</i> to <i>ending-interface-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit interfaces <a href="#">interface-range</a> <i>interface-range-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Specify the names of the first and last members of a sequence of interfaces belonging to an interface range.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>Range:</b> <i>Starting interface-name</i> to <i>ending interface-name</i> —The name of the first member and the name of the last member in the interface sequence.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure)</a></li><li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li><li>• <a href="#">Understanding Interface Ranges on EX Series Switches</a></li><li>• <a href="#">Understanding Interface Ranges on EX Series Switches on page 2598</a></li><li>• <a href="#">EX Series Switches Interfaces Overview on page 2577</a></li><li>• <a href="#">Junos OS Interfaces Fundamentals Configuration Guide</a></li></ul> |



## members

|                            |                                                                                                                                                                                                  |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>members [(all   <i>names</i>   <i>vlan-ids</i>)];</code>                                                                                                                                   |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>family</b> ethernet-switching <b>vlan</b> ]                                                                     |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated with enhanced ? (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches. |
| <b>Description</b>         | For trunk interfaces, configure the VLANs that can carry traffic.                                                                                                                                |



**TIP:** To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after `vlan` or `vlands` in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.



**NOTE:** The number of VLANs supported per switch varies for each model. Use the configuration-mode command `set vlans id vlan-id ?` to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum.

On an EX Series switch that runs Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style, the maximum number of VLAN members allowed on the switch is 8 times the maximum number of VLANs the switch supports (`vmember limit = vlan max * 8`). If the switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (`eswd`) due to memory allocation failure.

On an EX Series switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs the switch supports (`vmember limit = vlan max * 24`). If the configuration of one of these switches exceeds the recommended VLAN member maximum, a warning message appears in the system log (`syslog`).

**Options** `all`—Specifies that this trunk interface is a member of all the VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.



**NOTE:** Since VLAN members are limited, specifying all could cause the number of VLAN members to exceed the limit at some point.

**names**—Name of one or more VLANs. VLAN IDs are applied automatically in this case.



**NOTE:** **all** cannot be a VLAN name.

**vlan-ids**—Numeric identifier of one or more VLANs. For a series of tagged VLANs, specify a range; for example, 10–20 or 10–20 23 27–30.



**NOTE:** Each configured VLAN must have a specified VLAN ID to successfully commit the configuration; otherwise, the configuration commit fails.

**Required Privilege  
Level**

interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related  
Documentation**

- *show ethernet-switching interfaces*
- [show ethernet-switching interface on page 2442](#)
- *show vlans*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- [Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch on page 2281](#)
- *Example: Connecting an Access Switch to a Distribution Switch*
- [Example: Connecting Access Switches to a Distribution Switch on page 2291](#)
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 2619](#)
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- [Configuring VLANs for EX Series Switches \(CLI Procedure\) on page 2337](#)
- *Creating a Series of Tagged VLANs (CLI Procedure)*
- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

## minimum-interval (Liveness Detection)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-interval <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <code>iccp peer liveness-detection</code> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure simultaneously the minimum interval at which the peer transmits liveness detection requests and the minimum interval at which the peer expects to receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately by using the <code>transmit-interval</code> <code>minimal-interval</code> and <code>minimum-receive-interval</code> statements, respectively. |
| <b>Options</b>                  | <i>milliseconds</i> —Specify the minimum interval value for Bidirectional Forwarding Detection (BFD).<br><b>Range:</b> 1 through 255,000                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                        |

## minimum-receive-interval (Liveness Detection)

|                                 |                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-receive-interval <i>milliseconds</i>;</code>                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <code>iccp peer liveness-detection</code> ]                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches. |
| <b>Description</b>              | Configure the minimum interval at which the peer must receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session.                                                  |
| <b>Options</b>                  | <i>milliseconds</i> —Specify the minimum interval value.<br><b>Range:</b> 1 through 255,000                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                  |

## mtu

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>mtu bytes;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>     | <pre> [edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit interfaces <i>interface-range name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>] [edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols l2vpn interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols l2vpn site <i>site-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls] </pre> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Layer 2 VPNs and VPLS introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Support at the <code>[set interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>]</code> hierarchy level introduced in Junos OS Release 12.3R3 for MX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | <p>Specify the maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.</p> <p>To route jumbo data packets on an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI) on EX Series switches, you must configure the jumbo MTU size on the member physical interfaces of the VLAN that you have associated with the IRB interface or RVI, as well as on the IRB interface or RVI itself (the interface named <code>irb</code> or <code>vlan</code>, respectively).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |



**CAUTION:** For EX Series switches, setting or deleting the jumbo MTU size on an IRB interface or RVI while the switch is transmitting packets might cause packets to be dropped.



## NOTE:

The MTU for an IRB interface is calculated by removing the Ethernet header overhead [6(DMAC)+6(SMAC)+2(EtherType)]. Because, the MTU is the lower value of the MTU configured on the IRB interface and the MTU configured on the IRB's associated bridge domain IFDs or IFLs, the IRB MTU is calculated as follows:

- In case of Layer 2 IFL configured with the `flexible-vlan-tagging` statement, the IRB MTU is calculated by including 8 bytes overhead (SVLAN+CVLAN).
- In case of Layer 2 IFL configured with the `vlan-tagging` statement, the IRB MTU is calculated by including a single VLAN 4 bytes overhead.



## NOTE:

- If a packet whose size is larger than the configured MTU size is received on the receiving interface, the packet is eventually dropped. The value considered for MRU (maximum receive unit) size is also the same as the MTU size configured on that interface.
- Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for management Ethernet interfaces (fxp0, em0, or me0) or for loopback, multilink, and multicast tunnel devices.
- On ACX Series routers, you can configure the protocol MTU by including the `mtu` statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet] or [edit interfaces *interface-name* unit *logical-unit-number* family inet6] hierarchy level.
  - If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.
  - If you are configuring the protocol MTU for both inet and inet6 families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for inet and inet6 families that are configured on the same logical interface.

For more information about configuring MTU for specific interfaces and router or switch combinations, see [“Configuring the Media MTU” on page 2638](#).

|                                 |                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <b>bytes</b> —MTU size.<br><b>Range:</b> 256 through 9192 bytes, 256 through 9216 (EX Series switch interfaces), 256 through 9500 bytes (Junos OS 12.1X48R2 for PTX Series routers)<br><b>Default:</b> 1500 bytes (INET, INET6, and ISO families), 1448 bytes (MPLS), 1514 bytes (EX Series switch interfaces) |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring the Media MTU on page 2638</a></li><li>• <i>Configuring the MTU for Layer 2 Interfaces</i></li><li>• <a href="#">Setting the Protocol MTU on page 2651</a></li></ul>  |

---

## multi-chassis

---

|                                 |                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>multi-chassis {<br/>    multi-chassis-protection peer-ip-address {<br/>        interface interface-name;<br/>    }<br/>}</pre>                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit]                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.                                   |
| <b>Description</b>              | Configure an interchassis link-protection link (ICL-PL) between the two peers that host a multichassis link aggregation group (MC-LAG). You can configure either an aggregated Ethernet interface or a 10-Gigabit Ethernet interface to be an ICL-PL. |
| <b>Options</b>                  | <b>interface interface-name</b> —Specify the logical interface name of the peer.<br><br>The remaining statements are explained separately.                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                               |

## multi-chassis-protection

---

**Syntax**    `multi-chassis-protection peer-ip-address {  
                  interface interface-name;  
                  }`

**Hierarchy Level**    [edit [multi-chassis](#)]

**Release Information**    Statement introduced in Junos OS Release 9.6 for MX Series routers.  
Statement introduced in Junos OS Release 12.2 for the QFX Series.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description**    Configure multichassis link protection between the two peers that host a multichassis link aggregation group (MC-LAG). If the Interchassis Control Protocol (ICCP) connection is up and the interchassis link (ICL) comes up, the peer configured as standby brings up the multichassis aggregated Ethernet (MC-AE) interfaces shared with the peer. Multichassis protection must be configured on one interface for each peer.

The remaining statements are explained separately.

**Options**    **interface** *interface-name*—Specify the logical interface name of the peer.

The remaining statements are explained separately.

**Required Privilege Level**    **interface**—To view this statement in the configuration.  
                  **interface-control**—To add this statement to the configuration.

## native-vlan-id

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>native-vlan-id <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit interfaces <i>ge-fpc/pic/port</i> ],<br>[edit interfaces <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure mixed tagging support for untagged packets on a port for the following:</p> <ul style="list-style-type: none"><li>• M Series routers with Gigabit Ethernet IQ PICs with SFP and Gigabit Ethernet IQ2 PICs with SFP configured for 802.1Q flexible VLAN tagging</li><li>• MX Series routers with Gigabit Ethernet DPCs and MICs, Tri-Rate Ethernet DPCs and MICs, and 10-Gigabit Ethernet DPCs and MICs and MPCs configured for 802.1Q flexible VLAN tagging</li><li>• T4000 routers with 100-Gigabit Ethernet Type 5 PIC with CFP</li><li>• EX Series switches with Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces</li></ul> <p>When the <b>native-vlan-id</b> statement is included with the <a href="#">flexible-vlan-tagging</a> statement, untagged packets are accepted on the same mixed VLAN-tagged port.</p> <p>The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface. To configure the logical interface, include the <b>vlan-id</b> statement (matching the <b>native-vlan-id</b> statement on the physical interface) at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p> <p>When the <b>native-vlan-id</b> statement is included with the <a href="#">interface-mode</a> statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.</p> |
| <b>Options</b>                  | <p><b><i>number</i></b>—VLAN ID number.</p> <p><b>Range:</b> (ACX Series routers and EX Series switches) 0 through 4094.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Mixed Tagging Support for Untagged Packets</a></li><li>• <a href="#">Configuring a Logical Interface for Access Mode</a></li><li>• <a href="#">Configuring the Native VLAN Identifier (CLI Procedure) on page 2342</a></li><li>• <a href="#">Understanding Bridging and VLANs on EX Series Switches on page 2245</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



- [flexible-vlan-tagging on page 2387](#)
- [Understanding Q-in-Q Tunneling on EX Series Switches on page 2269](#)

## no-gratuitous-arp-request

---

|                                 |                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-gratuitous-arp-request;                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ]                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.           |
| <b>Description</b>              | For Ethernet interfaces and pseudowire logical interfaces, do not respond to gratuitous ARP requests.                                                                    |
| <b>Default</b>                  | Gratuitous ARP responses are enabled on all Ethernet interfaces.                                                                                                         |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gratuitous ARP on page 2663</a></li> <li>• <a href="#">gratuitous-arp-reply on page 2776</a></li> </ul> |

## no-redirects

---

|                                 |                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-redirects;                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> ]                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                    |
| <b>Description</b>              | Do not send protocol redirect messages on the interface.<br><br>To disable the sending of protocol redirect messages for the entire router or switch, include the <b>no-redirects</b> statement at the <b>[edit system]</b> hierarchy level. |
| <b>Default</b>                  | Interfaces send protocol redirect messages.                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Disabling the Transmission of Redirect Messages on an Interface on page 2665</a></li> <li>• <i>Junos OS Administration Library for Routing Devices</i></li> </ul>                       |

## peer (ICCP)

---

**Syntax** `peer ip-address {  
    authentication-key string;  
    backup-liveness-detection {  
        backup-peer-ip ip-address;  
    }  
    liveness-detection {  
        detection-time {  
            threshold milliseconds;  
        }  
        minimum-interval milliseconds;  
        minimum-receive-interval milliseconds;  
        multiplier number;  
        no-adaptation;  
        transmit-interval {  
            minimum-interval milliseconds;  
            threshold milliseconds;  
        }  
        version (1 | automatic);  
    }  
    local-ip-addr ipv4-address;  
    session-establishment-hold-time seconds;  
}`

**Hierarchy Level** [edit protocols [iccp](#)]

**Release Information** Statement introduced in Junos OS Release 10.0 for MX Series routers.  
Statement introduced in Junos OS Release 12.2 for the QFX Series.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** Configure the peers that host a multichassis link aggregation group (MC-LAG). You must configure Interchassis Control Protocol (ICCP) for both peers that host the MC-LAG.



**NOTE:** Backup liveness detection is not supported on MX Series routers.

---

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.


## periodic

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>periodic interval;</code>                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit interfaces aex aggregated-ether-options <a href="#">lcp</a> ],<br>[edit interfaces interface-range <i>name</i> aggregated-ether-options <a href="#">lcp</a> ]                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | For aggregated Ethernet interfaces only, configure the interval for periodic transmission of LACP packets.                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><i>interval</i>—Interval for periodic transmission of LACP packets.</p> <ul style="list-style-type: none"> <li><b>fast</b>—Transmit packets every second.</li> <li><b>slow</b>—Transmit packets every 30 seconds.</li> </ul> <p><b>Default:</b> <b>fast</b></p>                                                                                                                    |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring LACP for Aggregated Ethernet Interfaces</i></li> <li><a href="#">Configuring Aggregated Ethernet LACP (CLI Procedure) on page 2671</a></li> <li><i>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li> </ul> |

## preferred

---

|                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                             |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                     | preferred;                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>                                                                                                                                                            | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i><br>family <i>family</i> address <i>address</i> ] |
| <b>Release Information</b>                                                                                                                                                        | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                              |
| <b>Description</b>                                                                                                                                                                | Configure this address to be the preferred address on the interface. If you configure more than one address on the same subnet, the preferred source address is chosen by default as the source address when you initiate frame transfers to destinations on the subnet.                    |
| <div> <b>NOTE:</b> The edit logical-systems hierarchy is not available on QFabric systems.</div> |                                                                                                                                                                                                                                                                                             |
| <b>Default</b>                                                                                                                                                                    | The lowest-numbered address on the subnet is the preferred address.                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                                   | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                      | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Interface Address on page 2632</a></li></ul>                                                                                                                                                                            |

## primary (Address on Interface)

|                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                             |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                       | primary;                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>                                                                                                                                                              | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> ]                    |
| <b>Release Information</b>                                                                                                                                                          | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                  | Configure this address to be the primary address of the protocol on the interface. If the logical unit has more than one address, the primary address is used by default as the source address when packet transfer originates from the interface and the destination address does not indicate the subnet. |
| <div>  <b>NOTE:</b> The edit logical-systems hierarchy is not available on QFabric systems. </div> |                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>                                                                                                                                                                      | For unicast traffic, the primary address is the lowest non-127 (in other words, non-loopback) preferred address on the unit.                                                                                                                                                                                |
| <b>Required Privilege Level</b>                                                                                                                                                     | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                        | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Interface Address on page 2632</a></li> </ul>                                                                                                                                                                                          |

## proxy-arp

---

|                            |                                                                                                                                                                                                                                                                            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>proxy-arp (restricted   unrestricted);</code>                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ]                                                                           |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.6 for EX Series switches.<br><b>restricted</b> added in Junos OS Release 10.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series. |
| <b>Description</b>         | For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.                                                                                       |



**NOTE:** You must configure the IP address and the `inet` family for the interface when you enable proxy ARP.

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <ul style="list-style-type: none"><li>• <b>none</b>—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.</li><li>• <b>restricted</b>—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address.</li><li>• <b>unrestricted</b>—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.</li></ul> |
|                                 | <b>Default:</b> <code>unrestricted</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Restricted and Unrestricted Proxy ARP on page 2665</a></li><li>• <a href="#">Configuring Proxy ARP (CLI Procedure)</a></li><li>• <a href="#">Configuring Proxy ARP (CLI Procedure) on page 2360</a></li><li>• <a href="#">Example: Configuring Proxy ARP on an EX Series Switch on page 2331</a></li><li>• <a href="#">Configuring Gratuitous ARP on page 2663</a></li></ul>                                                                                                                                                                                                                                                                         |

## rpf-check

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rpf-check;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6]                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>On EX3200 and EX4200 switches, enable a reverse-path forwarding (RPF) check on unicast traffic (except ECMP packets) on all ingress interfaces.</p> <p>On EX4300 switches, enable a reverse-path forwarding (RPF) check on unicast traffic, including ECMP packets, on all ingress interfaces.</p> <p>On EX8200 and EX6200 switches, enable an RPF check on unicast traffic, including ECMP packets, on the selected ingress interfaces.</p> <p>On QFX Series switches, enable an RPF check on unicast traffic (except ECMP packets) on the selected ingress interfaces.</p> |
| <b>Default</b>                  | Unicast RPF is disabled on all interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Unicast RPF on an EX Series Switch</i></li> <li>• <a href="#">Configuring Unicast RPF (CLI Procedure) on page 2689</a></li> <li>• <a href="#">Disabling Unicast RPF (CLI Procedure) on page 2691</a></li> <li>• <a href="#">Understanding Unicast RPF on page 2592</a></li> </ul>                                                                                                                                                                                                                              |

## session-establishment-hold-time

---

|                                 |                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>session-establishment-hold-time <i>seconds</i>;</code>                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer</a> ],<br>[edit protocols <a href="#">iccp</a> ]                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches. |
| <b>Description</b>              | Specify the time during which an Interchassis Control Protocol (ICCP) connection must be established between peers.                                                                                                  |
| <b>Options</b>                  | <b>seconds</b> —Time (in seconds) within which a successful ICCP connection must be established.                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                  |



## speed (Ethernet)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>speed (10m   100m   1g   auto   auto-10m-100m);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ],<br>[edit interfaces <i>ge-pim/0/0</i> switch-options switch-port <i>port-number</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Configure the interface speed. This statement applies to the management Ethernet interface (<b>fxp0</b> or <b>em0</b>), Fast Ethernet 12-port and 48-port PICs, the built-in Fast Ethernet port on the FIC (M7i router), the built-in Ethernet interfaces on J Series Services Routers, Combo Line Rate DPCs and Tri-Rate Ethernet Copper interfaces on MX Series routers, Gigabit Ethernet ports on J Series Services Routers with uPIMs installed and configured for access switching mode, and Gigabit Ethernet interfaces on EX Series switches.</p> <p>When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled. When you configure 100BASE-FX SFP, you must set the port speed at 100 Mbps.</p>                       |
| <b>Options</b>                  | <p>You can specify the speed as either <b>10m</b> (10 Mbps), <b>100m</b> (100 Mbps), or on J Series routers with uPIMs installed and on MX Series routers, <b>1g</b> (1 Gbps). You can also specify the <b>auto</b> option on MX Series routers.</p> <p>For Gigabit Ethernet interfaces on EX Series switches, you can specify one of the following options:</p> <ul style="list-style-type: none"> <li>• <b>10m</b>—10 Mbps</li> <li>• <b>100m</b>—100 Mbps</li> <li>• <b>1g</b>—1 Gbps</li> <li>• <b>auto</b>—Automatically negotiate the speed (10 Mbps, 100 Mbps, or 1 Gbps) based on the speed of the other end of the link.</li> <li>• <b>auto-10m-100m</b>—Automatically negotiate the speed (10 Mbps or 100 Mbps) based on the speed of the other end of the link.</li> </ul> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Interface Speed</i></li> <li>• <i>Configuring the Interface Speed on Ethernet Interfaces</i></li> <li>• <i>Configuring Gigabit Ethernet Autonegotiation</i></li> <li>• <i>Configuring Gigabit Ethernet Interfaces on J Series Services Routers</i></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                             |

## traceoptions (Individual Interfaces)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>name</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i>;<br/>    match;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>         | <p>Define tracing operations for individual interfaces.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>The interfaces <b>traceoptions</b> statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system <b>syslog</b> file in the directory <b>/var/log/dcd</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>             | If you do not include this statement, no interface-specific tracing operations are performed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>             | <p><b>file name</b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log/dcd</b>. By default, interface process tracing output is placed in the file <b>files number</b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>match</b>—(Optional) Regular expression for lines to be traced.</p> <p><b>no-world-readable</b>—(Optional) Prevent any user from reading the log file.</p> <p><b>world-readable</b>—(Optional) Allow any user to read the log file.</p> <p><b>size size</b>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <b>trace-file</b> reaches this size, it is renamed <b>trace-file.0</b>. When the <b>trace-file</b> again reaches its maximum size, <b>trace-file.0</b> is renamed <b>trace-file.1</b> and <b>trace-file</b> is renamed <b>trace-file.0</b>. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none"><li>• <b>all</b>—All interface tracing operations</li><li>• <b>event</b>—Interface events</li><li>• <b>ipc</b>—Interface interprocess communication (IPC) messages</li></ul> |

- **media**—Interface media changes
- **q921**—Trace ISDN Q.921 frames
- **q931**—Trace ISDN Q.931 frames

**Required Privilege** interface—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

**Related Documentation**

- *Tracing Operations of an Individual Router Interface*

## traceoptions (Interface Process)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file &lt;filename&gt; &lt;files number&gt; &lt;match regular-expression&gt; &lt;size size&gt; &lt;world-readable  <br/>        no-world-readable&gt;;<br/>    flag flag &lt;disable&gt;;<br/>    no-remote-trace;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>     | [edit interfaces]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>         | Define tracing operations for the interface process (dcd).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>             | If you do not include this statement, no interface-specific tracing operations are performed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>filename</b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. By default, interface process tracing output is placed in the file <b>dcd</b>.</p> <p><b>files number</b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b></li><li>• <b>change-events</b>—Log changes that produce configuration events</li><li>• <b>config-states</b>—Log the configuration state machine changes</li><li>• <b>kernel</b>—Log configuration IPC messages to kernel</li><li>• <b>kernel-detail</b>—Log details of configuration messages to kernel</li></ul> <p><b>no-world-readable</b>—(Optional) Disallow any user to read the log file.</p> |

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify kilobytes, **xm** to specify megabytes, or **xg** to specify gigabytes

**Range:** 10 KB through the maximum file size supported on your router

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

**match regex**—(Optional) Refine the output to include only those lines that match the given regular expression.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Tracing Operations of the Interface Process*

## transmit-interval (Liveness Detection)

**Syntax** transmit-interval {  
    **minimum-interval** *milliseconds*;  
    **threshold** *milliseconds*;  
}

**Hierarchy Level** [edit protocols **iccp** **peer** **liveness-detection**]

**Release Information** Statement introduced in Junos OS Release 10.0 for MX Series routers.  
Statement introduced in Junos OS Release 12.2 for the QFX Series.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** Configure the Bidirectional Forwarding Detection (BFD) transmit interval. The negotiated transmit interval for a peer is the interval between the sending of BFD liveness detection requests to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

## traps

---

|                                 |                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (traps   no-traps);                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ],<br>[edit interfaces interface-range <i>name</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.                                                                              |
| <b>Description</b>              | Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.                                                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling or Disabling SNMP Notifications on Physical Interfaces</i></li><li>• <a href="#">Enabling or Disabling SNMP Notifications on Logical Interfaces on page 2666</a></li></ul>                                                                   |

## unit

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> unit <i>logical-unit-number</i> {   <i>accounting-profile name</i>;   <i>bandwidth rate</i>;   <i>description text</i>;   <i>disable</i>;   <i>family family-name</i> {...}   <i>proxy-arp</i> (restricted   unrestricted);   (<i>traps</i>   no-traps);   <i>vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number</i>; }</pre>                                                                                                                                     |
| <b>Hierarchy Level</b>          | <pre> [edit interfaces <i>interface-name</i>], [edit interfaces interface-range <i>name</i>]</pre>                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b><i>logical-unit-number</i></b>—Number of the logical unit.</p> <p><b>Range:</b> 0 through 16,384</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure)</a></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li> <li>• <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure) on page 2667</a></li> <li>• <a href="#">EX Series Switches Interfaces Overview on page 2577</a></li> <li>• <a href="#">Junos OS Ethernet Interfaces Configuration Guide</a></li> </ul> |

## vlan (802.1Q Tagging)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vlan {<br/>    members [(all   names   vlan-ids)];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>family</b> ethernet-switching]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Bind an 802.1Q VLAN tag ID to a logical interface.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>show ethernet-switching interfaces</i></li><li>• <a href="#">show ethernet-switching interface on page 2442</a></li><li>• <i>Example: Setting Up Bridging with Multiple VLANs for EX Series Switches</i></li><li>• <i>Configuring Routed VLAN Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Integrated Routing and Bridging Interfaces (CLI Procedure) on page 2340</a></li><li>• <a href="#">Understanding Bridging and VLANs on EX Series Switches on page 2245</a></li><li>• <i>Junos OS Ethernet Interfaces Configuration Guide</i></li></ul> |



## vlan-id (VLAN Tagging and Layer 3 Subinterfaces)

|                            |                                                                                 |
|----------------------------|---------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>vlan-id <i>vlan-id-number</i>;</code>                                     |
| <b>Hierarchy Level</b>     | [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.2 for EX Series switches.            |
| <b>Description</b>         | Bind an 802.1Q VLAN tag ID to a logical interface.                              |



**NOTE:** The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <p><i>vlan-id-number</i>—A valid VLAN identifier.</p> <p><b>Range:</b> 1 through 4094</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">vlan-tagging on page 2834</a></li> <li>• <i>Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch</i></li> <li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 2616</a></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 2619</a></li> <li>• <a href="#">Configuring a Layer 3 Subinterface (CLI Procedure) on page 2689</a></li> <li>• <a href="#">Configuring Q-in-Q Tunneling (CLI Procedure) on page 2351</a></li> <li>• <i>Junos OS Ethernet Interfaces Configuration Guide</i></li> </ul> |

## vlan-tagging

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | vlan-tagging;                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> ]                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Routers.<br>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series. |
| <b>Description</b>              | For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.                                                                                                                                        |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring BGP Autodiscovery for LDP VPLS</i></li><li>• <a href="#">Configuring Tagged Aggregated Ethernet Interfaces on page 2688</a></li><li>• <i>Configuring Interfaces for VPLS Routing</i></li><li>• <i>Enabling VLAN Tagging</i></li><li>• <a href="#">802.1Q VLANs Overview on page 2614</a></li></ul>       |

# Administration

- Routine Monitoring on page 2835
- Operational Commands on page 2844

## Routine Monitoring

---

- Monitoring Interface Status and Traffic on page 2835
- Verifying the Status of a LAG Interface on page 2837
- Verifying That EEE Is Saving Energy on Configured Ports on page 2837
- Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2839
- Verifying That Layer 3 Subinterfaces Are Working on page 2840
- Verifying Unicast RPF Status on page 2841
- Verifying IP Directed Broadcast Status on page 2843

## Monitoring Interface Status and Traffic

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to view interface status or to monitor interface bandwidth utilization and traffic statistics on the EX Series switches.

The J-Web interface monitors interface bandwidth utilization and plots real-time charts to display input and output rates in bytes per second. In addition, the Interface monitoring page displays input and output packet counters and error counters in the form of charts.

Alternatively, you can enter the **show** commands in the CLI to view interface status and traffic statistics.



**NOTE:** For logical interfaces on EX Series switches, the traffic statistics fields in **show interfaces** commands show only control traffic; the traffic statistics do not include data traffic.



**NOTE:** EX Series switches do not support the collection and reporting of IPv6 transit statistics. Therefore, the IPv6 transit statistics field in the `show interfaces` commands displays all values as 0.

**Action** To view general interface information in the J-Web interface such as available interfaces, select **Monitor > Interfaces**. Click any interface to view details about its status.

To set up interface monitoring for Virtual Chassis and EX8200 switches, select a member from the **Port for Member** list. Details such as the admin status and link status are displayed in the table. For an EX8200 Virtual Chassis setup, select the member, **FPC**, and the required interface.



**NOTE:** By default, the details of the first member in the FPC list is displayed. In an EX8200 Virtual Chassis setup, details of the first member and the first FPC is displayed.

You have the following options:

- **Start/Stop**—Starts or stops monitoring the selected interface.
- **Show Graph**—Displays input and output packet counters and error counters in the form of charts. Click the pop-up icon to view the graph in a separate window.
- **Details**—Displays interface information such as general details, traffic statistics, I/O errors, CoS counters, and Ethernet statistics.
- **Refresh Interval (sec)**—Displays the time interval you have set for page refresh.
- **Clear Statistics**—Clears the statistics for the interface selected from the table.

Using the CLI:

- To view interface status for all the interfaces, enter `show interfaces xe-`.
- To view status and statistics for a specific interface, enter `show interfaces xe-interface-name`.
- To view status and traffic statistics for all interfaces, enter either `show interfaces xe-detail` or `show interfaces xe- extensive`.

**Meaning** In the J-Web interface the charts displayed are:

- Bar charts—Display the input and output error counters.
- Pie charts—Display the number of broadcast, unicast, and multicast packet counters.

For details about output from the CLI commands, see `show interfaces ge-` (Gigabit Ethernet) or `show interfaces xe-` (10-Gigabit Ethernet).

- Related Documentation**
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 2619](#)
  - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
  - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)

## Verifying the Status of a LAG Interface

**Purpose** Verify that a LAG (ae0) has been created on the switch.

**Action** Enter the following command:

```
user@switch> show interfaces ae0 terse
Interface      Admin  Link  Proto  Local      Remote
ae0            up     up
ae0.0          up     up     inet   10.10.10.2/24
```

**Meaning** The output confirms that the ae0 link is up and shows the family and IP address assigned to this link.

- Related Documentation**
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 2667](#)
  - [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\) on page 2668](#)
  - *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

## Verifying That EEE Is Saving Energy on Configured Ports

**Purpose** Verify that enabling EEE saves energy on Base-T Copper Ethernet ports.

**Action** You can see the amount of energy saved by EEE on an EX Series switch using the [show chassis power-budget-statistics](#) command.

1. View the power budget of an EX Series switch before enabling EEE.

- On an EX6210 switch:

```
user@switch> show chassis power-budget-statistics
PSU 2 (EX6200-PWR-AC2500) : 2500 W Online
PSU 3 ) : 0 W Offline
Total Power supplied by all Online PSUs : 2500 W
Power Redundancy Configuration : N+1
Power Reserved for the Chassis : 500 W

Fan Tray Statistics
FTC 0 : Base power 300 W Power Used nan W
FPC Statistics
power Priority Base power Power Used PoE
FPC 3 (EX6200-48T) : 150 W 61.54 W
0 W 9
FPC 4 (EX6200-SRE64-4XS) : 100 W 48.25 W
0 W 0
FPC 5 (EX6200-SRE64-4XS) : 100 W 48.00 W
0 W 0
FPC 7 (EX6200-48T) : 150 W 63.11 W
0 W 9
```

```

FPC 8 (EX6200-48T) : 150 W 12.17 W
0 W 9

Total (non-PoE) Power allocated : 950 W
Total Power allocated for PoE : 0 W
Power Available (Redundant case) : 0 W
Total Power Available : 1550 W

```

- On an EX4300 switch:

```

user@switch> show chassis power-budget-statistics fpc 1
PSU 1 (JPSU-1100-AC-AF0-A) : 1100 W Online
Power redundancy configuration : N+0
Total power supplied by all online PSUs : 1100 W
Base power reserved : 175 W
Non-PoE power being consumed : 95 W
Total Power allocated for PoE : 925 W
Total PoE power consumed : 0 W
Total PoE power remaining : 925 W

```

2. Enable EEE on Base-T Copper Ethernet ports and save the configuration.

3. View the power budget of the switch after enabling EEE.

- On an EX6210 switch:

```

user@switch> show chassis power-budget-statistics
PSU 2 (EX6200-PWR-AC2500) : 2500 W Online
PSU 3 : 0 W Offline
Total Power supplied by all Online PSUs : 2500 W
Power Redundancy Configuration : N+1
Power Reserved for the Chassis : 500 W
Fan Tray Statistics
FTC 0 : Base power 300 W Power Used nan W
FPC Statistics
power Priority Base power Power Used PoE
FPC 3 (EX6200-48T) : 150 W 50.36 W
0 W 9
FPC 4 (EX6200-SRE64-4XS) : 100 W 48.60 W
0 W 0
FPC 5 (EX6200-SRE64-4XS) : 100 W 48.09 W
0 W 0
FPC 7 (EX6200-48T) : 150 W 51.38 W
0 W 9
FPC 8 (EX6200-48T) : 150 W 12.17 W
0 W 9

Total (non-PoE) Power allocated : 950 W
Total Power allocated for PoE : 0 W
Power Available (Redundant case) : 0 W
Total Power Available : 1550 W

```

- On an EX4300 switch:

```

user@switch> show chassis power-budget-statistics fpc 1
PSU 1 (JPSU-1100-AC-AF0-A) : 1100 W Online
Power redundancy configuration : N+0
Total power supplied by all online PSUs : 1100 W
Base power reserved : 175 W
Non-PoE power being consumed : 86 W
Total Power allocated for PoE : 925 W

```

```

Total PoE power consumed      :    0 W
Total PoE power remaining    :   925 W

```

**Meaning** On an EX6210 switch, the **Power Used** field in the output shows the actual power being consumed by the line card or SRE module, including PoE power. If you compare the values in the **Power Used** field before and after enabling EEE for FPC 3 and FPC 7, you will notice that power is saved when EEE is enabled.



**NOTE:** The **Power Used** field is displayed in the output only for EX6210 switches.

On an EX4300 switch, if you compare the values in the **Non-PoE power being consumed** field before and after enabling EEE, you will notice that power is saved when EEE is enabled.

- Related Documentation**
- [Configuring Energy Efficient Ethernet on Interfaces \(CLI Procedure\) on page 2684](#)
  - [Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces on page 2590](#)

## Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets

Verify that LACP has been set up correctly and that the bundle members are transmitting LACP protocol packets.

1. [Verifying the LACP Setup on page 2839](#)
2. [Verifying That LACP Packets Are Being Exchanged on page 2840](#)

### Verifying the LACP Setup

**Purpose** Verify that the LACP has been set up correctly.

**Action** To verify that LACP has been enabled as active on one end:

```

user@switch> show lacp interfaces xe-0/1/0
Aggregated interface: ae0

```

| LACP state:    | Role          | Exp | Def            | Dist | Col       | Syn | Aggr | Timeout | Activity |
|----------------|---------------|-----|----------------|------|-----------|-----|------|---------|----------|
| xe-0/1/0       | Actor         | No  | Yes            | No   | No        | No  | Yes  | Fast    | Active   |
| xe-0/1/0       | Partner       | No  | Yes            | No   | No        | No  | Yes  | Fast    | Passive  |
| LACP protocol: | Receive State |     | Transmit State |      | Mux State |     |      |         |          |
| xe-0/1/0       | Defaulted     |     | Fast periodic  |      | Detached  |     |      |         |          |

**Meaning** This output shows that LACP has been configured with one side as active and the other as passive. When LACP is enabled, at least one side must be set as active for the bundled link to be up.

### Verifying That LACP Packets Are Being Exchanged

---

**Purpose** Verify that LACP packets are being exchanged between interfaces.

**Action** Use the **show interfaces aex statistics** command to display LACP BPDU exchange information.

**show interfaces ae0 statistics**

```
Physical interface: ae0, Enabled, Physical link is Down
Interface index: 153, SNMP ifIndex: 30
Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
Minimum bandwidth needed: 0
Device flags   : Present Running
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Current address: 02:19:e2:50:45:e0, Hardware address: 02:19:e2:50:45:e0
Last flapped   : Never
Statistics last cleared: Never
  Input packets : 0
  Output packets: 0
Input errors: 0, Output errors: 0

Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
Statistics      Packets      pps      Bytes      bps
Bundle:
  Input :           0           0           0           0
  Output:           0           0           0           0
Protocol inet,
Flags: None
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

**Meaning** The output here shows that the link is down and that no PDUs are being exchanged (when there is no other traffic flowing on the link).

- Related Documentation**
- [Configuring Aggregated Ethernet LACP](#)
  - [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 2671](#)
  - [Verifying the Status of a LAG Interface](#)
  - [Verifying the Status of a LAG Interface on page 2837](#)

### Verifying That Layer 3 Subinterfaces Are Working

**Purpose** After configuring Layer 3 subinterfaces, verify they are set up properly and transmitting data.



- Action** 1. Use the **show interfaces** command to determine whether you successfully created the subinterfaces and the links are up:

```
user@switch> show interfaces interface-name terse
Interface           Admin Link Proto  Local           Remote
ge-0/0/0            up    up
ge-0/0/0.0          up    up   inet   1.1.1.1/24
ge-0/0/0.1          up    up   inet   2.1.1.1/24
ge-0/0/0.2          up    up   inet   3.1.1.1/24
ge-0/0/0.3          up    up   inet   4.1.1.1/24
ge-0/0/0.4          up    up   inet   5.1.1.1/24
ge-0/0/0.32767      up    up
```

2. Use the **ping** command from a device on one subnet to an address on another subnet to determine whether packets were transmitted correctly on the subinterface VLANs:

```
user@switch> ping ip-address
PING 1.1.1.1 (1.1.1.1): 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=64 time=0.157 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=64 time=0.238 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=64 time=0.255 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=64 time=0.128 ms
--- 1.1.1.1 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
```

**Meaning** The output confirms that the subinterfaces are created and the links are up.

- Related Documentation**
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 2689](#)
  - *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*

## Verifying Unicast RPF Status

**Purpose** Verify that unicast reverse-path forwarding (RPF) is enabled and is working on the interface.

- Action** Use one of the **show interfaces interface-name** commands with either the **extensive** or **detail** options to verify that unicast RPF is enabled and working on the switch. The following example displays output from the **show interfaces ge- extensive** command.

```
user@switch> show interfaces ge-1/0/10 extensive
Physical interface: ge-1/0/10, Enabled, Physical link is Down
Interface index: 139, SNMP ifIndex: 58, Generation: 140
Link-level type: Ethernet, MTU: 1514, Speed: Auto, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
Auto-negotiation: Enabled, Remote fault: Online
Device flags   : Present Running
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:19:e2:50:95:ab, Hardware address: 00:19:e2:50:95:ab
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes   :                               0                0 bps
```

```

Output bytes :                0                0 bps
Input packets:                0                0 pps
Output packets:               0                0 pps
IPv6 transit statistics:
  Input bytes :                0
  Output bytes :               0
  Input packets:              0
  Output packets:             0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort                0                0                0

  1 assured-forw              0                0                0

  5 expedited-fo              0                0                0

  7 network-cont              0                0                0

Active alarms : LINK
Active defects : LINK
MAC statistics:
  Receive      Transmit
  Total octets      0                0
  Total packets    0                0
  Unicast packets  0                0
  Broadcast packets 0                0
  Multicast packets 0                0
  CRC/Align errors 0                0
  FIFO errors      0                0
  MAC control frames 0                0
  MAC pause frames  0                0
  Oversized frames  0
  Jabber frames     0
  Fragment frames   0
  VLAN tagged frames 0
  Code violations   0
Filter statistics:
  Input packet count      0
  Input packet rejects    0
  Input DA rejects        0
  Input SA rejects        0
  Output packet count      0
  Output packet pad count  0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 1

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:

```

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Protocol inet, Generation: 144, Route table: 0
Flags: uRPF
Addresses, Flags: Is-Preferred Is-Primary

```

**Meaning** The `show interfaces ge-1/0/10 extensive` command (and the `show interfaces ge-1/0/10 detail` command) displays in-depth information about the interface. The **Flags:** output field near the bottom of the display reports the unicast RPF status. If unicast RPF has not been enabled, the **uRPF** flag is not displayed.

On EX3200, EX4200, and EX4300 switches, unicast RPF is implicitly enabled on *all* switch interfaces, including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs) when you enable unicast RPF on a single interface. However, the unicast RPF status is shown as enabled only on interfaces for which you have explicitly configured unicast RPF. Thus, the **uRPF** flag is not displayed on interfaces for which you have not explicitly configured unicast RPF even though unicast RPF is implicitly enabled on all interfaces on EX3200 and EX4200 switches.

- Related Documentation**
- [show interfaces xe- on page 2911](#)
  - [Example: Configuring Unicast RPF on an EX Series Switch](#)
  - [Configuring Unicast RPF \(CLI Procedure\) on page 2689](#)
  - [Disabling Unicast RPF \(CLI Procedure\) on page 2691](#)
  - [Troubleshooting Unicast RPF on page 2933](#)

## Verifying IP Directed Broadcast Status

**Purpose** Verify that IP directed broadcast is enabled and is working on the subnet.

**Action** Use the **show vlans extensive** command to verify that IP directed broadcast is enabled and working on the subnet as shown in *Example: Configuring IP Directed Broadcast on an EX Series Switch*.

**Related Documentation**

- [Configuring IP Directed Broadcast \(CLI Procedure\)](#)
- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 2691](#)
- [Example: Configuring IP Directed Broadcast on an EX Series Switch](#)

---

## Operational Commands

- [monitor interface](#)
- [request diagnostics tdr](#)
- [show diagnostics tdr](#)
- [show forwarding-options enhanced-hash-key](#)
- [show interfaces diagnostics optics](#)
- [show interfaces ge-](#)
- [show interfaces irb](#)
- [show interfaces mc-ae](#)
- [show interfaces me0](#)
- [show interfaces queue](#)
- [show interfaces xe-](#)
- [show lacp interfaces](#)
- [test interface restart-auto-negotiation](#)

## monitor interface

**Syntax** monitor interface  
 <interface-name> | traffic <detail>>

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display real-time statistics about interfaces, updating the statistics every second. Check for and display common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors.



**NOTE:** This command is not supported on the QFX3000 QFabric system.

**Options** none—Display real-time statistics for all interfaces.

**detail**—(Optional) With traffic option only, display detailed output.

**interface-name**—(Optional) Display real-time statistics for the specified interface. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified line-card chassis (LCC) only.

**traffic**—(Optional) Display traffic data for all active interfaces. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified LCC only.

**Additional Information** The output of this command shows how much each field has changed since you started the command or since you cleared the counters by pressing the c key. For a description of the statistical information provided in the output of this command, see the **show interfaces extensive** command for a particular interface type in the [CLI Explorer](#). To control the output of the **monitor interface** command while it is running, use the keys listed in [Table 300 on page 2845](#). The keys are not case-sensitive.

**Table 300: Output Control Keys for the monitor interface Command**

| Key | Action                                                                                                                                                                                                                     |
|-----|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| c   | Clears (returns to zero) the delta counters since <b>monitor interface</b> was started. This does not clear the accumulative counter. To clear the accumulative counter, use the <b>clear interfaces interval</b> command. |
| f   | Freezes the display, halting the display of updated statistics and delta counters.                                                                                                                                         |
| i   | Displays information about a different interface. The command prompts you for the name of a specific interface.                                                                                                            |

**Table 300: Output Control Keys for the monitor interface Command** (*continued*)

| Key      | Action                                                                                                                                                                                         |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| n        | Displays information about the next interface. The <b>monitor interface</b> command displays the physical or logical interfaces in the same order as the <b>show interfaces terse</b> command. |
| q or Esc | Quits the command and returns to the command prompt.                                                                                                                                           |
| t        | Thaws the display, resuming the update of the statistics and delta counters.                                                                                                                   |

To control the output of the **monitor interface traffic** command while it is running, use the keys listed in [Table 301 on page 2846](#). The keys are not case-sensitive.

**Table 301: Output Control Keys for the monitor interface traffic Command**

| Key      | Action                                                                                                               |
|----------|----------------------------------------------------------------------------------------------------------------------|
| b        | Displays the statistics in units of bits and bits per second (bps).                                                  |
| c        | Clears (return to 0) the delta counters in the <b>Current Delta</b> column. The statistics counters are not cleared. |
| d        | Displays the <b>Current Delta</b> column (instead of the rate column) in Bps or packets per second (pps).            |
| p        | Displays the statistics in units of packets and packets per second (pps).                                            |
| q or Esc | Quits the command and returns to the command prompt.                                                                 |
| r        | Displays the rate column (instead of the <b>Current Delta</b> column) in Bps and pps.                                |

**Required Privilege Level**

trace

**List of Sample Output**

[monitor interface \(Physical\) on page 2848](#)  
[monitor interface \(OTN Interface\) on page 2849](#)  
[monitor interface \(MX2020 Routers with MPC6E and OTN MICInterface\) on page 2850](#)  
[monitor interface \(Logical\) on page 2851](#)  
[monitor interface \(QFX3500 Switch\) on page 2851](#)  
[monitor interface traffic on page 2852](#)  
[monitor interface traffic \(QFX3500 Switch\) on page 2852](#)  
[monitor interface traffic detail \(QFX3500 Switch\) on page 2853](#)

**Output Fields**

[Table 302 on page 2847](#) describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 302: monitor interface Output Fields

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>routerl</b>           | Hostname of the router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |
| <b>Seconds</b>           | How long the monitor interface command has been running or how long since you last cleared the counters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| <b>Time</b>              | Current time (UTC).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels      |
| <b>Delay x/y/z</b>       | Time difference between when the statistics were displayed and the actual clock time. <ul style="list-style-type: none"> <li>• <b>x</b>—Time taken for the last polling (in milliseconds).</li> <li>• <b>y</b>—Minimum time taken across all pollings (in milliseconds).</li> <li>• <b>z</b>—Maximum time taken across all pollings (in milliseconds).</li> </ul>                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>Interface</b>         | Short description of the interface, including its name, status, and encapsulation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels      |
| <b>Link</b>              | State of the link: <b>Up</b> , <b>Down</b> , or <b>Test</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | All levels      |
| <b>Current delta</b>     | Cumulative number for the counter in question since the time shown in the Seconds field, which is the time since you started the command or last cleared the counters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels      |
| <b>Local Statistics</b>  | (Logical interfaces only) Number and rate of bytes and packets destined to the router or switch through the specified interface. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | All levels      |
| <b>Remote Statistics</b> | (Logical interfaces only) Statistics for traffic transiting the router or switch. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                | All levels      |

Table 302: monitor interface Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the interface. These statistics are the sum of the local and remote statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | All levels      |
| Description        | With the <b>traffic</b> option, displays the interface description configured at the <b>[edit interfaces <i>interface-name</i>]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail          |

## Sample Output

### monitor interface (Physical)

```

user@host> monitor interface so-0/0/0
router1                               Seconds: 19                      Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: 0C48
Traffic statistics:
    Input packets:                6045 (0 pps)
    Input bytes:                  6290065 (0 bps)
    Output packets:               10376 (0 pps)
    Output bytes:                 10365540 (0 bps)
Encapsulation statistics:
    Input keepalives:             1901
    Output keepalives:           1901
    NCP state: Opened
    LCP state: Opened
Error statistics:
    Input errors:                 0
    Input drops:                  0
    Input framing errors:         0
    Policed discards:             0
    L3 incompletes:              0
    L2 channel errors:           0
    L2 mismatch timeouts:        0
    Carrier transitions:          1
    Output errors:               0
    Output drops:                0
    Aged packets:                0
Active alarms : None
Active defects: None
SONET error counts/seconds:
    LOS count                    1
    LOF count                    1
    SEF count                    1
    ES-S                         0
    SES-S                        0
SONET statistics:
    BIP-B1                      458871

```



```

BIP-B2                      460072          [0]
REI-L                      465610          [0]
BIP-B3                      458978          [0]
REI-P                      458773          [0]

```

Received SONET overhead:

```

F1      : 0x00 J0      : 0x00 K1      : 0x00
K2      : 0x00 S1      : 0x00 C2      : 0x00
C2(cmp) : 0x00 F2      : 0x00 Z3      : 0x00
Z4      : 0x00 S1(cmp) : 0x00

```

Transmitted SONET overhead:

```

F1      : 0x00 J0      : 0x01 K1      : 0x00
K2      : 0x00 S1      : 0x00 C2      : 0xcf
F2      : 0x00 Z3      : 0x00 Z4      : 0x00

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

### monitor interface (OTN Interface)

```
user@host> monitor interface ge-7/0/0
```

```

Interface: ge-7/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
  Input bytes:                0 (0 bps)
  Output bytes:               0 (0 bps)
  Input packets:              0 (0 pps)
  Output packets:             0 (0 pps)
Error statistics:
  Input errors:                0
  Input drops:                 0
  Input framing errors:        0
  Policed discards:            0
  L3 incompletes:              0
  L2 channel errors:           0
  L2 mismatch timeouts:        0
  Carrier transitions:         5
  Output errors:               0
  Output drops:                0
  Aged packets:                0
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
  Unicast packets              0
  Broadcast packets            0
  Multicast packets            0
  Oversized frames             0
  Packet reject count          0
  DA rejects                   0
  SA rejects                   0
Output MAC/Filter Statistics:
  Unicast packets              0
  Broadcast packets            0
  Multicast packets            0
  Packet pad count             0
  Packet error count           0
OTN Link 0
  OTN Alarms: OTU_BDI, OTU_TTIM, ODU_BDI
  OTN Defects: OTU_BDI, OTU_TTIM, ODU_BDI, ODU_TTIM
  OTN OC - Seconds
    LOS                        2

```

```

LOF 9
OTN OTU - FEC Statistics
  Corr err ratio N/A
  Corr bytes 0
  Uncorr words 0
OTN OTU - Counters
  BIP 0
  BBE 0
  ES 0
  SES 0
  UAS 422
OTN ODU - Counters
  BIP 0
  BBE 0
  ES 0
  SES 0
  UAS 422
OTN ODU - Received Overhead APSPCC 0-3: 0

```

#### monitor interface (MX2020 Routers with MPC6E and OTN MICInterface)

```

user@host> monitor interface xe-3/0/0
host name Seconds: 67 Time: 23:46:46
Delay: 0/0/13

Interface: xe-3/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics: Current delta
  Input bytes: 0 (0 bps) [0]
  Output bytes: 0 (0 bps) [0]
  Input packets: 0 (0 pps) [0]
  Output packets: 0 (0 pps) [0]
Error statistics:
  Input errors: 0 [0]
  Input drops: 0 [0]
  Input framing errors: 0 [0]
  Policed discards: 0 [0]
  L3 incompletes: 0 [0]
  L2 channel errors: 0 [0]
  L2 mismatch timeouts: 0 [0]
  Carrier transitions: 3 [0]
  Output errors: 0 [0]
  Output drops: 0 [0]
  Aged packets: 0 [0]
OTN Link 0
OTN Alarms:
OTN Defects:
OTN OC - Seconds
  LOS 0 [0]
  LOF 0 [0]
OTN OTU - FEC Statistics
  Corr err ratio N/A
  Corr bytes 0 [0]
  Uncorr words 0 [0]
OTN OTU - Counters
  BIP 0 [0]
  BBE 0 [0]
  ES 0 [0]
  SES 0 [0]
  UAS 0 [0]
OTN ODU - Counters
  BIP 0 [0]

```

```

BBE                                0                                [0]
ES                                0                                [0]
SES                                0                                [0]
UAS                                0                                [0]
OTN ODU - Received Overhead       [0]
APSPCC 0-3:                        00 00 00 00

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

### monitor interface (Logical)

```

user@host> monitor interface so-1/0/0.0
host name                Seconds: 16                Time: 15:33:39
  Delay: 0/0/1

Interface: so-1/0/0.0, Enabled, Link is Down
Flags: Hardware-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Local statistics:
Input bytes:              0                                [0]
Output bytes:             0                                [0]
Input packets:            0                                [0]
Output packets:           0                                [0]
Remote statistics:
Input bytes:              0 (0 bps)                       [0]
Output bytes:             0 (0 bps)                       [0]
Input packets:            0 (0 pps)                       [0]
Output packets:           0 (0 pps)                       [0]
Traffic statistics:
Destination address: 192.168.8.193, Local: 192.168.8.21

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

```

### monitor interface (QFX3500 Switch)

```

user@switch> monitor interface ge-0/0/0
Interface: ge-0/0/0, Enabled, Link is Down
Encapsulation: Ethernet, Speed: Unspecified
Traffic statistics:
Input bytes:              0 (0 bps)                       [0]
Output bytes:             0 (0 bps)                       [0]
Input packets:            0 (0 pps)                       [0]
Output packets:           0 (0 pps)                       [0]
Error statistics:
Input errors:             0                                [0]
Input drops:              0                                [0]
Input framing errors:     0                                [0]
Policed discards:         0                                [0]
L3 incompletes:           0                                [0]
L2 channel errors:        0                                [0]
L2 mismatch timeouts:     0                                [0]
Carrier transitions:      0                                [0]
Output errors:            0                                [0]
Output drops:             0                                [0]
Aged packets:             0                                [0]
Active alarms : LINK
Active defects: LINK
Input MAC/Filter statistics:
Unicast packets           0                                [0]
Broadcast packets         0 Multicast packet             [0]

```

Interface warnings:  
 o Outstanding LINK alarm

### monitor interface traffic

```
user@host> monitor interface traffic
host name                               Seconds: 15                               Time: 12:31:09
```

| Interface | Link | Input packets | (pps) | Output packets | (pps) |
|-----------|------|---------------|-------|----------------|-------|
| so-1/0/0  | Down | 0             | (0)   | 0              | (0)   |
| so-1/1/0  | Down | 0             | (0)   | 0              | (0)   |
| so-1/1/1  | Down | 0             | (0)   | 0              | (0)   |
| so-1/1/2  | Down | 0             | (0)   | 0              | (0)   |
| so-1/1/3  | Down | 0             | (0)   | 0              | (0)   |
| t3-1/2/0  | Down | 0             | (0)   | 0              | (0)   |
| t3-1/2/1  | Down | 0             | (0)   | 0              | (0)   |
| t3-1/2/2  | Down | 0             | (0)   | 0              | (0)   |
| t3-1/2/3  | Down | 0             | (0)   | 0              | (0)   |
| so-2/0/0  | Up   | 211035        | (1)   | 36778          | (0)   |
| so-2/0/1  | Up   | 192753        | (1)   | 36782          | (0)   |
| so-2/0/2  | Up   | 211020        | (1)   | 36779          | (0)   |
| so-2/0/3  | Up   | 211029        | (1)   | 36776          | (0)   |
| so-2/1/0  | Up   | 189378        | (1)   | 36349          | (0)   |
| so-2/1/1  | Down | 0             | (0)   | 18747          | (0)   |
| so-2/1/2  | Down | 0             | (0)   | 16078          | (0)   |
| so-2/1/3  | Up   | 0             | (0)   | 80338          | (0)   |
| at-2/3/0  | Up   | 0             | (0)   | 0              | (0)   |
| at-2/3/1  | Down | 0             | (0)   | 0              | (0)   |

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D

### monitor interface traffic (QFX3500 Switch)

```
user@switch> monitor interface traffic
switch                               Seconds: 7                               Time: 16:04:37
```

| Interface | Link | Input packets | (pps) | Output packets | (pps) |
|-----------|------|---------------|-------|----------------|-------|
| ge-0/0/0  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/1  | Up   | 392187        | (0)   | 392170         | (0)   |
| ge-0/0/2  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/3  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/4  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/5  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/6  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/7  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/8  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/9  | Up   | 392184        | (0)   | 392171         | (0)   |
| ge-0/0/10 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/11 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/12 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/13 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/14 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/15 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/16 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/17 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/18 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/19 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/20 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/21 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/22 | Up   | 392172        | (0)   | 392187         | (0)   |
| ge-0/0/23 | Up   | 392185        | (0)   | 392173         | (0)   |

|       |      |   |     |         |     |
|-------|------|---|-----|---------|-----|
| vcp-0 | Down | 0 |     | 0       |     |
| vcp-1 | Down | 0 |     | 0       |     |
| ae0   | Down | 0 | (0) | 0       | (0) |
| bme0  | Up   | 0 |     | 1568706 |     |

### monitor interface traffic detail (QFX3500 Switch)

user@switch> monitor interface traffic detail  
switch

Seconds: 74

Time: 16:03:02

| Interface<br>Description | Link | Input packets | (pps) | Output packets | (pps) |
|--------------------------|------|---------------|-------|----------------|-------|
| ge-0/0/0                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/1                 | Up   | 392183        | (0)   | 392166         | (0)   |
| ge-0/0/2                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/3                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/4                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/5                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/6                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/7                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/8                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/9                 | Up   | 392181        | (0)   | 392168         | (0)   |
| ge-0/0/10                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/11                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/12                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/13                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/14                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/15                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/16                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/17                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/18                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/19                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/20                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/21                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/22                | Up   | 392169        | (0)   | 392184         | (1)   |
| ge-0/0/23                | Up   | 392182        | (0)   | 392170         | (0)   |
| vcp-0                    | Down | 0             |       | 0              |       |
| vcp-1                    | Down | 0             |       | 0              |       |
| ae0                      | Down | 0             | (0)   | 0              | (0)   |
| bme0                     | Up   | 0             |       | 1568693        |       |

## request diagnostics tdr

---

**Syntax**    request diagnostics tdr (abort | start) interface *interface-name*

**Release Information**    Command introduced in Junos OS Release 9.0 for EX Series switches.

**Description**    Start a time domain reflectometry (TDR) diagnostic test on the specified interface. This test characterizes and locates faults on twisted-pair Ethernet cables. For example, it can detect a broken twisted pair and provide the approximate distance to the break. It can also detect polarity swaps, pair swaps, and excessive skew.

The TDR test is supported on the following switches and interfaces:

- EX2200, EX3200, EX3300, and EX4200 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.
- EX6200 and EX8200 switches—RJ-45 interfaces on line cards.



**NOTE:** We recommend running the TDR test when there is no traffic on the interface under test.

---

You view the results of the TDR test with the [show diagnostics tdr](#) command.

**Options**    **abort**—Stop the TDR test currently in progress on the specified interface. No results are reported, and previous results, if any, are cleared.

**interface-name**—The name of the interface.

**start**—Start a TDR test on the specified interface.

**Required Privilege Level**    maintenance

**Related Documentation**

- [show diagnostics tdr on page 2856](#)
- [Diagnosing a Faulty Twisted-Pair Cable \(CLI Procedure\) on page 2934](#)

**List of Sample Output**    [request diagnostics tdr start interface ge-0/0/19 on page 2855](#)

**Output Fields**    [Table 303 on page 2855](#) lists the output fields for the **request diagnostics tdr** command. Output fields are listed in the approximate order in which they appear.

Table 303: request diagnostics tdr Output Fields

| Field Name  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Test Status | <p>Information about the status of the TDR test request:</p> <ul style="list-style-type: none"><li>• <b>Admin Down <i>interface-name</i></b>—The interface is administratively down. The TDR test cannot run on interfaces that are administratively down.</li><li>• <b>Interface <i>interface-name</i> not found</b>—The interface does not exist.</li><li>• <b>Test successfully executed <i>interface-name</i></b>—The test has successfully started on the interface. You can view the test results with the <b>show diagnostics tdr</b> command.</li><li>• <b>VCT not supported on <i>interface-name</i></b>—The TDR test is not supported on the interface.</li></ul> |

## Sample Output

request diagnostics tdr start interface ge-0/0/19

```
user@switch> request diagnostics tdr start interface ge-0/0/19
```

Interface TDR detail:

Test status : Test successfully executed ge-0/0/19

## show diagnostics tdr

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show diagnostics tdr</code><br><code>&lt;interface <i>interface-name</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Display the results of a time domain reflectometry (TDR) diagnostic test run on an interface. A TDR test characterizes and locates faults on twisted-pair Ethernet cables. For example, it can detect a broken twisted pair and provide the approximate distance to the break. It can also detect polarity swaps, pair swaps, and excessive skew.</p> <p>The TDR test is supported on the following switches and interfaces:</p> <ul style="list-style-type: none"><li>• EX2200, EX3200, EX3300, and EX4200 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.</li><li>• EX6200 and EX8200 switches— RJ-45 interfaces on line cards.</li></ul> <p>Use the <a href="#">request diagnostics tdr</a> command to request a TDR test on a specified interface. Use the <b>show diagnostic tdr</b> command to display the last TDR test results for a specified interface or the last TDR test results for all network interfaces on the switch that support the TDR test.</p> |
| <b>Options</b>                  | <p><b>none</b>—Show summarized last results for all interfaces on the switch that support the TDR test.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Show detailed last results for the specified interface or a range of interfaces. Specify a range of interfaces by entering the beginning and ending interface in the range, separated by a dash—for example, <b>ge-0/0/15-ge-0/0/20</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">request diagnostics tdr on page 2854</a></li><li>• <a href="#">Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure) on page 2934</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show diagnostics tdr interface ge-0/0/19 (Normal Cable) on page 2858</a><br><a href="#">show diagnostics tdr interface ge-2/0/2 (Faulty Cable) on page 2859</a><br><a href="#">show diagnostics tdr (All Supported Interfaces) on page 2859</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | <a href="#">Table 304 on page 2857</a> lists the output fields for the <b>show diagnostics tdr</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |



Table 304: show diagnostics tdr Output Fields

| Field Name                                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface name or Interface</b>          | Name of interface for which TDR test results are being reported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Test status</b>                          | <p>Status of TDR test:</p> <ul style="list-style-type: none"> <li>• <b>Aborted</b>—Test was terminated by operator before it was complete.</li> <li>• <b>Failed</b>—Test was not completed successfully.</li> <li>• <b>Interface <i>interface-name</i> not found</b>—Specified interface does not exist.</li> <li>• <b>Not Started</b>—No TDR test results are available for the interface.</li> <li>• <b>Passed</b>—Test completed successfully. The cable, however, might still have a fault—see the <b>Cable status</b> field for information on the cable.</li> <li>• <b>Started</b>—Test is currently running and not yet complete.</li> <li>• <b>VCT not supported on <i>interface-name</i></b>—TDR test is not supported on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                          |
| <b>Link status</b>                          | Operating status of link: <b>UP</b> or <b>Down</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>MDI pair</b>                             | Twisted pair for which test results are being reported, identified by pin numbers. (Displayed only when the <b>interface</b> option is used.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Cable status</b>                         | <p>When detailed information is displayed, status for a twisted pair:</p> <ul style="list-style-type: none"> <li>• <b>Failed</b>—TDR test failed on the cable pair.</li> <li>• <b>Impedance Mismatch</b>—Impedance on the twisted pair is not correct. Possible reasons for an impedance mismatch include: <ul style="list-style-type: none"> <li>• The twisted pair is not connected properly.</li> <li>• The twisted pair is damaged.</li> <li>• The connector is faulty.</li> </ul> </li> <li>• <b>Normal</b>—No cable fault detected for the twisted pair.</li> <li>• <b>Open</b>—Lack of continuity between the pins at each end of the twisted-pair.</li> <li>• <b>Short on Pair-<i>n</i></b>—A short-circuit was detected on the twisted pair.</li> </ul> <p>When summary information for all interfaces is displayed, status for the cable as a whole:</p> <ul style="list-style-type: none"> <li>• <b>Fault</b>—A fault was detected on one or more of the twisted-pairs.</li> <li>• <b>OK</b>—No fault was detected on any of the twisted pairs.</li> </ul> |
| <b>Distance fault or Max distance fault</b> | <p>Distance to the fault in whole meters. If there is no fault, this value is 0.</p> <p>When summary information for all interfaces is displayed, this value is the distance to the most distant fault if there is more than one twisted pair with a fault.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 304: show diagnostics tdr Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Polarity swap</b> | <p>Indicates the polarity status of the twisted pair:</p> <ul style="list-style-type: none"> <li>• <b>Normal</b>—Polarity is normal. Each conductor in the twisted pair has been connected the same pins at the both ends of the connection. For example, a conductor connected to pin 1 at the near end of the connection is connected to pin 1 at the far end.</li> <li>• <b>Reversed</b>—Polarity has been reversed. For the twisted pair, the conductors have switched which pins they are connected to at the near and far ends of the connection. For example, the conductor connected to pin 1 at the near end is connected to pin 2 at the far end.</li> </ul> <p>(Not available on EX8200 switches.) (Displayed only when the <b>interface</b> option is used)</p> |
| <b>Skew time</b>     | <p>Difference in nanoseconds between the propagation delay on this twisted pair and the twisted pair with the shortest propagation delay. (Not available on EX8200 switches.) (Displayed only when the <b>interface</b> option is used.)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Channel Pair</b>  | <p>Number of the 10/100BASE-T transmit/receive pair being reported on.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Pair Swap</b>     | <p>Indicates whether or not the twisted pairs are swapped:</p> <ul style="list-style-type: none"> <li>• <b>MDI</b>—The pairs are not swapped (straight-through cable).</li> <li>• <b>MDIX</b>—The pairs are swapped (cross-over cable).</li> </ul> <p>(Displayed only when the <b>interface</b> option is used.)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Downshift</b>     | <p>Indicates whether the connection speed is being downshifted:</p> <ul style="list-style-type: none"> <li>• <b>No Downshift</b>—No downshifting of connection speed.</li> <li>• <b>Downshift occurs</b>—Connection speed is downshifted to 10 or 100 Mbs. This occurs if the cable is a two-pair cable rather than the four-pair cable required by Gigabit Ethernet.</li> </ul> <p>(Displayed only when the <b>interface</b> option is used.)</p>                                                                                                                                                                                                                                                                                                                          |

## Sample Output

### show diagnostics tdr interface ge-0/0/19 (Normal Cable)

```

user@switch> show diagnostics tdr interface ge-0/0/19
Interface TDR detail:
Interface name           : ge-0/0/19
Test status              : Passed
Link status              : UP
MDI pair                 : 1-2
  Cable status           : Normal
  Distance fault         : 0 Meters
  Polartiy swap          : Normal
  Skew time              : 0 ns
MDI pair                 : 3-6

```

```

Cable status           : Normal
Distance fault         : 0 Meters
Polartiy swap          : Normal
Skew time              : 8 ns
MDI pair               : 4-5
Cable status           : Normal
Distance fault         : 0 Meters
Polartiy swap          : Normal
Skew time              : 8 ns
MDI pair               : 7-8
Cable status           : Normal
Distance fault         : 0 Meters
Polartiy swap          : Normal
Skew time              : 8 ns
Channel pair           : 1
Pair swap              : MDI
Channel pair           : 2
Pair swap              : MDI
Downshift              : No Downshift

```

#### show diagnostics tdr interface ge-2/0/2 (Faulty Cable)

```

user@switch> show diagnostics tdr interface ge-2/0/2
Interface TDR detail:
Interface name         : ge-2/0/2
Test status            : Passed
Link status            : Down
MDI Pair              : 1-2
  Cable status         : 1-2
  Distance fault       : 2 Meters
  Polartiy swap        : N/A
  Skew time            : N/A
MDI Pair              : 3-6
  Cable status         : Impedance Mismatch
  Distance fault       : 3 Meters
  Polartiy swap        : N/A
  Skew time            : N/A
MDI Pair              : 4-5
  Cable status         : Impedance Mismatch
  Distance fault       : 3 Meters
  Polartiy swap        : N/A
  Skew time            : N/A
MDI Pair              : 7-8
  Cable status         : Short on Pair-2
  Distance fault       : 3 Meters
  Polartiy swap        : N/A
  Skew time            : N/A
Channel pair          : 1
Pair swap             : N/A
Channel pair          : 2
Pair swap             : N/A
Downshift             : N/A

```

#### show diagnostics tdr (All Supported Interfaces)

```

user@switch> show diagnostics tdr

```

| Interface | Test status | Link status | Cable status | Max distance fault |
|-----------|-------------|-------------|--------------|--------------------|
| ge-0/0/0  | Not Started | N/A         | N/A          | N/A                |
| ge-0/0/1  | Not Started | N/A         | N/A          | N/A                |
| ge-0/0/2  | Started     | N/A         | N/A          | N/A                |
| ge-0/0/3  | Started     | N/A         | N/A          | N/A                |

|           |        |      |       |     |
|-----------|--------|------|-------|-----|
| ge-0/0/4  | Passed | UP   | OK    | 0   |
| ge-0/0/5  | Passed | UP   | Fault | 173 |
| ge-0/0/6  | Passed | UP   | OK    | 0   |
| ge-0/0/7  | Passed | UP   | OK    | 0   |
| ge-0/0/8  | Passed | UP   | OK    | 0   |
| ge-0/0/9  | Passed | UP   | OK    | 0   |
| ge-0/0/10 | Passed | UP   | OK    | 0   |
| ge-0/0/11 | Passed | UP   | OK    | 0   |
| ge-0/0/12 | Passed | UP   | OK    | 0   |
| ge-0/0/13 | Passed | UP   | OK    | 0   |
| ge-0/0/14 | Passed | UP   | OK    | 0   |
| ge-0/0/15 | Passed | UP   | OK    | 0   |
| ge-0/0/16 | Passed | UP   | OK    | 0   |
| ge-0/0/17 | Passed | UP   | OK    | 0   |
| ge-0/0/18 | Passed | UP   | OK    | 0   |
| ge-0/0/19 | Passed | UP   | OK    | 0   |
| ge-0/0/20 | Passed | Down | Fault | 0   |
| ge-0/0/21 | Passed | Down | Fault | 5   |
| ge-0/0/22 | Passed | UP   | OK    | 0   |
| ge-0/0/23 | Passed | UP   | OK    | 0   |

## show forwarding-options enhanced-hash-key

|                                 |                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show forwarding-options enhanced-hash-key</b>                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p><b>Fabric Load Balancing Options</b> output fields introduced in Junos OS Release 14.1X53-D10.</p>                                                                                                |
| <b>Description</b>              | <p>Display information about which packet fields are used by the hashing algorithm to make hashing decisions.</p> <p>You can configure the fields that are inspected by the hashing algorithm to make hashing decisions for traffic entering a LAG bundle using the <b>forwarding-options enhanced-hash-key</b> statement.</p>                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2686</a></li> <li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2585</a></li> <li>• <a href="#">enhanced-hash-key on page 2764</a></li> </ul> |
| <b>List of Sample Output</b>    | <p><a href="#">show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode) on page 2862</a></p> <p><a href="#">show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode) on page 2863</a></p> <p><a href="#">show forwarding-options enhanced-hash-key (Fabric Load Balancing Options) on page 2863</a></p>                                         |
| <b>Output Fields</b>            | <p><a href="#">Table 305 on page 2861</a> lists the output fields for the <b>show forwarding-options enhanced-hash-key</b> command. Output fields are listed in the approximate order in which they first appear.</p>                                                                                                                                                    |

**Table 305: show forwarding-options enhanced-hash-key Output Fields**

| Field Name                   | Field Description                                                                                        |
|------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Hash-Mode</b>             | Current hash mode: Layer 2 header or Layer 2 payload.                                                    |
| <b>Protocol</b>              | Indicates whether the Protocol field is or is not used by the hashing algorithm: Yes or No.              |
| <b>Destination L4 Port</b>   | Indicates whether the Destination L4 Port field is or is not used by the hashing algorithm: Yes or No.   |
| <b>Source L4 Port</b>        | Indicates whether the Source L4 Port field is or is not used by the hashing algorithm: Yes or No.        |
| <b>Destination IPv4 Addr</b> | Indicates whether the Destination IPv4 Addr field is or is not used by the hashing algorithm: Yes or No. |

Table 305: show forwarding-options enhanced-hash-key Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                    |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Source IPv4 Addr</b>        | Indicates whether the Source IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                  |
| <b>Vlan id</b>                 | Indicates whether the Vlan id field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                           |
| <b>Next Hdr</b>                | Indicates whether the Next Hdr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                          |
| <b>Destination IPv6 Addr</b>   | Indicates whether the Destination IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                             |
| <b>Source IPv6 Addr</b>        | Indicates whether the Source IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                  |
| <b>Ether Type</b>              | Indicates whether the Ether Type field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                        |
| <b>Destination MAC Address</b> | Indicates whether the Destination MAC Address field is or is not used by the hashing algorithm: Yes or No.                                                                                                                           |
| <b>Source MAC Address</b>      | Indicates whether the Source MAC Address field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                |
| <b>Load Balancing Method</b>   | Indicates the load balancing method for adaptive load balancing (ALB): flowlet or per-packet.<br><br>The load balancing method is flowlet by default, and can be configured using the <a href="#">fabric-load-balance</a> statement. |
| <b>Fabric Link Scale</b>       | Indicates the fabric link scale, in mbps.                                                                                                                                                                                            |
| <b>Inactivity Interval</b>     | Indicates the fabric load balance inactivity interval, in microseconds (us).<br><br>The inactivity interval is 16 microseconds by default, and can be configured using the <a href="#">inactivity-interval</a> statement.            |
| <b>Hash Region Size/Trunk</b>  | Indicates the hash region size, in buckets per fabric trunk.                                                                                                                                                                         |

## Sample Output

### show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0
```

```
Current Hash Settings
-----
```

```
Hash-Mode                               :layer2-payload
```

```
inet Hash settings-
```

```
-----
```

```
inet packet fields
```

```
Protocol                               : Yes
Destination L4 Port                    : Yes
Source L4 Port                         : Yes
Destination IPv4 Addr                  : Yes
Source IPv4 Addr                       : Yes
Vlan id                               : No
```

```
inet6 Hash settings-
```

```
-----
```

```
inet6 packet fields
```

```
Next Hdr                             : Yes
Destination L4 Port                    : Yes
Source L4 Port                         : Yes
Destination IPv6 Addr                  : Yes
Source IPv6 Addr                       : Yes
Vlan id                               : No
```

#### show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0
```

```
Current Hash Settings
```

```
-----
```

```
Hash-Mode                               : layer2-header
```

```
layer2 Hash settings-
```

```
-----
```

```
layer2 packet fields
```

```
Ether Type                           : Yes
Destination MAC Address                : Yes
Source MAC Address                     : Yes
VLAN ID                               : No
```

#### show forwarding-options enhanced-hash-key (Fabric Load Balancing Options)

```
user@switch> show forwarding-options enhanced-hash-key
<some output removed for brevity>
```

```
Fabric Load Balancing Options
```

```
-----
```

```
Load Balancing Method : Flowlet
Fabric Link Scale      : 40960 (mbps)
Inactivity Interval    : 16 (us)
Hash Region Size/Trunk : 1024 (buckets)
```

## show interfaces diagnostics optics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show interfaces diagnostics optics <i>interface-name</i></code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP, SFP+, XFP, QSFP+, or CFP) installed in EX Series or QFX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information. For a list of transceivers supported on EX Series switches and their specifications, including DOM support, see <i>Pluggable Transceivers Supported on EX Series Switches</i>.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.</p> |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface associated with the port in which the transceiver is installed: <i>ge-fpc/pic/port</i> , <i>xe-fpc/pic/port</i> , or <i>et-fpc/pic/port</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring Interface Status and Traffic on page 2835</a></li> <li>• <i>Monitoring Interface Status and Traffic</i></li> <li>• <i>Installing a Transceiver in an EX Series Switch</i></li> <li>• <i>Installing a Transceiver in a QFX Series Device</i></li> <li>• <i>Removing a Transceiver from an EX Series Switch</i></li> <li>• <i>Removing a Transceiver from a QFX Series Device</i></li> <li>• <a href="#">Junos OS Ethernet Interfaces Configuration Guide</a></li> </ul>                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver) on page 2871</a><br><a href="#">show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver) on page 2872</a><br><a href="#">show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver) on page 2873</a><br><a href="#">show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver) on page 2874</a><br><a href="#">show interfaces diagnostics optics et-4/1/0 (CFP Transceiver) on page 2875</a>                                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | <a href="#">Table 306 on page 2864</a> lists the output fields for the <b>show interfaces diagnostics optics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

Table 306: show interfaces diagnostics optics Output Fields

| Field Name         | Field Description                            |
|--------------------|----------------------------------------------|
| Physical interface | Displays the name of the physical interface. |



Table 306: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                           | Field Description                                                                                                                                                  |
|------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Laser bias current</b>                                                                            | Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents. |
| <b>Laser output power</b><br>(Not available for QSFP+ transceivers)                                  | Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).                                                                         |
| <b>Laser temperature</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)               | Displays the laser temperature, in Celsius and Fahrenheit.                                                                                                         |
| <b>Module temperature</b>                                                                            | Displays the temperature, in Celsius and Fahrenheit.                                                                                                               |
| <b>Module voltage</b><br>(Not available for XFP transceivers)                                        | Displays the voltage, in Volts.                                                                                                                                    |
| <b>Laser rx power</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)                  | Displays the laser received optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).                                                               |
| <b>Receiver signal average optical power</b><br>(Not available for XFP, QSFP+, and CFP transceivers) | Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).                                                      |
| <b>Laser bias current high alarm</b>                                                                 | Displays whether the laser bias power setting high alarm is <b>On</b> or <b>Off</b> .                                                                              |
| <b>Laser bias current low alarm</b>                                                                  | Displays whether the laser bias power setting low alarm is <b>On</b> or <b>Off</b> .                                                                               |
| <b>Laser bias current high warning</b>                                                               | Displays whether the laser bias power setting high warning is <b>On</b> or <b>Off</b> .                                                                            |
| <b>Laser bias current low warning</b>                                                                | Displays whether the laser bias power setting low warning is <b>On</b> or <b>Off</b> .                                                                             |
| <b>Laser output power high alarm</b><br>(Not available for QSFP+ transceivers)                       | Displays whether the laser output power high alarm is <b>On</b> or <b>Off</b> .                                                                                    |
| <b>Laser output power low alarm</b><br>(Not available for QSFP+ transceivers)                        | Displays whether the laser output power low alarm is <b>On</b> or <b>Off</b> .                                                                                     |
| <b>Laser output power high warning</b><br>(Not available for QSFP+ transceivers)                     | Displays whether the laser output power high warning is <b>On</b> or <b>Off</b> .                                                                                  |

Table 306: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                          | Field Description                                                                 |
|-----------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|
| <b>Laser output power low warning</b><br>(Not available for QSFP+ transceivers)                     | Displays whether the laser output power low warning is <b>On</b> or <b>Off</b> .  |
| <b>Laser temperature high alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays whether the laser temperature high alarm is <b>On</b> or <b>Off</b> .    |
| <b>Laser temperature low alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)    | Displays whether the laser temperature low alarm is <b>On</b> or <b>Off</b> .     |
| <b>Laser temperature high warning</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers) | Displays whether the laser temperature high warning is <b>On</b> or <b>Off</b> .  |
| <b>Laser temperature low warning</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)  | Displays whether the laser temperature low warning is <b>On</b> or <b>Off</b> .   |
| <b>Module temperature high alarm</b><br>(Not available for QSFP+ transceivers)                      | Displays whether the module temperature high alarm is <b>On</b> or <b>Off</b> .   |
| <b>Module temperature low alarm</b><br>(Not available for QSFP+ transceivers)                       | Displays whether the module temperature low alarm is <b>On</b> or <b>Off</b> .    |
| <b>Module temperature high warning</b><br>(Not available for QSFP+ transceivers)                    | Displays whether the module temperature high warning is <b>On</b> or <b>Off</b> . |
| <b>Module temperature low warning</b><br>(Not available for QSFP+ transceivers)                     | Displays whether the module temperature low warning is <b>On</b> or <b>Off</b> .  |
| <b>Module voltage high alarm</b><br>(Not available for XFP and QSFP+ transceivers)                  | Displays whether the module voltage high alarm is <b>On</b> or <b>Off</b> .       |
| <b>Module voltage low alarm</b><br>(Not available for XFP and QSFP+ transceivers)                   | Displays whether the module voltage low alarm is <b>On</b> or <b>Off</b> .        |
| <b>Module voltage high warning</b><br>(Not available for XFP and QSFP+ transceivers)                | Displays whether the module voltage high warning is <b>On</b> or <b>Off</b> .     |
| <b>Module voltage low warning</b><br>(Not available for XFP and QSFP+ transceivers)                 | Displays whether the module voltage low warning is <b>On</b> or <b>Off</b> .      |

Table 306: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                                  | Field Description                                                                                                                            |
|-------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Laser rx power high alarm</b><br>(Not available for QSFP+ and CFP transceivers)                          | Displays whether the receive laser power high alarm is <b>On</b> or <b>Off</b> .                                                             |
| <b>Laser rx power low alarm</b><br>(Not available for QSFP+ and CFP transceivers)                           | Displays whether the receive laser power low alarm is <b>On</b> or <b>Off</b> .                                                              |
| <b>Laser rx power high warning</b><br>(Not available for QSFP+ and CFP transceivers)                        | Displays whether the receive laser power high warning is <b>On</b> or <b>Off</b> .                                                           |
| <b>Laser rx power low warning</b><br>(Not available for QSFP+ and CFP transceivers)                         | Displays whether the receive laser power low warning is <b>On</b> or <b>Off</b> .                                                            |
| <b>Laser bias current high alarm threshold</b><br>(Not available for QSFP+ transceivers)                    | Displays the vendor-specified threshold for the laser bias current high alarm.                                                               |
| <b>Module not ready alarm</b><br>(Not available for SFP, SFP+, and QSFP+ transceivers)                      | Displays whether the module not ready alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , the module has an operational fault. |
| <b>Module low power alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                 | Displays whether the module low power alarm is <b>On</b> or <b>Off</b> .                                                                     |
| <b>Module initialization incomplete alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers) | Displays whether the module initialization incomplete alarm is <b>On</b> or <b>Off</b> .                                                     |
| <b>Module fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                     | Displays whether the module fault alarm is <b>On</b> or <b>Off</b> .                                                                         |
| <b>PLD Flash initialization fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays whether the PLD Flash initialization fault alarm is <b>On</b> or <b>Off</b> .                                                       |
| <b>Power supply fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)               | Displays whether the power supply fault alarm is <b>On</b> or <b>Off</b> .                                                                   |
| <b>Checksum fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                   | Displays whether the checksum fault alarm is <b>On</b> or <b>Off</b> .                                                                       |
| <b>Tx laser disabled alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                | Displays whether the Tx laser disabled alarm is <b>On</b> or <b>Off</b> .                                                                    |

Table 306: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                   | Field Description                                                                                                                                                                                |
|----------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Module power down alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers) | Displays whether the module power down alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , module is in a limited power mode, low for normal operation.                            |
| <b>Tx data not ready alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers) | Any condition leading to invalid data on the transmit path. Displays whether the Tx data not ready alarm is <b>On</b> or <b>Off</b> .                                                            |
| <b>Tx not ready alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)      | Any condition leading to invalid data on the transmit path. Displays whether the Tx not ready alarm is <b>On</b> or <b>Off</b> .                                                                 |
| <b>Tx laser fault alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)    | Laser fault condition. Displays whether the Tx laser fault alarm is <b>On</b> or <b>Off</b> .                                                                                                    |
| <b>Tx CDR loss of lock alarm</b><br>(Not available for SFP, SFP+, and QSFP+ transceivers)    | Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays whether the Tx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .                   |
| <b>Rx not ready alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)      | Any condition leading to invalid data on the receive path. Displays whether the Rx not ready alarm is <b>On</b> or <b>Off</b> .                                                                  |
| <b>Rx loss of signal alarm</b><br>(Not available for SFP and SFP+ transceivers)              | Receive loss of signal alarm. When the output is <b>On</b> , indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is <b>On</b> or <b>Off</b> . |
| <b>Rx CDR loss of lock alarm</b><br>(Not available for SFP, SFP+, and QSFP+ transceivers)    | Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays whether the Rx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .                                               |
| <b>Laser bias current low alarm threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the laser bias current low alarm.                                                                                                                    |
| <b>Laser bias current high warning threshold</b><br>(Not available for QSFP+ transceivers)   | Displays the vendor-specified threshold for the laser bias current high warning.                                                                                                                 |
| <b>Laser bias current low warning threshold</b><br>(Not available for QSFP+ transceivers)    | Displays the vendor-specified threshold for the laser bias current low warning.                                                                                                                  |
| <b>Laser output power high alarm threshold</b><br>(Not available for QSFP+ transceivers)     | Displays the vendor-specified threshold for the laser output power high alarm.                                                                                                                   |
| <b>Laser output power low alarm threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the laser output power low alarm.                                                                                                                    |

Table 306: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                     | Field Description                                                                |
|------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| <b>Laser output power high warning threshold</b><br>(Not available for QSFP+ transceivers)     | Displays the vendor-specified threshold for the laser output power high warning. |
| <b>Laser output power low warning threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the laser output power low warning.  |
| <b>Module temperature high alarm threshold</b><br>(Not available for QSFP+ transceivers)       | Displays the vendor-specified threshold for the module temperature high alarm.   |
| <b>Module temperature low alarm threshold</b><br>(Not available for QSFP+ transceivers)        | Displays the vendor-specified threshold for the module temperature low alarm.    |
| <b>Module temperature high warning threshold</b><br>(Not available for QSFP+ transceivers)     | Displays the vendor-specified threshold for the module temperature high warning. |
| <b>Module temperature low warning threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the module temperature low warning.  |
| <b>Module voltage high alarm threshold</b><br>(Not available for XFP and QSFP+ transceivers)   | Displays the vendor-specified threshold for the module voltage high alarm.       |
| <b>Module voltage low alarm threshold</b><br>(Not available for XFP and QSFP+ transceivers)    | Displays the vendor-specified threshold for the module voltage low alarm.        |
| <b>Module voltage high warning threshold</b><br>(Not available for XFP and QSFP+ transceivers) | Displays the vendor-specified threshold for the module voltage high warning.     |
| <b>Module voltage low warning threshold</b><br>(Not available for XFP and QSFP+ transceivers)  | Displays the vendor-specified threshold for the module voltage low warning.      |
| <b>Laser rx power high alarm threshold</b><br>(Not available for QSFP+ transceivers)           | Displays the vendor-specified threshold for the laser rx power high alarm.       |
| <b>Laser rx power low alarm threshold</b><br>(Not available for QSFP+ transceivers)            | Displays the vendor-specified threshold for the laser rx power low alarm.        |
| <b>Laser rx power high warning threshold</b><br>(Not available for QSFP+ transceivers)         | Displays the vendor-specified threshold for the laser rx power high warning.     |

Table 306: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                                    | Field Description                                                                                          |
|---------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| <b>Laser rx power low warning threshold</b><br>(Not available for QSFP+ transceivers)                         | Displays the vendor-specified threshold for the laser rx power low warning.                                |
| <b>Laser temperature high alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays the vendor-specified threshold for the laser temperature high alarm, in Celsius and Fahrenheit.   |
| <b>Laser temperature low alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)    | Displays the vendor-specified threshold for the laser temperature low alarm, in Celsius and Fahrenheit.    |
| <b>Laser temperature high warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers) | Displays the vendor-specified threshold for the laser temperature high warning, in Celsius and Fahrenheit. |
| <b>Laser temperature low warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)  | Displays the vendor-specified threshold for the laser temperature low warning, in Celsius and Fahrenheit.  |
| <b>SOA bias current high alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)    | Displays the vendor-specified threshold for SOA bias current high alarm.                                   |
| <b>SOA bias current low alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)     | Displays the vendor-specified threshold for SOA bias current low alarm.                                    |
| <b>SOA bias current high warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)  | Displays the vendor-specified threshold for SOA bias current high warning.                                 |
| <b>SOA bias current low warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays the vendor-specified threshold for SOA bias current low warning.                                  |
| <b>Laser receiver power high alarm</b><br>(Not available for SFP, SFP+, and XFP transceivers)                 | Displays whether the laser receiver power high alarm is <b>On</b> or <b>Off</b> .                          |
| <b>Laser receiver power low alarm</b><br>(Not available for SFP, SFP+, and XFP transceivers)                  | Displays whether the laser receiver power low alarm is <b>On</b> or <b>Off</b> .                           |
| <b>Laser receiver power high warning</b><br>(Not available for SFP, SFP+, and XFP transceivers)               | Displays whether the laser receiver power high warning is <b>On</b> or <b>Off</b> .                        |
| <b>Laser receiver power low warning</b><br>(Not available for SFP, SFP+, and XFP transceivers)                | Displays whether the laser receiver power low warning is <b>On</b> or <b>Off</b> .                         |

Table 306: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name                                                                                          | Field Description                                                                            |
|-----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| <b>Laser receiver power</b><br>(Not available for SFP, SFP+, and XFP transceivers)                  | Displays the laser receiver power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm). |
| <b>Tx loss of signal functionality alarm</b><br>(Not available for SFP, SFP+, and XFP transceivers) | Displays whether the Tx loss of signal functionality alarm is <b>On</b> or <b>Off</b> .      |
| <b>APD supply fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)         | Displays whether the APD supply fault alarm is <b>On</b> or <b>Off</b> .                     |
| <b>TEC fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                | Displays whether the TEC fault alarm is <b>On</b> or <b>Off</b> .                            |
| <b>Wavelength unlocked alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)      | Displays whether the Wavelength unlocked alarm is <b>On</b> or <b>Off</b> .                  |

## Sample Output

### show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver)

```

user@switch> show interfaces diagnostics optics ge-0/1/0
Physical interface: ge-0/1/0
  Laser bias current           : 5.444 mA
  Laser output power          : 0.3130 mW / -5.04 dBm
  Module temperature          : 36 degrees C / 97 degrees F
  Module voltage              : 3.2120 V
  Receiver signal average optical power : 0.3840 mW / -4.16 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm  : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm : Off
  Module temperature low alarm  : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Module voltage high alarm     : Off
  Module voltage low alarm      : Off
  Module voltage high warning   : Off
  Module voltage low warning    : Off
  Laser rx power high alarm     : Off
  Laser rx power low alarm      : Off
  Laser rx power high warning   : Off
  Laser rx power low warning    : Off
  Laser bias current high alarm threshold : 15.000 mA
  Laser bias current low alarm threshold  : 1.000 mA
  Laser bias current high warning threshold : 12.000 mA

```

```

Laser bias current low warning threshold : 2.000 mA
Laser output power high alarm threshold : 0.6300 mW / -2.01 dBm
Laser output power low alarm threshold : 0.0660 mW / -11.80 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0780 mW / -11.08 dBm
Module temperature high alarm threshold : 109 degrees C / 228 degrees F
Module temperature low alarm threshold : -29 degrees C / -20 degrees F
Module temperature high warning threshold : 103 degrees C / 217 degrees F
Module temperature low warning threshold : -13 degrees C / 9 degrees F
Module voltage high alarm threshold : 3.900 V
Module voltage low alarm threshold : 2.700 V
Module voltage high warning threshold : 3.700 V
Module voltage low warning threshold : 2.900 V
Laser rx power high alarm threshold : 1.2589 mW / 1.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7939 mW / -1.00 dBm
Laser rx power low warning threshold : 0.0157 mW / -18.04 dBm

```

## Sample Output

### show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver)

```

user@switch> show interfaces diagnostics optics xe-0/1/0
Physical interface: xe-0/1/0
Laser bias current : 4.968 mA
Laser output power : 0.4940 mW / -3.06 dBm
Module temperature : 27 degrees C / 81 degrees F
Module voltage : 3.2310 V
Receiver signal average optical power : 0.0000
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : On
Laser rx power high warning : Off
Laser rx power low warning : On
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F

```



```

Module voltage high alarm threshold      : 3.630 V
Module voltage low alarm threshold       : 2.970 V
Module voltage high warning threshold    : 3.465 V
Module voltage low warning threshold     : 3.135 V
Laser rx power high alarm threshold      : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold       : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold    : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold     : 0.1023 mW / -9.90 dBm

```

## Sample Output

### show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver)

```

user@switch> show interfaces diagnostics optics xe-0/1/0
Physical interface: xe-0/1/0
Laser bias current                : 8.029 mA
Laser output power                : 0.6430 mW / -1.92 dBm
Module temperature                : 4 degrees C / 39 degrees F
Laser rx power                    : 0.0012 mW / -29.21 dBm
Laser bias current high alarm     : Off
Laser bias current low alarm      : Off
Laser bias current high warning   : Off
Laser bias current low warning    : Off
Laser output power high alarm     : Off
Laser output power low alarm      : Off
Laser output power high warning   : Off
Laser output power low warning    : Off
Module temperature high alarm     : Off
Module temperature low alarm      : Off
Module temperature high warning   : Off
Module temperature low warning    : Off
Laser rx power high alarm         : Off
Laser rx power low alarm          : On
Laser rx power high warning       : Off
Laser rx power low warning        : On
Module not ready alarm            : On
Module power down alarm           : Off
Tx data not ready alarm           : Off
Tx not ready alarm                : Off
Tx laser fault alarm              : Off
Tx CDR loss of lock alarm         : Off
Rx not ready alarm                : On
Rx loss of signal alarm           : On
Rx CDR loss of lock alarm         : On
Laser bias current high alarm threshold : 13.000 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 12.000 mA
Laser bias current low warning threshold : 3.000 mA
Laser output power high alarm threshold : 0.8310 mW / -0.80 dBm
Laser output power low alarm threshold : 0.1650 mW / -7.83 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 90 degrees C / 194 degrees F
Module temperature low alarm threshold : 0 degrees C / 32 degrees F
Module temperature high warning threshold : 85 degrees C / 185 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Laser rx power high alarm threshold : 0.8912 mW / -0.50 dBm
Laser rx power low alarm threshold : 0.0912 mW / -10.40 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

## Sample Output

### show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver)

```

user@switch> show interfaces diagnostics optics et-3/0/0
Physical interface: et-3/0/0
  Module temperature                : 33 degrees C / 92 degrees F
  Module voltage                    : 3.3060 V
Lane 0
  Laser bias current                : 7.182 mA
  Laser receiver power              : 0.743 mW / -1.29 dBm
  Laser bias current high alarm     : Off
  Laser bias current low alarm      : Off
  Laser bias current high warning   : Off
  Laser bias current low warning    : Off
  Laser receiver power high alarm    : Off
  Laser receiver power low alarm     : Off
  Laser receiver power high warning  : Off
  Laser receiver power low warning   : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm           : Off
Lane 1
  Laser bias current                : 7.326 mA
  Laser receiver power              : 0.752 mW / -1.24 dBm
  Laser bias current high alarm     : Off
  Laser bias current low alarm      : Off
  Laser bias current high warning   : Off
  Laser bias current low warning    : Off
  Laser receiver power high alarm    : Off
  Laser receiver power low alarm     : Off
  Laser receiver power high warning  : Off
  Laser receiver power low warning   : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm           : Off
Lane 2
  Laser bias current                : 7.447 mA
  Laser receiver power              : 0.790 mW / -1.03 dBm
  Laser bias current high alarm     : Off
  Laser bias current low alarm      : Off
  Laser bias current high warning   : Off
  Laser bias current low warning    : Off
  Laser receiver power high alarm    : Off
  Laser receiver power low alarm     : Off
  Laser receiver power high warning  : Off
  Laser receiver power low warning   : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm           : Off
Lane 3
  Laser bias current                : 7.734 mA
  Laser receiver power              : 0.768 mW / -1.15 dBm
  Laser bias current high alarm     : Off
  Laser bias current low alarm      : Off
  Laser bias current high warning   : Off
  Laser bias current low warning    : Off
  Laser receiver power high alarm    : Off
  Laser receiver power low alarm     : Off
  Laser receiver power high warning  : Off
  Laser receiver power low warning   : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm           : Off

```

## Sample Output

### show interfaces diagnostics optics et-4/1/0 (CFP Transceiver)

```

user@switch> show interfaces diagnostics optics et-4/1/0
Physical interface: et-4/1/0
  Module temperature                : 38 degrees C / 101 degrees F
  Module voltage                    : 3.2500 V
  Module temperature high alarm     : Off
  Module temperature low alarm      : Off
  Module temperature high warning   : Off
  Module temperature low warning    : Off
  Module voltage high alarm         : Off
  Module voltage low alarm          : Off
  Module voltage high warning       : Off
  Module voltage low warning        : Off
  Module not ready alarm            : Off
  Module low power alarm            : Off
  Module initialization incomplete alarm : Off
  Module fault alarm                : Off
  PLD Flash initialization fault alarm : Off
  Power supply fault alarm          : Off
  Checksum fault alarm              : Off
  Tx laser disabled alarm           : Off
  Tx loss of signal functionality alarm : Off
  Tx CDR loss of lock alarm         : Off
  Rx loss of signal alarm           : Off
  Rx CDR loss of lock alarm         : Off
  Module temperature high alarm threshold : 75 degrees C / 167 degrees F
  Module temperature low alarm threshold : -5 degrees C / 23 degrees F
  Module temperature high warning threshold : 70 degrees C / 158 degrees F
  Module temperature low warning threshold : 0 degrees C / 32 degrees F
  Module voltage high alarm threshold : 3.5000 V
  Module voltage low alarm threshold : 3.0990 V
  Module voltage high warning threshold : 3.4000 V
  Module voltage low warning threshold : 3.2000 V
  Laser bias current high alarm threshold : 250.000 mA
  Laser bias current low alarm threshold : 37.500 mA
  Laser bias current high warning threshold : 225.000 mA
  Laser bias current low warning threshold : 50.000 mA
  Laser output power high alarm threshold : 3.9800 mW / 6.00 dBm
  Laser output power low alarm threshold : 0.4670 mW / -3.31 dBm
  Laser output power high warning threshold : 3.5480 mW / 5.50 dBm
  Laser output power low warning threshold : 0.5240 mW / -2.81 dBm
  Laser rx power high alarm threshold : 3.5481 mW / 5.50 dBm
  Laser rx power low alarm threshold : 0.0616 mW / -12.10 dBm
  Laser rx power high warning threshold : 3.1622 mW / 5.00 dBm
  Laser rx power low warning threshold : 0.0691 mW / -11.61 dBm
  Laser temperature high alarm threshold : 67 degrees C / 153 degrees F
  Laser temperature low alarm threshold : 35 degrees C / 95 degrees F
  Laser temperature high warning threshold : 62 degrees C / 144 degrees F
  Laser temperature low warning threshold : 40 degrees C / 104 degrees F
  SOA bias current high alarm threshold : 0.000 mA
  SOA bias current low alarm threshold : 0.000 mA
  SOA bias current high warning threshold : 0.000 mA
  SOA bias current low warning threshold : 0.000 mA
Lane 0
  Laser bias current                : 131.684 mA
  Laser output power                 : 1.002 mW / 0.01 dBm
  Laser temperature                  : 54 degrees C / 128 degrees F
  Laser receiver power               : 0.497 mW / -3.03 dBm

```

```

Laser bias current high alarm      : Off
Laser bias current low alarm       : Off
Laser bias current high warning    : Off
Laser bias current low warning     : Off
Laser output power high alarm      : Off
Laser output power low alarm       : Off
Laser output power high warning    : Off
Laser output power low warning     : Off
Laser temperature high alarm       : Off
Laser temperature low alarm        : Off
Laser temperature high warning     : Off
Laser temperature low warning      : Off
Laser receiver power high alarm    : Off
Laser receiver power low alarm     : Off
Laser receiver power high warning  : Off
Laser receiver power low warning   : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm          : Off
Rx loss of signal alarm            : Off
Rx CDR loss of lock alarm          : Off
APD supply fault alarm             : Off
TEC fault alarm                   : Off
Wavelength unlocked alarm          : Off

Lane 1
Laser bias current                 : 122.345 mA
Laser output power                 : 1.002 mW / 0.01 dBm
Laser temperature                 : 51 degrees C / 124 degrees F
Laser receiver power              : 0.611 mW / -2.14 dBm
Laser bias current high alarm      : Off
Laser bias current low alarm       : Off
Laser bias current high warning    : Off
Laser bias current low warning     : Off
Laser output power high alarm      : Off
Laser output power low alarm       : Off
Laser output power high warning    : Off
Laser output power low warning     : Off
Laser temperature high alarm       : Off
Laser temperature low alarm        : Off
Laser temperature high warning     : Off
Laser temperature low warning      : Off
Laser receiver power high alarm    : Off
Laser receiver power low alarm     : Off
Laser receiver power high warning  : Off
Laser receiver power low warning   : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm          : Off
Rx loss of signal alarm            : Off
Rx CDR loss of lock alarm          : Off
APD supply fault alarm             : Off
TEC fault alarm                   : Off
Wavelength unlocked alarm          : Off

Lane 2
Laser bias current                 : 112.819 mA
Laser output power                 : 1.000 mW / 0.00 dBm
Laser temperature                 : 50 degrees C / 122 degrees F
Laser receiver power              : 0.540 mW / -2.67 dBm
Laser bias current high alarm      : Off
Laser bias current low alarm       : Off
Laser bias current high warning    : Off
Laser bias current low warning     : Off
Laser output power high alarm      : Off

```

```

Laser output power low alarm           : Off
Laser output power high warning        : Off
Laser output power low warning         : Off
Laser temperature high alarm           : Off
Laser temperature low alarm            : Off
Laser temperature high warning         : Off
Laser temperature low warning          : Off
Laser receiver power high alarm        : Off
Laser receiver power low alarm         : Off
Laser receiver power high warning      : Off
Laser receiver power low warning       : Off
Tx loss of signal functionality alarm   : Off
Tx CDR loss of lock alarm              : Off
Rx loss of signal alarm                : Off
Rx CDR loss of lock alarm              : Off
APD supply fault alarm                 : Off
TEC fault alarm                       : Off
Wavelength unlocked alarm              : Off

Lane 3
Laser bias current                     : 100.735 mA
Laser output power                     : 1.002 mW / 0.01 dBm
Laser temperature                      : 50 degrees C / 122 degrees F
Laser receiver power                   : 0.637 mW / -1.96 dBm
Laser bias current high alarm          : Off
Laser bias current low alarm           : Off
Laser bias current high warning        : Off
Laser bias current low warning         : Off
Laser output power high alarm          : Off
Laser output power low alarm           : Off
Laser output power high warning        : Off
Laser output power low warning         : Off
Laser temperature high alarm           : Off
Laser temperature low alarm            : Off
Laser temperature high warning         : Off
Laser temperature low warning          : Off
Laser receiver power high alarm        : Off
Laser receiver power low alarm         : Off
Laser receiver power high warning      : Off
Laser receiver power low warning       : Off
Tx loss of signal functionality alarm   : Off
Tx CDR loss of lock alarm              : Off
Rx loss of signal alarm                : Off
Rx CDR loss of lock alarm              : Off
APD supply fault alarm                 : Off
TEC fault alarm                       : Off
Wavelength unlocked alarm              : Off

```

## show interfaces ge-

---

**Syntax** `show interfaces ge-fpc/pic/port`  
`<brief | detail | extensive | terse>`  
`<media>`  
`<statistics>`

**Release Information** Command introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Display status information about the specified Gigabit Ethernet interface.



**NOTE:** You must have a transceiver plugged into an SFP or SFP+ port before information about the interface can be displayed.

**Options** `ge-fpc/pic/port`—Display standard information about the specified Gigabit Ethernet interface.

`brief | detail | extensive | terse`—(Optional) Display the specified level of output.

`media`—(Optional) Display media-specific information about network interfaces.

`statistics`—(Optional) Display static interface statistics.

**Required Privilege Level** view

- Related Documentation**
- [Monitoring Interface Status and Traffic on page 2835](#)
  - [Troubleshooting Network Interfaces on EX3200 Switches](#)
  - [Troubleshooting Network Interfaces on EX4200 Switches](#)
  - [Troubleshooting an Aggregated Ethernet Interface on page 2931](#)
  - [Junos OS Ethernet Interfaces Configuration Guide](#)

**List of Sample Output** [show interfaces ge-0/0/0 on page 2885](#)  
[show interfaces ge-0/0/0 brief on page 2885](#)  
[show interfaces ge-0/0/0 brief \(with EEE Enabled on the EEE-capable Base-T copper Ethernet interfaces\) on page 2886](#)  
[show interfaces ge-0/0/0 detail on page 2886](#)  
[show interfaces ge-0/0/4 extensive on page 2887](#)

**Output Fields** [Table 307 on page 2879](#) lists the output fields for the `show interfaces ge-` command. Output fields are listed in the approximate order in which they appear.

Table 307: show interfaces ge- Output Fields

| Field Name                                    | Field Description                                                                                                                                                                                                     | Level of Output               |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| <b>Physical Interface</b>                     |                                                                                                                                                                                                                       |                               |
| <b>Physical interface</b>                     | Name of the physical interface.                                                                                                                                                                                       | All levels                    |
| <b>Enabled</b>                                | State of the interface: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                           | All levels                    |
| <b>Interface index</b>                        | Index number of the physical interface, which reflects its initialization sequence.                                                                                                                                   | <b>detail extensive none</b>  |
| <b>SNMP ifIndex</b>                           | SNMP index number for the physical interface.                                                                                                                                                                         | <b>detail extensive none</b>  |
| <b>Generation</b>                             | Unique number for use by Juniper Networks technical support only.                                                                                                                                                     | <b>detail extensive</b>       |
| <b>Description</b>                            | Optional user-specified description.                                                                                                                                                                                  | <b>brief detail extensive</b> |
| <b>Link-level type</b>                        | Encapsulation being used on the physical interface.                                                                                                                                                                   | All levels                    |
| <b>MTU</b>                                    | Maximum transmission unit size on the physical interface. Default is 1514.                                                                                                                                            | All levels                    |
| <b>Speed</b>                                  | Speed of the interface: Auto if autonegotiation of speed is enabled; speed in megabits per second if the interface speed is explicitly configured.                                                                    | All levels                    |
| <b>Duplex</b>                                 | Link mode of the interface: Auto if autonegotiation of link mode is enabled; Full-Duplex or Half-Duplex if the link mode is explicitly configured.                                                                    | All levels                    |
| <b>Loopback</b>                               | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                                                                                        | All levels                    |
| <b>Source filtering</b>                       | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                          | All levels                    |
| <b>Flow control</b>                           | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                              | All levels                    |
| <b>Auto-negotiation</b>                       | Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                           | All levels                    |
| <b>Remote-fault</b>                           | Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul> | All levels                    |
| <b>IEEE 802.3az Energy Efficient Ethernet</b> | IEEE 802.3az Energy Efficient Ethernet status: <b>Enabled</b> or <b>Disabled</b> (appears only for EEE-capable Base-T copper Ethernet interfaces).                                                                    | All levels                    |
| <b>Device flags</b>                           | Information about the physical device.                                                                                                                                                                                | All levels                    |
| <b>Interface flags</b>                        | Information about the interface.                                                                                                                                                                                      | All levels                    |
| <b>Link flags</b>                             | Information about the link.                                                                                                                                                                                           | All levels                    |
| <b>CoS queues</b>                             | Number of CoS queues configured.                                                                                                                                                                                      | <b>detail extensive none</b>  |

Table 307: show interfaces ge- Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output              |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b>      |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive none</b> |
| <b>Hardware address</b>        | MAC address of the hardware.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: <i>year-month-day hour:minute:second timezone (hour:minute:second ago)</i></b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive none</b> |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b>      |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled on the switch.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>      |
| <b>Input errors</b>            | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 sanity checks of the headers. For example, a frame with less than 20 bytes of available IP header is discarded.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>             |



Table 307: show interfaces ge- Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output              |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Output errors</b>                    | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>             |
| <b>Egress queues</b>                    | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>      |
| <b>Queue counters (Egress )</b>         | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>      |
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain time, it is promoted to an alarm. Based on the switch configuration, a defect can activate the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive none</b> |

Table 307: show interfaces ge- Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output  |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>MAC statistics</b>    | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Unicast packets</b>, <b>Broadcast packets</b>, and <b>Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of frames that exceed 1518 octets.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul> | <b>extensive</b> |
| <b>Filter Statistics</b> | Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b> |

Table 307: show interfaces ge- Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Autonegotiation information            | <p>Information about link autonegotiation:</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Complete</b>—The autonegotiation process between the local and remote Ethernet interfaces was successful.</li> <li>• <b>Incomplete</b>—Remote Ethernet interface has the speed or link mode configured or does not perform autonegotiation.</li> <li>• <b>No autonegotiation</b>—Local Ethernet interface has autonegotiation disabled and the link mode and speed are manually configured.</li> </ul> </li> <li>• <b>Link partner</b>—Information from the link partner: <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>. If the link mode of the remote device cannot be determined, the value is <b>Unknown</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, the types are: <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit); <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit); and <b>Symmetric/Asymmetric</b> (link partner supports <b>PAUSE</b> on both receive and transmit or <b>PAUSE</b> only on receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> <li>• <b>Link partner speed</b>—Speed of the link partner.</li> </ul> </li> <li>• <b>Local resolution</b>—Resolution of the autonegotiation process on the local interface: <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Type of flow control that is used by the local interface. For Gigabit Ethernet interfaces, the types are: <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit); <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit); and <b>Symmetric/Asymmetric</b> (link partner supports <b>PAUSE</b> on both receive and transmit or <b>PAUSE</b> only on receive).</li> <li>• <b>Link mode</b>—Link mode of local interface: either <b>Full-duplex</b> or <b>Half-duplex</b>. Displayed when <b>Negotiation status</b> is <b>Incomplete</b>.</li> <li>• <b>Local link speed</b>—Speed of the local interface. Displayed when <b>Negotiation status</b> is <b>Incomplete</b>.</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul> | extensive       |
| Packet Forwarding Engine configuration | <p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number: <ul style="list-style-type: none"> <li>• On standalone switches with built-in interfaces, the slot number refers to the switch itself and is always 0.</li> <li>• On Virtual Chassis composed of switches with built-in interfaces, the slot number refers to the member ID of the switch.</li> <li>• On switches with line cards or on Virtual Chassis composed of switches with line cards, the slot number refers to the line card slot number on the switch or Virtual Chassis.</li> </ul> </li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | extensive       |

#### Logical Interface

Table 307: show interfaces ge- Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output              |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Logical interface</b>       | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                        | All levels                   |
| <b>Index</b>                   | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>            | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>      |
| <b>Flags</b>                   | Information about the logical interface.                                                                                                                                                                                                                                                                                                                                                                              | All levels                   |
| <b>Encapsulation</b>           | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                               | All levels                   |
| <b>Protocol</b>                | Protocol family.                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b> |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.<br><br><i>NOTE:</i> For logical interfaces on EX Series switches, the traffic statistics fields in <b>show interfaces</b> commands show only control traffic; the traffic statistics do not include data traffic.                                                                                          | <b>detail extensive</b>      |
| <b>IPv6 transit statistics</b> | EX Series switches do not support the collection and reporting of IPv6 transit statistics.                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>             |
| <b>Local statistics</b>        | Number and rate of bytes and packets destined to and from the switch.                                                                                                                                                                                                                                                                                                                                                 | <b>extensive</b>             |
| <b>Transit statistics</b>      | Number and rate of bytes and packets transiting the switch.                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>             |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>      |
| <b>Route Table</b>             | Route table in which the logical interface address is located. For example, 0 refers to the routing table <b>inet.0</b> .                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Input Filters</b>           | Names of any input filters applied to this interface.                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b>      |
| <b>Output Filters</b>          | Names of any output filters applied to this interface.                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>      |
| <b>Flags</b>                   | Information about protocol family flags.<br><br>If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled. | <b>detail extensive</b>      |
| <b>protocol-family</b>         | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                                              | <b>brief</b>                 |
| <b>Flags</b>                   | Information about the address flags.                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive none</b> |

Table 307: show interfaces ge- Output Fields (*continued*)

| Field Name         | Field Description                                                 | Level of Output              |
|--------------------|-------------------------------------------------------------------|------------------------------|
| <b>Destination</b> | IP address of the remote side of the connection.                  | <b>detail extensive none</b> |
| <b>Local</b>       | IP address of the logical interface.                              | <b>detail extensive none</b> |
| <b>Broadcast</b>   | Broadcast address of the logical interlace.                       | <b>detail extensive none</b> |
| <b>Generation</b>  | Unique number for use by Juniper Networks technical support only. | <b>detail extensive</b>      |

## Sample Output

### show interfaces ge-0/0/0

```

user@switch> show interfaces ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Interface index: 129, SNMP ifIndex: 21
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:19:e2:50:3f:41, Hardware address: 00:19:e2:50:3f:41
  Last flapped   : 2008-01-16 11:40:53 UTC (4d 02:30 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms   : None
  Active defects  : None

Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 22)
  Flags: SNMP-Traps
  Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch
  Flags: None

```

### show interfaces ge-0/0/0 brief

```

user@switch> show interfaces ge-0/0/0 brief
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Description: voice priority and tcp and icmp traffic rate-limiting filter at i
  ngress port
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags      : None

Logical interface ge-0/0/0.0

```

Flags: Device-Down SNMP-Traps Encapsulation: ENET2  
eth-switch

#### show interfaces ge-0/0/0 brief (with IEEE Enabled on the IEEE-capable Base-T copper Ethernet interfaces)

```
user@switch> show interfaces ge-0/0/0 brief
Physical interface: ge-0/0/0, Enabled, Physical link is Up
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
Auto-negotiation: Enabled, Remote fault: Online,
IEEE 802.3az Energy Efficient Ethernet: Enabled, NO LPI
Device flags : Present Running
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags : None
```

#### show interfaces ge-0/0/0 detail

```
user@switch> show interfaces ge-0/0/0 detail
Physical interface: ge-0/0/0, Enabled, Physical link is Up
Interface index: 193, SNMP ifIndex: 206, Generation: 196
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
Last flapped : 2009-05-05 06:03:05 UTC (00:22:13 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 0 0 0
1 assured-forw 0 0 0
5 expedited-fo 0 0 0
7 network-cont 0 0 0

Active alarms : None
Active defects : None

Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
Flags: SNMP-Traps Encapsulation: ENET2
Bandwidth: 0
Traffic statistics:
Input bytes : 0
```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, Generation: 146, Route table: 0
Flags: Is-Primary
Input Filters: f1,
Output Filters: f2,,,,

```

### show interfaces ge-0/0/4 extensive

```

user@switch> show interfaces ge-0/0/4 extensive
Physical interface: ge-0/0/4, Enabled, Physical link is Up
Interface index: 165, SNMP ifIndex: 152, Generation: 168
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:33:65:44, Hardware address: 00:1f:12:33:65:44
Last flapped : 2008-09-17 11:02:25 UTC (16:32:54 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 2989761 984 bps
Input packets: 0 0 pps
Output packets: 24307 1 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 0 0 0
1 assured-forw 0 0 0
5 expedited-fo 0 0 0

```

```

7 network-cont                                0                24307                0

Active alarms : None
Active defects : None
MAC statistics:
    Receive          Transmit
Total octets        0          2989761
Total packets       0          24307
Unicast packets     0           0
Broadcast packets   0           0
Multicast packets   0          24307
CRC/Align errors    0           0
FIFO errors         0           0
MAC control frames  0           0
MAC pause frames    0           0
Oversized frames    0
Jabber frames       0
Fragment frames     0
Code violations      0

Autonegotiation information:
Negotiation status: Complete
Link partner:
    Link mode: Full-duplex, Flow control: None, Remote fault: OK,
    Link partner Speed: 1000 Mbps
Local resolution:
    Flow control: None, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 0
Direction : Output
CoS transmit queue      Bandwidth      Buffer Priority
Limit
    %      bps      %      usec
0 best-effort          95      950000000      95      NA      low
none
7 network-control      5       50000000      5      NA      low
none

Logical interface ge-0/0/4.0 (Index 82) (SNMP ifIndex 184) (Generation 147)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
Output bytes : 4107883
Input packets: 0
Output packets: 24307
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 4107883
Input packets: 0
Output packets: 24307
Transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0

```



```
Input packets:          0
Output packets:         0
Protocol eth-switch, Generation: 159, Route table: 0
Flags: None
Input Filters: f2,
Output Filters: f1,,,,
```

## show interfaces irb

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces irb &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance <i>instance-name</i>&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 12.3R2.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for the QFX Series</p>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Display integrated routing and bridging interfaces information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance <i>instance-name</i></b>—(Optional) Display information for the interface with the specified SNMP index.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the interface with the specified SNMP index.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Additional Information</b>   | Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another VLAN that has a Layer 3 protocol configured.                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces irb extensive on page 2894</a></p> <p><a href="#">show interfaces irb snmp-index on page 2895</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | <a href="#">Table 247 on page 2454</a> lists the output fields for the <b>show interfaces irb</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Table 308: show interfaces irb Output Fields**

| Field Name                | Field Description                                                                                                                             | Level of Output |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Physical Interface</b> |                                                                                                                                               |                 |
| <b>Physical interface</b> | Name of the physical interface.                                                                                                               | All levels      |
| <b>Enabled</b>            | State of the physical interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels      |

Table 308: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                          | Level of Output                    |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| <b>Proto</b>                   | Protocol configured on the interface.                                                                                                                                                                                                                      | <b>terse</b>                       |
| <b>Interface index</b>         | Physical interface index number, which reflects its initialization sequence.                                                                                                                                                                               | <b>detail extensive none</b>       |
| <b>SNMP ifIndex</b>            | SNMP index number for the physical interface.                                                                                                                                                                                                              | <b>detail extensive none</b>       |
| <b>Type</b>                    | Physical interface type.                                                                                                                                                                                                                                   | <b>detail extensive none</b>       |
| <b>Link-level type</b>         | Encapsulation being used on the physical interface.                                                                                                                                                                                                        | <b>detail extensive brief none</b> |
| <b>MTU</b>                     | MTU size on the physical interface.                                                                                                                                                                                                                        | <b>detail extensive brief none</b> |
| <b>Clocking</b>                | Reference clock source: <b>Internal</b> or <b>External</b> . Always unspecified on IRB interfaces.                                                                                                                                                         | <b>detail extensive brief</b>      |
| <b>Speed</b>                   | Speed at which the interface is running. Always unspecified on IRB interfaces.                                                                                                                                                                             | <b>detail extensive brief</b>      |
| <b>Device flags</b>            | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .                                                                                                         | <b>detail extensive brief none</b> |
| <b>Interface flags</b>         | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .                                                                                                            | <b>detail extensive brief none</b> |
| <b>Link type</b>               | Physical interface link type: <b>full duplex</b> or <b>half duplex</b> .                                                                                                                                                                                   | <b>detail extensive none</b>       |
| <b>Link flags</b>              | Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .                                                                                                                     | <b>detail extensive none</b>       |
| <b>Physical Info</b>           | Physical interface information.                                                                                                                                                                                                                            | All levels                         |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                        | <b>detail extensive</b>            |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                    | <b>detail extensive none</b>       |
| <b>Hardware address</b>        | MAC address of the hardware.                                                                                                                                                                                                                               | <b>detail extensive none</b>       |
| <b>Alternate link address</b>  | Backup address of the link.                                                                                                                                                                                                                                | <b>detail extensive</b>            |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> . | <b>detail extensive none</b>       |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                          | <b>detail extensive</b>            |

Table 308: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output         |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> |
| <b>IPv6 transit statistics</b> | <p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Input errors</b>            | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>           | <b>detail extensive</b> |
| <b>Output errors</b>           | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the DPC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>detail extensive</b> |

#### Logical Interface

Table 308: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output                 |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Logical interface</b>       | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels                      |
| <b>Index</b>                   | Index number of the logical interface (which reflects its initialization sequence).                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>            | SNMP interface index number of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b><br>none |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive</b>         |
| <b>Flags</b>                   | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b>         |
| <b>Encapsulation</b>           | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>         |
| <b>Bandwidth</b>               | Speed at which the interface is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b>         |
| <b>Routing Instance</b>        | Routing instance IRB is configured under.                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>         |
| <b>Bridging Domain</b>         | Bridging domain IRB is participating in.                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b>         |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the logical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                             | <b>detail extensive</b>         |
| <b>IPv6 transit statistics</b> | <p>Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b>         |
| <b>Local statistics</b>        | Statistics for traffic received from and transmitted to the Routing Engine.                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>         |
| <b>Transit statistics</b>      | Statistics for traffic transiting the router.                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b>         |
| <b>Protocol</b>                | Protocol family configured on the local interface. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>         |
| <b>MTU</b>                     | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b>         |
| <b>Maximum labels</b>          | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b><br>none |

Table 308: show interfaces irb Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                               | Level of Output         |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                                                               | <b>detail extensive</b> |
| <b>Route table</b>      | Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.                                             | <b>detail extensive</b> |
| <b>Addresses, Flags</b> | Information about address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .                 | <b>detail extensive</b> |
| <b>Policer</b>          | The policer that is to be evaluated when packets are received or transmitted on the interface.                                                                  | <b>detail extensive</b> |
| <b>Flags</b>            | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | <b>detail extensive</b> |

## Sample Output

### show interfaces irb extensive

```

user@host> show interfaces irb extensive
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0
    Output packets: 0
  IPv6 transit statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0
    Output packets: 0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Giants: 0, Policed discards:
0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0

```

```

Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Protocol inet, MTU: 1500, Generation: 154, Route table: 0
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255,
    Generation: 155
Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
  Flags: Is-Primary
  Policer: Input: __default_arp_policer

```

### show interfaces irb snmp-index

```

user@host> show interfaces irb snmp-index 25
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 25
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags : Present Running
  Interface flags: SNMP-Traps
  Link type : Full-Duplex
  Link flags : None
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Last flapped : Never
    Input packets : 0
    Output packets: 0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1500
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
  Protocol multiservice, MTU: 1500
    Flags: Is-Primary

```

## show interfaces mc-ae

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show interfaces mc-ae id <i>identifier</i> unit <i>number</i></b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | On peers with multi-chassis aggregated Ethernet ( <b>mc-aeX</b> ) interfaces, use this command to display information about the <b>mc-aeX</b> interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <b>identifier</b> —(Optional) Name of the multichassis aggregated Ethernet interface.<br><b>number</b> —(Optional) Specify the logical interface by unit number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Multichassis Link Aggregation (QFX Series Switches)</a></li> <li>• <a href="#">Understanding Multichassis Link Aggregation on page 2599 (EX Series Switches)</a></li> <li>• <a href="#">Configuring Multichassis Link Aggregation (QFX Series Switches)</a></li> <li>• <a href="#">Configuring Multichassis Link Aggregation on page 2681 (EX Series Switches)</a></li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation (QFX Series Switches)</a></li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation with Layer 3 MAC Address Synchronization (QFX Series Switches)</a></li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast using MAC Address Synchronization (QFX Series Switches)</a></li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol (VRRP) (QFX Series Switches)</a></li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using VRRP on EX9200 Switches (EX Series Switches)</a></li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP (QFX Series Switches)</a></li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on EX9200 Switches (EX Series Switches)</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show interfaces mc-ae on page 2897</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 309 on page 2896</a> lists the output fields for the <b>show interfaces mc-ae</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

**Table 309: show interfaces mc-ae Output Fields**

| Output Field Name                    | Field Description                                               |
|--------------------------------------|-----------------------------------------------------------------|
| <b>Current State Machine's State</b> | Specifies the state of the MC-LAG initialization state machine. |



Table 309: show interfaces mc-ae Output Fields (*continued*)

| Output Field Name        | Field Description                                                                                                                                                                                |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Member Link</b>       | Specifies the identifiers of the configured multichassis link aggregated interface members.                                                                                                      |
| <b>Local Status</b>      | Specifies the status of the local link: <b>active</b> or <b>standby</b> .                                                                                                                        |
| <b>Peer Status</b>       | Specifies the status of the peer link: <b>active</b> or <b>standby</b> .                                                                                                                         |
| <b>Peer State</b>        | Specifies the status of the local and peer links in an <b>active/active</b> MC-LAG configuration                                                                                                 |
| <b>Logical Interface</b> | Specifies the identifier and unit of the AE interface.                                                                                                                                           |
| <b>Topology Type</b>     | Specifies the bridge configured on the AE.                                                                                                                                                       |
| <b>Local State</b>       | Specifies if the local device is up or down.                                                                                                                                                     |
| <b>Peer State</b>        | Specifies if the peer device is up or down.                                                                                                                                                      |
| <b>Peer Ip/MCP/State</b> | Specifies the multichassis protection (MCP) link or the interchassis link-protection link (ICL-PL) for all of the multichassis aggregated Ethernet (MC-AE) interfaces that are part of the peer. |

## Sample Output

### show interfaces mc-ae

```

user@host> show interfaces mc-ae ae1 512
Member Link           : ae0
Current State Machine's State: mcae active state
Local Status          : active
Local State           : up
Peer Status           : active
Peer State            : up
    Logical Interface  : ae0.0
    Topology Type      : bridge
    Local State        : up
    Peer State         : up
    Peer Ip/MCP/State  : 3.3.3.2 ae1.0 up

```

## show interfaces me0

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces me0 &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Display status information about the management Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about the management Ethernet interface.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance</b>—(Optional) Display the name of the routing instance.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch on page 4791</a></li> <li>• <a href="#">Configuring Firewall Filters (CLI Procedure) on page 4804</a></li> </ul>                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces me0 on page 2902</a></p> <p><a href="#">show interfaces me0 brief on page 2902</a></p> <p><a href="#">show interfaces me0 detail on page 2902</a></p> <p><a href="#">show interfaces me0 extensive on page 2903</a></p>                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | <a href="#">Table 310 on page 2898</a> lists the output fields for the <b>show interfaces me0</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                |

**Table 310: show interfaces me0 Output Fields**

| Field Name                | Field Description                                                                   | Level of Output              |
|---------------------------|-------------------------------------------------------------------------------------|------------------------------|
| <b>Physical Interface</b> |                                                                                     |                              |
| <b>Physical interface</b> | Name of the physical interface.                                                     | All levels                   |
| <b>Enabled</b>            | State of the interface: <b>Enabled</b> or <b>Disabled</b> .                         | All levels                   |
| <b>Interface index</b>    | Index number of the physical interface, which reflects its initialization sequence. | <b>detail extensive none</b> |

Table 310: show interfaces me0 Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                   | Level of Output        |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| SNMP ifIndex            | SNMP index number for the physical interface.                                                                                                                                                                                                                       | detail extensive none  |
| Generation              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                   | detail extensive       |
| Description             | Optional user-specified description.                                                                                                                                                                                                                                | brief detail extensive |
| Type                    | Information about the type of functional interface.                                                                                                                                                                                                                 | All levels             |
| Link-level type         | Encapsulation being used on the physical interface.                                                                                                                                                                                                                 | All levels             |
| MTU                     | Maximum transmission unit size on the physical interface. The default is 1514.                                                                                                                                                                                      | All levels             |
| Clocking                | Interface that acts as a clock source. This field is not supported on EX Series switches and the default value is always <b>Unspecified</b> .                                                                                                                       | detail extensive       |
| Speed                   | Speed at which the interface is running.                                                                                                                                                                                                                            | All levels             |
| Device flags            | Information about the physical device.                                                                                                                                                                                                                              | All levels             |
| Interface flags         | Information about the interface.                                                                                                                                                                                                                                    | All levels             |
| Link type               | Information about whether the link is duplex and whether the negotiation is manual or automatic.                                                                                                                                                                    | detail extensive none  |
| Physical info           | Information about the device dependent physical interface selector. This field is applied only when a clocking option is specified. This field is not supported on EX Series switches and the default value is always <b>Unspecified</b> .                          | detail extensive       |
| Hold-times              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                 | detail extensive       |
| Current address         | Configured MAC address.                                                                                                                                                                                                                                             | detail extensive none  |
| Hardware address        | MAC address of the hardware.                                                                                                                                                                                                                                        | detail extensive none  |
| Alternate link address  | Information about alternate hardware address.                                                                                                                                                                                                                       | detail extensive       |
| Last flapped            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second timezone (weeksw:daysdhour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3w:3d 22:58 ago)</b> . | detail extensive none  |
| Statistics last cleared | Time when the statistics for the interface was last set to zero. The format is <b>Last flapped: year-month-day hour:minute:second timezone (weeksw:daysdhour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3w:3d 22:58 ago)</b> . | detail extensive       |

Table 310: show interfaces me0 Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output         |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <p>Following are fields in <b>Traffic statistics</b>:</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>IPv6 transit statistics</b> | <p>Number and rate of bytes and IPv6 packets received and transmitted on the physical interface.</p> <p>Following are fields in <b>IPv6 transit statistics</b>:</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Input errors</b>            | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and frame checksum (FCS) errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid FCS.</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of packets that exceed the size for the medium. For example, if the medium is Ethernet, the <b>Giant</b> field shows the count of packets with size greater than 1518 bytes.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>                                 | <b>extensive</b>        |
| <b>Output errors</b>           | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly. It increases only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increment quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>        |

Table 310: show interfaces me0 Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                           | Level of Output              |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------|------------------------------|
| Logical Interface              |                                                                                                                             |                              |
| <b>Logical interface</b>       | Name of the logical interface.                                                                                              | All levels                   |
| <b>Index</b>                   | Index number of the logical interface, which reflects its initialization sequence.                                          | <b>detail extensive</b> none |
| <b>SNMP ifIndex</b>            | SNMP interface index number for the logical interface.                                                                      | <b>detail extensive</b> none |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                           | <b>detail extensive</b>      |
| <b>Flags</b>                   | Information about the logical interface.                                                                                    | All levels                   |
| <b>Encapsulation</b>           | Encapsulation on the logical interface.                                                                                     | All levels                   |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.                  | <b>detail extensive</b>      |
| <b>IPv6 transit statistics</b> | If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface. | <b>detail extensive</b>      |
| <b>Local statistics</b>        | Number and rate of bytes and packets destined to and exiting from the switch.                                               | <b>extensive</b>             |
| <b>Protocol</b>                | Protocol family.                                                                                                            | <b>detail extensive</b> none |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                           | <b>detail extensive</b>      |
| <b>Route Table</b>             | Routing table in which the logical interface address is located. For example, 0 refers to the routing table <b>inet.0</b> . | <b>detail extensive</b>      |
| <b>Flags</b>                   | Information about protocol family flags.                                                                                    | <b>detail extensive</b>      |
| <b>Input Filter</b>            | Ingress filter name.                                                                                                        | <b>extensive</b>             |
| <b>Output Filter</b>           | Egress filter name.                                                                                                         | <b>extensive</b>             |
| <b>Addresses</b>               | Information about the management interface addresses.                                                                       | <b>detail extensive</b> none |
| <b>Flags</b>                   | Information about the address flags.                                                                                        | <b>detail extensive</b> none |
| <b>Destination</b>             | IP address of the remote side of the connection.                                                                            | <b>detail extensive</b> none |
| <b>Local</b>                   | IP address of the logical interface.                                                                                        | <b>detail extensive</b> none |
| <b>Broadcast</b>               | Broadcast address of the logical interface.                                                                                 | <b>detail extensive</b> none |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                           | <b>detail extensive</b>      |

## Sample Output

### show interfaces me0

```
user@switch> show interfaces me0
Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:1f:12:35:3c:bf, Hardware address: 00:1f:12:35:3c:bf
  Last flapped   : 2010-07-31 23:45:50 PDT (5d 00:32 ago)
    Input packets : 1661830
    Output packets: 3200

Logical interface me0.0 (Index 3) (SNMP ifIndex 34)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 1661830
  Output packets: 3200
  Protocol inet
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.204.32/20, Local: 10.204.33.103,
      Broadcast: 10.204.47.255
  Protocol inet6
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::21f:12ff:fe35:3cbf
```

### show interfaces me0 brief

```
user@switch> show interfaces me0 brief
Physical interface: me0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

Logical interface me0.0
  Flags: SNMP-Traps Encapsulation: ENET2
  inet 10.204.33.103/20
  inet6 fe80::21f:12ff:fe35:3cbf/64
```

### show interfaces me0 detail

```
user@switch> show interfaces me0 detail
Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info   : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:35:3c:bf, Hardware address: 00:1f:12:35:3c:bf
  Alternate link address: Unspecified
  Last flapped   : 2010-07-31 23:45:50 PDT (5d 00:37 ago)
  Statistics last cleared: Never
```

```

Traffic statistics:
Input bytes :          366663167
Output bytes :          498590
Input packets:         1664031
Output packets:         3259
IPv6 transit statistics:
Input bytes :          0
Output bytes :          0
Input packets:         0
Output packets:         0

Logical interface me0.0 (Index 3) (SNMP ifIndex 34) (Generation 1)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes :          366665637
Output bytes :          500569
Input packets:         1664048
Output packets:         3275
IPv6 transit statistics:
Input bytes :          0
Output bytes :          0
Input packets:         0
Output packets:         0
Local statistics:
Input bytes :          366665637
Output bytes :          500569
Input packets:         1664048
Output packets:         3275
Protocol inet, Generation: 1, Route table: 0
Flags: Is-Primary
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.204.32/20, Local: 10.204.33.103, Broadcast: 10.204.47.255,
Generation: 1
Protocol inet6, Generation: 2, Route table: 0
Flags: Is-Primary
Addresses, Flags: Is-Preferred
Destination: fe80::/64, Local: fe80::21f:12ff:fe35:3cbf
Generation: 2

```

#### show interfaces me0 extensive

```

user@switch> show interfaces me0 extensive
Physical interface: me0, Enabled, Physical link is Up
Interface index: 1, SNMP ifIndex: 33, Generation: 1
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
Device flags : Present Running
Interface flags: SNMP-Traps
Link type : Full-Duplex
Physical info : Unspecified
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:38:58:bf, Hardware address: 00:1f:12:38:58:bf
Alternate link address: Unspecified
Last flapped : 2010-08-15 06:27:33 UTC (03:06:22 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :          82310392
Output bytes :          1966952
Input packets:         110453
Output packets:         17747
IPv6 transit statistics:

```

```
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
Policed discards: 0, Resource errors: 0
Output errors:
Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0,
Resource errors: 0

Logical interface me0.0 (Index 3) (SNMP ifIndex 34) (Generation 1)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes : 82310392
Output bytes : 1966952
Input packets: 110453
Output packets: 17747
Local statistics:
Input bytes : 82310392
Output bytes : 1966952
Input packets: 110453
Output packets: 17747
Protocol inet, Generation: 1, Route table: 0
Flags: Is-Primary
Input Filters: mgmt_filter,
Addresses, Flags: Is-Default Is-Preferred Is-Primary
Destination: 10.204.96/20, Local: 10.204.96.234,
Broadcast: 10.204.111.255, Generation: 1
```



## show interfaces queue

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show interfaces queue<br><both-ingress-egress><br><egress><br><forwarding-class <i>forwarding-class</i> ><br><ingress><br>< <i>interface-name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Display class-of-service (CoS) queue information for physical interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>none</b>—Show detailed CoS queue statistics for all physical interfaces.</p> <p><b>both-ingress-egress</b>—(Optional) Show both ingress and egress queue statistics. (Ingress statistics are not available for all interfaces.)</p> <p><b>egress</b>—(Optional) Show egress queue statistics only.</p> <p><b>forwarding-class <i>forwarding-class</i></b>—(Optional) Show queue statistics only for the specified forwarding class.</p> <p><b>ingress</b>—(Optional) Show ingress queue statistics only. (Ingress statistics are not available for all interfaces.)</p> <p><b><i>interface-name</i></b>—(Optional) Show queue statistics for the specified interface.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring Interface Status and Traffic on page 2835</a></li> <li>• <a href="#">Monitoring Interfaces That Have CoS Components on page 2160</a></li> <li>• <a href="#">Defining CoS Schedulers and Scheduler Maps (CLI Procedure) on page 2109</a></li> <li>• <a href="#">Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure) on page 2126</a></li> </ul>                                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces queue ge-0/0/0 (EX2200 Switch) on page 2907</a></p> <p><a href="#">show interfaces queue xe-6/0/39 (Line Card with Oversubscribed Ports in an EX8200 Switch) on page 2908</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | Table 311 on page 2905 lists the output fields for the <b>show interfaces queue</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**Table 311: show interfaces queue Output Fields**

| Field Name                                                 | Field Description               |
|------------------------------------------------------------|---------------------------------|
| <b>Physical Interface and Forwarding Class Information</b> |                                 |
| Physical interface                                         | Name of the physical interface. |

Table 311: show interfaces queue Output Fields (*continued*)

| Field Name                                                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Enabled</b>                                                   | State of the interface. Possible values are: <ul style="list-style-type: none"> <li>• <b>Administratively down, Physical link is Down</b>—The interface is turned off, and the physical link is inoperable.</li> <li>• <b>Administratively down, Physical link is Up</b>—The interface is turned off, but the physical link is operational and can pass packets when it is enabled.</li> <li>• <b>Enabled, Physical link is Down</b>—The interface is turned on, but the physical link is inoperable and cannot pass packets.</li> <li>• <b>Enabled, Physical link is Up</b>—The interface is turned on, and the physical link is operational and can pass packets.</li> </ul> |
| <b>Interface index</b>                                           | Index number of the physical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>SNMP ifIndex</b>                                              | SNMP index number for the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                               | User-configured interface description.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Forwarding classes</b>                                        | Number of forwarding classes supported and in use for the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Ingress Queues Information (not shown for all interfaces)</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Ingress queues</b>                                            | Number of input queues supported and in use on the specified interface. For an interface on a line card with oversubscribed ports, the ingress queue handles low priority traffic on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Transmitted</b>                                               | Transmission statistics for the queue: <ul style="list-style-type: none"> <li>• <b>Packets</b>—Number of packets transmitted by this queue.</li> <li>• <b>Bytes</b>—Number of bytes transmitted by this queue.</li> <li>• <b>Tail-dropped packets</b>—Number of packets dropped because the queue buffers were full.</li> </ul>                                                                                                                                                                                                                                                                                                                                                |
| <b>PFE chassis queues</b>                                        | For an interface on a line card with oversubscribed ports, the number of Packet Forwarding Engine chassis queues supported and in use for the port group to which the interface belongs. The Packet Forwarding Engine chassis queue for a port group handles high priority traffic from all the interfaces in the port group.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Egress Queues Information</b>                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Egress queues</b>                                             | Number of output queues supported and in use on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Queue</b>                                                     | CoS queue number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Queued</b>                                                    | This counter is not supported on EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 311: show interfaces queue Output Fields (*continued*)

| Field Name                                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Transmitted</b>                             | <p>Number of packets and bytes transmitted by this queue. Information on transmitted packets and bytes can include:</p> <ul style="list-style-type: none"> <li>• <b>Packets</b>—Number of packets transmitted.</li> <li>• <b>Bytes</b>—Number of bytes transmitted.</li> <li>• <b>Tail-dropped packets</b>—Number of arriving packets dropped because output queue buffers were full.</li> <li>• <b>RED-dropped packets</b>—Number of packets dropped because of random early detection (RED). <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low loss priority packets dropped because of RED.</li> <li>• <b>High</b>—Number of high loss priority packets dropped because of RED.</li> </ul> </li> <li>• <b>RED-dropped bytes</b>—Number of bytes dropped because of random early detection (RED). <ul style="list-style-type: none"> <li>• <b>Low</b>—Number of low loss priority bytes dropped because of RED.</li> <li>• <b>High</b>—Number of high loss priority bytes dropped because of RED.</li> </ul> </li> </ul> |
| <b>Packet Forwarding Engine Chassis Queues</b> | <p>For an interface on a line card with oversubscribed ports, the number of Packet Forwarding Engine chassis queues supported and in use for the port group to which the interface belongs. The queue statistics reflect the traffic flowing on all the interfaces in the port group.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### show interfaces queue ge-0/0/0 (EX2200 Switch)

```

user@switch> show interfaces queue ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Interface index: 130, SNMP ifIndex: 501
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Transmitted:
      Packets           :                0
      Bytes             :                0
      Tail-dropped packets :                0
Queue: 1, Forwarding classes: assured-forwarding
  Queued:
    Transmitted:
      Packets           :                0
      Bytes             :                0
      Tail-dropped packets :                0
Queue: 5, Forwarding classes: expedited-forwarding
  Queued:
    Transmitted:
      Packets           :                0
      Bytes             :                0
      Tail-dropped packets :                0
Queue: 7, Forwarding classes: network-control
  Queued:
    Transmitted:
      Packets           :                0

```

```
Bytes : 0
Tail-dropped packets : 0
```

### show interfaces queue xe-6/0/39 (Line Card with Oversubscribed Ports in an EX8200 Switch)

```
user@switch> show interfaces queue xe-6/0/39

Physical interface: xe-6/0/39, Enabled, Physical link is Up
  Interface index: 291, SNMP ifIndex: 1641
  Forwarding classes: 16 supported, 7 in use
  Ingress queues: 1 supported, 1 in use
    Transmitted:
      Packets : 337069086018
      Bytes : 43144843010304
      Tail-dropped packets : 8003867575
  PFE chassis queues: 1 supported, 1 in use
    Transmitted:
      Packets : 0
      Bytes : 0
      Tail-dropped packets : 0
  Forwarding classes: 16 supported, 7 in use
  Egress queues: 8 supported, 7 in use
  Queue: 0, Forwarding classes: best-effort
    Queued:
      Transmitted:
        Packets : 334481399932
        Bytes : 44151544791024
        Tail-dropped packets : 0
    Queue: 1, Forwarding classes: assured-forwarding
      Queued:
        Transmitted:
          Packets : 0
          Bytes : 0
          Tail-dropped packets : 0
    Queue: 2, Forwarding classes: mcast-be
      Queued:
        Transmitted:
          Packets : 274948977
          Bytes : 36293264964
          Tail-dropped packets : 0
    Queue: 4, Forwarding classes: mcast-ef
      Queued:
        Transmitted:
          Packets : 0
          Bytes : 0
          Tail-dropped packets : 0
    Queue: 5, Forwarding classes: expedited-forwarding
      Queued:
        Transmitted:
          Packets : 0
          Bytes : 0
          Tail-dropped packets : 0
    Queue: 6, Forwarding classes: mcast-af
      Queued:
        Transmitted:
          Packets : 0
          Bytes : 0
          Tail-dropped packets : 0
    Queue: 7, Forwarding classes: network-control
      Queued:
        Transmitted:
```

```

Packets          :          46714
Bytes            :          6901326
Tail-dropped packets :          0

Packet Forwarding Engine Chassis Queues:
Queues: 8 supported, 7 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
  Transmitted:
    Packets          :          739338141426
    Bytes            :          94635282101928
    Tail-dropped packets :          0
    RED-dropped packets :          5606426444
      Low            :          5606426444
      High           :          0
    RED-dropped bytes :          683262846464
      Low            :          683262846464
      High           :          0
Queue: 1, Forwarding classes: assured-forwarding
  Queued:
  Transmitted:
    Packets          :          0
    Bytes            :          0
    Tail-dropped packets :          0
    RED-dropped packets :          0
      Low            :          0
      High           :          0
    RED-dropped bytes :          0
      Low            :          0
      High           :          0
Queue: 2, Forwarding classes: mcast-be
  Queued:
  Transmitted:
    Packets          :          0
    Bytes            :          0
    Tail-dropped packets :          0
    RED-dropped packets :          0
      Low            :          0
      High           :          0
    RED-dropped bytes :          0
      Low            :          0
      High           :          0
Queue: 4, Forwarding classes: mcast-ef
  Queued:
  Transmitted:
    Packets          :          0
    Bytes            :          0
    Tail-dropped packets :          0
    RED-dropped packets :          0
      Low            :          0
      High           :          0
    RED-dropped bytes :          0
      Low            :          0
      High           :          0
Queue: 5, Forwarding classes: expedited-forwarding
  Queued:
  Transmitted:
    Packets          :          0
    Bytes            :          0
    Tail-dropped packets :          0
    RED-dropped packets :          0

```

|                   |   |   |
|-------------------|---|---|
| Low               | : | 0 |
| High              | : | 0 |
| RED-dropped bytes | : | 0 |
| Low               | : | 0 |
| High              | : | 0 |

Queue: 6, Forwarding classes: mcast-af

Queued:

Transmitted:

|                      |   |   |
|----------------------|---|---|
| Packets              | : | 0 |
| Bytes                | : | 0 |
| Tail-dropped packets | : | 0 |
| RED-dropped packets  | : | 0 |
| Low                  | : | 0 |
| High                 | : | 0 |
| RED-dropped bytes    | : | 0 |
| Low                  | : | 0 |
| High                 | : | 0 |

Queue: 7, Forwarding classes: network-control

Queued:

Transmitted:

|                      |   |          |
|----------------------|---|----------|
| Packets              | : | 97990    |
| Bytes                | : | 14987506 |
| Tail-dropped packets | : | 0        |
| RED-dropped packets  | : | 0        |
| Low                  | : | 0        |
| High                 | : | 0        |
| RED-dropped bytes    | : | 0        |
| Low                  | : | 0        |
| High                 | : | 0        |

## show interfaces xe-

**Syntax** `show interfaces xe-fpc/pic/port`  
`<brief | detail | extensive | terse>`  
`<media>`  
`<statistics>`

**Release Information** Command introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Display status information about the specified 10-Gigabit Ethernet interface.



**NOTE:** You must have a transceiver plugged into an SFP+ or an XFP port before information about the interface can be displayed.



**NOTE:** On an EX Series switch, the traffic statistics for a LAG might vary slightly from the cumulative traffic statistics of the member interfaces of the LAG. This difference is more likely to be seen when the traffic is bursty in nature, and because the statistics are not fetched from the LAG and the members in the same instant. For accurate traffic statistics for a LAG, use the aggregated Ethernet counters.

**Options** `xe-fpc/pic/port`—Display standard information about the specified 10-Gigabit Ethernet interface.

**brief | detail | extensive | terse**—(Optional) Display the specified level of output.

**media**—(Optional) Display media-specific information about network interfaces. For 10-Gigabit Ethernet interfaces, using the **media** option does not provide you with new or additional information. The output is the same as when the **media** option is not used.

**statistics**—(Optional) Display static interface statistics. For 10-Gigabit Ethernet interfaces, using the **statistics** option does not provide you with new or additional information. The output is the same as when the **statistics** option is not used.

**Required Privilege Level** view

**Related Documentation**

- [Monitoring Interface Status and Traffic on page 2835](#)
- [Troubleshooting Network Interfaces on EX3200 Switches](#)
- [Troubleshooting Network Interfaces on EX4200 Switches](#)
- [Troubleshooting an Aggregated Ethernet Interface on page 2931](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

**List of Sample Output** [show interfaces xe-4/1/0 on page 2920](#)  
[show interfaces xe-0/1/0 brief on page 2921](#)  
[show interfaces xe-4/1/0 detail on page 2921](#)  
[show interfaces xe-6/0/39 extensive on page 2922](#)

**Output Fields** [Table 312 on page 2912](#) lists the output fields for the **show interfaces xe-** command. Output fields are listed in the approximate order in which they appear.

**Table 312: show interfaces xe- Output Fields**

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output                                                  |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| Fields for the Terse Output Level Only |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                  |
| <b>Interface</b>                       | Name of the physical or logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>terse</b>                                                     |
| <b>Admin</b>                           | Administrative state of the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>terse</b>                                                     |
| <b>Link</b>                            | State of the physical link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>terse</b>                                                     |
| <b>Proto</b>                           | Protocol family configured on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>terse</b>                                                     |
| <b>Local</b>                           | Local IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>terse</b>                                                     |
| <b>Remote</b>                          | Remote IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>terse</b>                                                     |
| Fields for the Physical Interface      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                  |
| <b>Physical interface</b>              | Name of the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>brief</b><br><b>detail</b><br><b>extensive</b><br><b>none</b> |
| <b>Enabled</b>                         | State of the interface. Can be one of the following: <ul style="list-style-type: none"> <li><b>Administratively down, Physical link is Down</b>—The interface is turned off, and the physical link is inoperable and cannot pass packets even when it is enabled.</li> <li><b>Administratively down, Physical link is Up</b>—The interface is turned off, but the physical link is operational and can pass packets when it is enabled.</li> <li><b>Enabled, Physical link is Down</b>—The interface is turned on, but the physical link is inoperable and cannot pass packets.</li> <li><b>Enabled, Physical link is Up</b>—The interface is turned on, and the physical link is operational and can pass packets.</li> </ul> | <b>brief</b><br><b>detail</b><br><b>extensive</b><br><b>none</b> |
| <b>Interface index</b>                 | Index number of the physical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail</b><br><b>extensive</b><br><b>none</b>                 |
| <b>SNMP ifIndex</b>                    | SNMP index number for the physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b><br><b>extensive</b><br><b>none</b>                 |
| <b>Generation</b>                      | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail</b><br><b>extensive</b>                                |



Table 312: show interfaces xe- Output Fields (*continued*)

| Field Name               | Field Description                                                                                                              | Level of Output                                           |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| <b>Description</b>       | User-configured interface description.                                                                                         | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>Link-level type</b>   | Encapsulation being used on the physical interface.                                                                            | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>MTU</b>               | Maximum transmission unit size on the physical interface.                                                                      | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>Speed</b>             | Speed at which the interface is running.                                                                                       | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>Duplex</b>            | Duplex mode of the interface.                                                                                                  | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>BPDU Error</b>        | Not supported on EX Series switches.                                                                                           | <b>detail</b><br><b>extensive</b><br>none                 |
| <b>MAC-REWRITE Error</b> | Not supported on EX Series switches.                                                                                           | <b>detail</b><br><b>extensive</b><br>none                 |
| <b>Loopback</b>          | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> . | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>Source filtering</b>  | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                   | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>Flow control</b>      | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                       | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>Device flags</b>      | Information about the physical device.                                                                                         | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |

Table 312: show interfaces xe- Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output                                           |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| <b>Interface flags</b>         | Information about the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>Link flags</b>              | Information about the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>CoS queues</b>              | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b><br><b>extensive</b><br>none                 |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b><br><b>extensive</b>                         |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b><br><b>extensive</b><br>none                 |
| <b>Hardware address</b>        | Hardware MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b><br><b>extensive</b><br>none                 |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <i>year-month-day hour:minute:second timezone (weekswdaysd hours:minutes:seconds ago)</i> . For example, 2008-01-16 10:52:40 UTC (3d 22:58 ago).                                                                                                                                                                                                                                                                                                                                            | <b>detail</b><br><b>extensive</b><br>none                 |
| <b>Input Rate</b>              | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | none                                                      |
| <b>Output Rate</b>             | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | none                                                      |
| <b>Statistics last cleared</b> | Date, time, and how long ago the statistics for the interface were cleared. The format is <i>year-month-day hour:minute:second timezone (weekswdaysd hours:minutes:seconds ago)</i> . For example, 2010-05-17 07:51:28 PDT (00:04:33 ago).                                                                                                                                                                                                                                                                                                                                 | <b>detail</b><br><b>extensive</b>                         |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li><b>Input bytes</b>—Number of bytes received on the interface and rate in bits per second.</li> <li><b>Output bytes</b>—Number of bytes transmitted on the interface and rate in bits per second.</li> <li><b>Input packets</b>—Number of packets received on the interface and rate in packets per second.</li> <li><b>Output packets</b>—Number of packets transmitted on the interface and rate in packets per second.</li> </ul> | <b>detail</b><br><b>extensive</b>                         |

Table 312: show interfaces xe- Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output     |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| IPv6 transit statistics | EX Series switches do not support the collection and reporting of IPv6 transit statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail<br>extensive |
| Input errors            | Input errors on the interface: <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | extensive           |

Table 312: show interfaces xe- Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output                   |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------|
| <b>Output errors</b>     | <p>Output errors on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. A 10-Gigabit Ethernet interface supports only full-duplex operation, so for 10-Gigabit Ethernet interfaces, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>                  |
| <b>Ingress queues</b>    | Number of CoS ingress queues supported on the specified interface. Displayed only for an interface on a line card with oversubscribed ports.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail</b><br><b>extensive</b> |
| <b>Egress queues</b>     | Number of CoS egress queues supported on the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail</b><br><b>extensive</b> |
| <b>PFE Egress queues</b> | Number of Packet Forwarding Engine egress queues shared by the interfaces in a port group. Displayed only for an interface on a line card with oversubscribed ports.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b><br><b>extensive</b> |
| <b>Queue counters</b>    | <p>Statistics for queues:</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets. This counter is not supported on EX switches and always contains 0.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail</b><br><b>extensive</b> |

Table 312: show interfaces xe- Output Fields (*continued*)

| Field Name                                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output                                           |
|-----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| <b>Active alarms and Active defects</b>       | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <p><b>detail</b><br/><b>extensive</b><br/><b>none</b></p> |
| <b>MAC statistics</b>                         | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets.</li> <li>• <b>Unicast packets</b>, <b>Broadcast packets</b>, and <b>Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of frames that exceed 1518 octets.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul> | <b>extensive</b>                                          |
| <b>Packet Forwarding Engine configuration</b> | <p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number: <ul style="list-style-type: none"> <li>• On standalone switches with built-in interfaces, the slot number refers to the switch itself and is always 0.</li> <li>• On Virtual Chassis composed of switches with built-in interfaces, the slot number refers to the member ID of the switch.</li> <li>• On switches with line cards or on Virtual Chassis composed of switches with line cards, the slot number refers to the line card slot number on the switch or Virtual Chassis.</li> </ul> </li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>extensive</b>                                          |

Table 312: show interfaces xe- Output Fields (*continued*)

| Field Name                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output                   |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------|
| <b>CoS Information</b>               | <p>Scheduler information for the CoS egress queues on the physical interface:</p> <ul style="list-style-type: none"> <li>• <b>Direction</b>—Queue direction, always <b>Output</b>.</li> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth</b>—Information about bandwidth allocated to the queue: <ul style="list-style-type: none"> <li>• <b>%</b>—Bandwidth allocated to the queue as a percentage</li> <li>• <b>bps</b>—Bandwidth allocated to the queue in bps</li> </ul> </li> <li>• <b>Buffer</b>—Information about buffer space allocated to the queue: <ul style="list-style-type: none"> <li>• <b>%</b>—Buffer space allocated to the queue as a percentage.</li> <li>• <b>usec</b>—Buffer space allocated to the queue in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> </ul> </li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> | <b>extensive</b>                  |
| <b>Fields for MACsec statistics</b>  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                                   |
| <b>Protected Packets</b>             | The number of packets sent from the interface that were secured using MACsec when encryption was disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail</b><br><b>extensive</b> |
| <b>Encrypted Packets</b>             | The number of packets sent from the interface that were secured and encrypted using MACsec.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail</b><br><b>extensive</b> |
| <b>Protected Bytes</b>               | The number of bytes sent from the interface that were secured using MACsec, but not encrypted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail</b><br><b>extensive</b> |
| <b>Encrypted Bytes</b>               | The number of packets sent from the interface that were secured and encrypted using MACsec.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail</b><br><b>extensive</b> |
| <b>Accepted Packets</b>              | <p>The number of received packets that have been accepted on the interface. A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check.</p> <p>This counter increments for traffic that is and is not encrypted using MACsec.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b><br><b>extensive</b> |
| <b>Validated Bytes</b>               | <p>The number of bytes that have been validated by the MACsec integrity check and received on the interface.</p> <p>This counter does not increment when MACsec encryption is disabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail</b><br><b>extensive</b> |
| <b>Decrypted Bytes</b>               | The number of bytes received on the interface that have been decrypted. An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACSec.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b><br><b>extensive</b> |
| <b>Fields for Logical Interfaces</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                                   |

Table 312: show interfaces xe- Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                            | Level of Output                                           |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| <b>Logical interface</b>  | Name of the logical interface.                                                                                                                                                                                                                                                                                               | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>Index</b>              | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                           | <b>detail</b><br><b>extensive</b><br>none                 |
| <b>SNMP ifIndex</b>       | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                       | <b>detail</b><br><b>extensive</b><br>none                 |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                            | <b>detail</b><br><b>extensive</b>                         |
| <b>Description</b>        | User-configured description of the interface.                                                                                                                                                                                                                                                                                | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>Flags</b>              | Information about the logical interface.                                                                                                                                                                                                                                                                                     | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>Encapsulation</b>      | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                      | <b>brief</b><br><b>detail</b><br><b>extensive</b><br>none |
| <b>Traffic statistics</b> | Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.<br><br><b>NOTE:</b> For logical interfaces on EX Series switches, the traffic statistics fields in <b>show interfaces</b> commands show only control traffic; the traffic statistics do not include data traffic. | <b>detail</b><br><b>extensive</b>                         |
| <b>Local statistics</b>   | Number and rate of bytes and packets destined to and from the switch.                                                                                                                                                                                                                                                        | <b>extensive</b>                                          |
| <b>Transit statistics</b> | Number and rate of bytes and packets transiting the switch.                                                                                                                                                                                                                                                                  | <b>extensive</b>                                          |
| <b>Protocol</b>           | Protocol family.                                                                                                                                                                                                                                                                                                             | <b>detail</b><br><b>extensive</b><br>none                 |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                            | <b>detail</b><br><b>extensive</b>                         |

Table 312: show interfaces xe- Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output                           |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| <b>Route Table</b>            | Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .                                                                                                                                                                                                                                                                                      | <b>detail</b><br><b>extensive</b><br>none |
| <b>Input Filters</b>          | Names of any input filters applied to this interface.                                                                                                                                                                                                                                                                                                                                                                 | <b>detail</b><br><b>extensive</b>         |
| <b>Output Filters</b>         | Names of any output filters applied to this interface.                                                                                                                                                                                                                                                                                                                                                                | <b>detail</b><br><b>extensive</b>         |
| <b>Flags</b>                  | Information about protocol family flags.<br><br>If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled. | <b>detail</b><br><b>extensive</b>         |
| <b>Addresses, Flags</b>       | Information about the address flags.                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b><br><b>extensive</b><br>none |
| <i><b>protocol-family</b></i> | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                                              | <b>brief</b>                              |
| <b>Flags</b>                  | Information about the address flags.                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b><br><b>extensive</b><br>none |
| <b>Destination</b>            | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b><br><b>extensive</b><br>none |
| <b>Local</b>                  | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b><br><b>extensive</b><br>none |
| <b>Broadcast</b>              | Broadcast address of the logical interlace.                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b><br><b>extensive</b><br>none |
| <b>Generation</b>             | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                     | <b>detail</b><br><b>extensive</b>         |

## Sample Output

show interfaces xe-4/1/0

```

user@switch show interfaces xe-4/1/0
Physical interface: xe-4/1/0, Enabled, Physical link is Up
Interface index: 387, SNMP ifIndex: 369

```



```

Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Current address: 00:23:9c:03:8e:70, Hardware address: 00:23:9c:03:8e:70
Last flapped   : 2009-05-12 08:01:04 UTC (00:13:44 ago)
Input rate     : 36432 bps (3 pps)
Output rate    : 0 bps (0 pps)
Active alarms  : None
Active defects : None

```

```

Logical interface xe-4/1/0.0 (Index 66) (SNMP ifIndex 417)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch
  Flags: None

```

#### show interfaces xe-0/1/0 brief

```

user@switch> show interfaces xe-0/1/0 brief
Physical interface: xe-0/1/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None

Logical interface xe-0/1/0.0
  Flags: SNMP-Traps Encapsulation: ENET2
  eth-switch

```

#### show interfaces xe-4/1/0 detail

```

user@switch> show interfaces xe-4/1/0 detail
Physical interface: xe-4/1/0, Enabled, Physical link is Up
  Interface index: 387, SNMP ifIndex: 369, Generation: 390
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:23:9c:03:8e:70, Hardware address: 00:23:9c:03:8e:70
  Last flapped   : 2009-05-12 08:01:04 UTC (00:13:49 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :          4945644          48576 bps
    Output bytes  :              0          0 bps
    Input packets :          3258          4 pps
    Output packets:              0          0 pps
  IPv6 transit statistics:
    Input bytes   :              0
    Output bytes  :              0
    Input packets :              0
    Output packets:              0

```

```

Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort              0              0              0

  1 assured-forw             0              0              0

  5 expedited-fo             0              0              0

  7 network-cont             0              0              0

Active alarms  : None
Active defects : None

Logical interface xe-4/1/0.0 (Index 66) (SNMP ifIndex 417) (Generation 158)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input bytes  :              0
  Output bytes :              0
  Input packets:              0
  Output packets:            0
Local statistics:
  Input bytes  :              0
  Output bytes :              0
  Input packets:              0
  Output packets:            0
Transit statistics:
  Input bytes  :              0              0 bps
  Output bytes :              0              0 bps
  Input packets:              0              0 pps
  Output packets:            0              0 pps
Protocol eth-switch, Generation: 174, Route table: 0
Flags: None
Input Filters: f1,
Output Filters: f2,,,,

```

#### show interfaces xe-6/0/39 extensive

```

user@switch> show interfaces xe-6/0/39 extensive
Physical interface: xe-6/0/39, Enabled, Physical link is Up
Interface index: 291, SNMP ifIndex: 1641, Generation: 316
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:19:e2:72:f2:88, Hardware address: 00:19:e2:72:f2:88
Last flapped   : 2010-05-13 14:49:43 PDT (1d 00:14 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes  :      49625962140160      4391057408 bps
  Output bytes :      47686985710805      4258984960 bps
  Input packets:      387702829264      4288139 pps
  Output packets:      372554570944      4159166 pps
IPv6 transit statistics:
  Input bytes  :              0
  Output bytes :              0
  Input packets:              0

```

```

Output packets:                0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 2 supported, 2 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets
  Low priority              0          336342805223      7986622358
  High priority             0                      0              0
Egress queues: 8 supported, 8 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0 best-effort              0          333760130103      0
  1 assured-forw              0                      0              0
  2 mcast-be                  0          274948977        0
  3 queue3                    0                      0              0
  4 mcast-ef                  0                      0              0
  5 expedited-fo              0                      0              0
  6 mcast-af                  0                      0              0
  7 network-cont              0          46613          0
PFE Egress queues: 8 supported, 8 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0 best-effort              0          737867061290    5595302082
  1 assured-forw              0                      0              0
  2 mcast-be                  0                      0              0
  3 queue3                    0                      0              0
  4 mcast-ef                  0                      0              0
  5 expedited-fo              0                      0              0
  6 mcast-af                  0                      0              0
  7 network-cont              0          97800          0
Active alarms : None
Active defects : None
MAC statistics:
  Receive      Transmit
  Total octets  49625962140160  47686985710805
  Total packets 387702829264    372554570944
  Unicast packets 387702829264    372554518472
  Broadcast packets 0          2
  Multicast packets 0          52470
  CRC/Align errors 0          0
  FIFO errors      0          0
  MAC control frames 0          0
  MAC pause frames 0          0
  Oversized frames 0
  Jabber frames    0
  Fragment frames  0
  Code violations   0
Packet Forwarding Engine configuration:
  Destination slot: 6
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority  Limit
                           %      bps      %      usec
  0 best-effort            75      7500000000    75      0      low  none
  2 mcast-be                20      2000000000    20      0      low  none
  7 network-cont            5       500000000     5       0      low  none

```

Logical interface xe-6/0/39.0 (Index 1810) (SNMP ifIndex 2238) (Generation 1923)

```
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes :          0
  Output bytes :        9375416
  Input packets:          0
  Output packets:       48901
Local statistics:
  Input bytes :          0
  Output bytes :        9375416
  Input packets:          0
  Output packets:       48901
Transit statistics:
  Input bytes :          0          0 bps
  Output bytes :          0          0 bps
  Input packets:          0          0 pps
  Output packets:          0          0 pps
Protocol eth-switch, Generation: 1937, Route table: 0
  Flags: Trunk-Mode
```

## show lacp interfaces

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show lacp interfaces<br><interface-name>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet or Gigabit Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p>none—Display LACP information for all interfaces.</p> <p><i>interface-name</i>—(Optional) Display LACP information for the specified interface:</p> <ul style="list-style-type: none"> <li>• Aggregated Ethernet—<i>aex</i></li> <li>• Gigabit Ethernet—<i>ge-fpc/pic/port</i></li> <li>• 10-Gigabit Ethernet—<i>xe-fpc/pic/port</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li> <li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li> <li>• <i>Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch</i></li> <li>• <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure) on page 2667</a></li> <li>• <a href="#">Configuring Link Aggregation</a></li> <li>• <a href="#">Configuring Aggregated Ethernet LACP (CLI Procedure) on page 2671</a></li> <li>• <a href="#">Configuring Aggregated Ethernet LACP</a></li> <li>• <a href="#">Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 2672</a></li> <li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2582</a></li> <li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP</a></li> <li>• <a href="#">Junos OS Interfaces Fundamentals Configuration Guide</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show lacp interfaces (EX Series Switches) on page 2927</a><br><a href="#">show lacp interfaces (QFX Series) on page 2928</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 313 on page 2926</a> lists the output fields for the <b>show lacp interfaces</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

Table 313: show lacp interfaces Output Fields

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aggregated interface | Aggregated Ethernet interface name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| LACP State           | <p>LACP state information for each aggregated Ethernet interface:</p> <ul style="list-style-type: none"> <li>For a child interface configured with the <b>force-up</b> statement, LACP state displays <b>FUP</b> along with the interface name.</li> <li><b>Role</b>—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> <li><b>Actor</b>—Local device participating in the LACP negotiation.</li> <li><b>Partner</b>—Remote device participating in the LACP negotiation.</li> </ul> </li> <li><b>Exp</b>—Expired state. <b>Yes</b> indicates that the actor or partner is in an expired state. <b>No</b> indicates that the actor or partner is not in an expired state.</li> <li><b>Def</b>—Default. <b>Yes</b> indicates that the actor's receive machine is using the default operational partner information, which is administratively configured for the partner. <b>No</b> indicates that the operational partner information in use has been received in an LACP PDU.</li> <li><b>Dist</b>—Distribution of outgoing frames. <b>No</b> indicates that the distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is <b>Yes</b>.</li> <li><b>Col</b>—Collection of incoming frames. <b>Yes</b> indicates that the collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is <b>No</b>.</li> <li><b>Syn</b>—Synchronization. If the value is <b>Yes</b>, the link is considered to be synchronized. The link has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is <b>No</b>, the link is not synchronized. The link is currently not in the right aggregation.</li> <li><b>Aggr</b>—Ability of the aggregation port to aggregate (<b>Yes</b>) or to operate only as an individual link (<b>No</b>).</li> <li><b>Timeout</b>—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or a fast transmission rate, depending upon the expressed LACP timeout preference (<b>Long Timeout</b> or <b>Short Timeout</b>).</li> <li><b>Activity</b>—Actor's or partner's port activity. <b>Passive</b> indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is <b>Active</b>. <b>Active</b> indicates the port's preference to participate in the protocol regardless of the partner's control value.</li> </ul> |

Table 313: show lacp interfaces Output Fields (*continued*)

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LACP Protocol | <p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> <li>Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.</li> <li><b>Receive State</b>—One of the following values: <ul style="list-style-type: none"> <li><b>Current</b>—The state machine receives an LACP PDU and enters the <b>Current</b> state.</li> <li><b>Defaulted</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires a second time, the state machine enters the <b>Defaulted</b> state.</li> <li><b>Expired</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires once, the state machine enters the <b>Expired</b> state.</li> <li><b>Initialize</b>—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the <b>Initialize</b> state.</li> <li><b>LACP Disabled</b>—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to <b>LACP Disabled</b>. This state is similar to the <b>Defaulted</b> state, except that the port is forced to operate as an individual port.</li> <li><b>Port Disabled</b>—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the <b>Port Disabled</b> state.</li> </ul> </li> <li><b>Transmit State</b>—Transmit state of the state machine. The transmit state is one of the following values: <ul style="list-style-type: none"> <li><b>Fast periodic</b>—Periodic transmissions are enabled at a fast transmission rate.</li> <li><b>No periodic</b>—Periodic transmissions are disabled.</li> <li><b>Periodic timer</b>—Transitory state entered when the periodic timer expires.</li> <li><b>Slow periodic</b>—Periodic transmissions are enabled at a slow transmission rate.</li> </ul> </li> <li><b>Mux State</b>—State of the multiplexer state machine for the aggregation port. The state is one of the following values: <ul style="list-style-type: none"> <li><b>Attached</b>—The multiplexer state machine initiates the process of attaching the port to the selected aggregator.</li> <li><b>Collecting—Yes</b> indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. <b>No</b> indicates the receive function of this link is not enabled.</li> <li><b>Collecting distributing</b>—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.</li> <li><b>Detached</b>—Process of detaching the port from the aggregator is in progress.</li> <li><b>Distributing—Yes</b> indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames can be passed down from the aggregator's distribution function for transmission. <b>No</b> indicates the transmit function of this link is not enabled.</li> <li><b>Waiting</b>—The multiplexer state machine is in a holding process, awaiting an outcome.</li> </ul> </li> </ul> |

## Sample Output

### show lacp interfaces (EX Series Switches)

```

user@switch> show lacp interfaces ae5
Aggregated interface: ae5
  LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
    xe-2/0/7      Actor  No   No   Yes  Yes  Yes  Yes    Fast    Active
    xe-2/0/7      Partner No   No   Yes  Yes  Yes  Yes    Fast    Passive

```

|          |         |    |    |    |     |     |     |      |         |
|----------|---------|----|----|----|-----|-----|-----|------|---------|
| xe-4/0/7 | Actor   | No | No | No | No  | No  | Yes | Fast | Active  |
| xe-4/0/7 | Partner | No | No | No | Yes | Yes | Yes | Fast | Passive |

|                    |               |                |                         |
|--------------------|---------------|----------------|-------------------------|
| LACP protocol:     | Receive State | Transmit State | Mux State               |
| xe-2/0/7(Active)   | Current       | Fast periodic  | Collecting distributing |
| xe-34/0/7(Standby) | Current       | Fast periodic  | Waiting                 |

### show lacp interfaces (QFX Series)

```
user@switch> show lacp interfaces nodegroup1:ae0 extensive
```

```
Aggregated interface: nodegroup1:ae0
```

| LACP state:       | Role    | Exp | Def | Dist | Col | Syn | Aggr | Timeout | Activity |
|-------------------|---------|-----|-----|------|-----|-----|------|---------|----------|
| node1:xe-0/0/1FUP | Actor   |     | No  | Yes  | No  | No  | No   | Yes     | Fast     |
| Active            |         |     |     |      |     |     |      |         |          |
| node1xe-0/0/1FUP  | Partner |     | No  | Yes  | No  | No  | No   | Yes     | Fast     |
| Passive           |         |     |     |      |     |     |      |         |          |
| node2:xe-0/0/2    | Actor   |     | No  | Yes  | No  | No  | No   | Yes     | Fast     |
| Active            |         |     |     |      |     |     |      |         |          |
| node2:xe-0/0/2    | Partner |     | No  | Yes  | No  | No  | No   | Yes     | Fast     |
| Passive           |         |     |     |      |     |     |      |         |          |



|              | LACP protocol:           | Receive State | Transmit State | Mux State  |
|--------------|--------------------------|---------------|----------------|------------|
|              | node1:xe-0/0/1FUP        | Current       | Fast periodic  | Collecting |
| distributing | node2:xe-0/0/2           | Current       | Fast periodic  | Collecting |
| distributing | node1:xe-0/0/1 (active)  | Current       | Fast periodic  | Collecting |
| distributing | node2:xe-0/0/2 (standby) | Current       | Fast periodic  | WAITING    |

## test interface restart-auto-negotiation

---

|                                 |                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | test interface restart-auto-negotiation <i>interface-name</i>                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 7.6.<br>Command introduced in Junos OS Release 9.0 for EX Series switches. |
| <b>Description</b>              | Restarts auto-negotiation on a Fast Ethernet or Gigabit Ethernet interface.                                       |
| <b>Options</b>                  | <i>interface-name</i> —Interface name: <i>fe-fpc/pic/port</i> or <i>ge-fpc/pic/port</i> .                         |
| <b>Required Privilege Level</b> | view                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">test interface restart-auto-negotiation on page 2930</a>                                              |
| <b>Output Fields</b>            | Use the <code>show interfaces extensive</code> command to see the state for auto-negotiation.                     |

## Sample Output

### test interface restart-auto-negotiation

```
user@host> test interface restart-auto-negotiation fe-1/0/0
```

# Troubleshooting Procedures

- [Troubleshooting an Aggregated Ethernet Interface on page 2931](#)
- [Troubleshooting Interface Configuration and Cable Faults on page 2932](#)
- [Troubleshooting Unicast RPF on page 2933](#)
- [Diagnosing a Faulty Twisted-Pair Cable \(CLI Procedure\) on page 2934](#)

## Troubleshooting an Aggregated Ethernet Interface

---

Troubleshooting issues for aggregated Ethernet interfaces:

- [Show Interfaces Command Shows the LAG is Down on page 2931](#)
- [Logical Interface Statistics Do Not Reflect All Traffic on page 2931](#)
- [IPv6 Interface Traffic Statistics Are Not Supported on page 2932](#)
- [SNMP Counters ifHCInBroadcastPkts and ifInBroadcastPkts Are Always 0 on page 2932](#)

### Show Interfaces Command Shows the LAG is Down

**Problem**    **Description:** The `show interfaces terse` command shows that the LAG is down.

**Solution**    Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet—switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch (or the same Virtual Chassis).

### Logical Interface Statistics Do Not Reflect All Traffic

**Problem**    **Description:** The traffic statistics for a logical interface do not include all of the traffic.

**Solution** Traffic statistics fields for logical interfaces in **show interfaces** commands show only control traffic; the traffic statistics do not include data traffic. You can view the statistics for all traffic only per physical interface.

## IPv6 Interface Traffic Statistics Are Not Supported

**Problem** **Description:** The IPv6 transit statistics in the **show interfaces** command display all 0 values.

**Solution** EX Series switches do not support the collection and reporting of IPv6 transit statistics.

## SNMP Counters ifHCInBroadcastPkts and ifInBroadcastPkts Are Always 0

**Problem** **Description:** The values for the SNMP counters ifHCInBroadcastPkts and ifInBroadcastPkts are always 0.

**Solution** The SNMP counters ifHCInBroadcastPkts and ifInBroadcastPkts are not supported for aggregated Ethernet interfaces on EX Series switches.

**Related Documentation**

- [Verifying the Status of a LAG Interface on page 2837](#)
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

## Troubleshooting Interface Configuration and Cable Faults

---



**NOTE:** This topic applies only to the J-Web Application package.

Troubleshooting interface configuration and connectivity on the EX Series switch:

1. [Interface Configuration or Connectivity Is Not Working on page 2932](#)

## Interface Configuration or Connectivity Is Not Working

**Problem** **Description:**



**NOTE:** This topic applies only to the J-Web Application package.

You encounter errors when you attempt to configure an interface on the switch, or the interface is exhibiting connectivity problems.

**Solution** Use the port troubleshooter feature in the J-Web interface to identify and rectify port configuration and connectivity related problems.

To use the J-Web interface port troubleshooter:

1. Select the option **Troubleshoot** from the main menu.
2. Click **Troubleshoot Port**. The Port Troubleshooting wizard is displayed. Click **Next**.
3. Select the ports to troubleshoot.
4. Select the test cases to be executed on the selected port. Click **Next**.

When the selected test cases are executed, the final result and the recommended action is displayed.

If there is a cable fault, the port troubleshooter displays details and the recommended action. For example, the cable must be replaced.

If the port configuration needs to be modified, the port troubleshooter displays details and the recommended action.

**Related Documentation**

- [Monitoring Interface Status and Traffic on page 2835](#)
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 2619](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 2616](#)
- [Connecting and Configuring an EX Series Switch \(CLI Procedure\)](#)
- [Connecting and Configuring an EX Series Switch \(J-Web Procedure\)](#)

## Troubleshooting Unicast RPF

Troubleshooting issues for unicast reverse-path forwarding (RPF) on EX Series switches include:

1. [Legitimate Packets Are Discarded on page 2933](#)

### Legitimate Packets Are Discarded

**Problem** **Description:** The switch filters valid packets from legitimate sources, which results in the switch's discarding packets that should be forwarded.

**Solution** The interface or interfaces on which legitimate packets are discarded are asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, so the interface that receives a packet is not the same interface the switch uses to reply to the packet's source.

Unicast RPF works properly only on symmetrically routed interfaces. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. Unicast RPF filters packets by checking the forwarding table

for the best return path to the source of an incoming packet. If the best return path uses the same interface as the interface that received the packet, the switch forwards the packet. If the best return path uses a different interface than the interface that received the packet, the switch discards the packet.



**NOTE:** On EX3200, EX4200, and EX4300 switches, unicast RPF works properly only if all switch interfaces—including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs)—are symmetrically routed, because unicast RPF is enabled globally on all switch interfaces.

- Related Documentation**
- [Verifying Unicast RPF Status on page 2841](#)
  - [Understanding Unicast RPF on page 2592](#)

---

## Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure)

---

**Problem**    **Description:** A 10/100/1000BASE-T Ethernet interface has connectivity problems that you suspect might be caused by a faulty cable.

**Solution**    Use the time domain reflectometry (TDR) test to determine whether a twisted-pair Ethernet cable is faulty.

The TDR test:

- Detects and reports faults for each twisted pair in an Ethernet cable. Faults detected include open circuits, short circuits, and impedance mismatches.
- Reports the distance to fault to within 1 meter.
- Detects and reports pair swaps, pair polarity reversals, and excessive pair skew.

The TDR test is supported on the following switches and interfaces:

- EX2200, EX3200, EX3300, and EX4200 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.
- EX6200 and EX8200 switches—RJ-45 network interfaces on line cards.



**NOTE:** We recommend running the TDR test on an interface when there is no traffic on the interface.

To diagnose a cable problem by running the TDR test:

1. Run the `request diagnostics tdr` command.

```
user@switch> request diagnostics tdr start interface ge-0/0/10
```

```
Interface TDR detail:
```

```
Test status : Test successfully executed ge-0/0/10
```

2. View the results of the TDR test with the `show diagnostics tdr` command.

```
user@switch> show diagnostics tdr interface ge-0/0/10
```

```
Interface TDR detail:
```

```
Interface name : ge-0/0/10
```

```
Test status : Passed
```

```
Link status : Down
```

```
MDI pair : 1-2
```

```
  Cable status : Normal
```

```
  Distance fault : 0 Meters
```

```
  Polartiy swap : N/A
```

```
  Skew time : N/A
```

```
MDI pair : 3-6
```

```
  Cable status : Normal
```

```
  Distance fault : 0 Meters
```

```
  Polartiy swap : N/A
```

```
  Skew time : N/A
```

```
MDI pair : 4-5
```

```
  Cable status : Open
```

```
  Distance fault : 1 Meters
```

```
  Polartiy swap : N/A
```

```
  Skew time : N/A
```

```
MDI pair : 7-8
```

```
  Cable status : Normal
```

```
  Distance fault : 0 Meters
```

```
  Polartiy swap : N/A
```

```
  Skew time : N/A
```

```
Channel pair : 1
```

```
  Pair swap : N/A
```

```
Channel pair : 2
```

```
  Pair swap : N/A
```

```
Downshift : N/A
```

3. Examine the **Cable status** field for the four MDI pairs to determine if the cable has a fault. In the preceding example, the twisted pair on pins 4 and 5 is broken or cut at approximately one meter from the `ge-0/0/10` port connection.



**NOTE:** The **Test Status** field indicates the status of the TDR test, not the cable. The value **Passed** means the test completed—it does not mean that the cable has no faults.

The following is additional information about the TDR test:

- The TDR test can take some seconds to complete. If the test is still running when you execute the `show diagnostics tdr` command, the **Test status** field displays **Started**. For example:

```
user@switch> show diagnostics tdr interface ge-0/0/22
```

```
Interface TDR detail:
```

```
Interface name      : ge-0/0/22
Test status        : Started
```

- You can terminate a running TDR test before it completes by using the **request diagnostics tdr abort interface *interface-name*** command. The test terminates with no results, and the results from any previous test are cleared.
- You can display summary information about the last TDR test results for all interfaces on the switch that support the TDR test by not specifying an interface name with the **show diagnostics tdr** command. For example:

```
user@switch> show diagnostics tdr
```

| Interface | Test status | Link status | Cable status | Max distance | fault |
|-----------|-------------|-------------|--------------|--------------|-------|
| ge-0/0/0  | Passed      | UP          | OK           | 0            |       |
| ge-0/0/1  | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/2  | Passed      | UP          | OK           | 0            |       |
| ge-0/0/3  | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/4  | Passed      | UP          | OK           | 0            |       |
| ge-0/0/5  | Passed      | UP          | OK           | 0            |       |
| ge-0/0/6  | Passed      | UP          | OK           | 0            |       |
| ge-0/0/7  | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/8  | Passed      | Down        | OK           | 0            |       |
| ge-0/0/9  | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/10 | Passed      | Down        | Fault        | 1            |       |
| ge-0/0/11 | Passed      | UP          | OK           | 0            |       |
| ge-0/0/12 | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/13 | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/14 | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/15 | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/16 | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/17 | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/18 | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/19 | Passed      | Down        | OK           | 0            |       |
| ge-0/0/20 | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/21 | Not Started | N/A         | N/A          | N/A          |       |
| ge-0/0/22 | Passed      | UP          | OK           | 0            |       |
| ge-0/0/23 | Not Started | N/A         | N/A          | N/A          |       |

#### Related Documentation

- [Troubleshooting Interface Configuration and Cable Faults on page 2932](#)
- [request diagnostics tdr on page 2854](#)
- [show diagnostics tdr on page 2856](#)



## PART 15

# BGP

- [Overview on page 2939](#)
- [Configuration on page 2943](#)
- [Administration on page 3035](#)



## CHAPTER 48

# Overview

- [Layer 3 Protocols on page 2939](#)

## Layer 3 Protocols

---

- [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)
- [Layer 3 Protocols Not Supported on EX Series Switches on page 2940](#)

## Layer 3 Protocols Supported on EX Series Switches

EX Series switches support the Junos OS Layer 3 features and configuration statements listed in [Table 314 on page 2939](#):

**Table 314: Supported Junos OS Layer 3 Protocol Statements and Features**

| Protocol           | Notes                                                                                                                                            | For More Information                                             |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| BGP                | Fully supported.                                                                                                                                 | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| BFD                | Fully supported.                                                                                                                                 | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| ICMP               | Fully supported.                                                                                                                                 | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| IGMPv1, v2, and v3 | Fully supported.                                                                                                                                 | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |
| IS-IS              | Supported, with the exceptions noted in " <a href="#">Layer 3 Protocols Not Supported on EX Series Switches</a> " on <a href="#">page 2940</a> . | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| MLD                | Fully supported (MLD versions 1 and 2).                                                                                                          | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |
| MPLS               | Supported, with the exceptions noted in " <a href="#">Layer 3 Protocols Not Supported on EX Series Switches</a> " on <a href="#">page 2940</a> . | <a href="#">Junos OS MPLS Applications Configuration Guide</a>   |
| OSPFv1, v2 and v3  | Supported, with the exceptions noted in " <a href="#">Layer 3 Protocols Not Supported on EX Series Switches</a> " on <a href="#">page 2940</a> . | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| PIM                | Fully supported.                                                                                                                                 | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |

Table 314: Supported Junos OS Layer 3 Protocol Statements and Features (*continued*)

| Protocol | Notes                                                                                                | For More Information                                                                                                     |
|----------|------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| PPM      | Supported. See <i>EX Series Switch Software Features Overview</i> for specific platform information. | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| RIP      | Fully supported.                                                                                     | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| RIPng    | Fully supported.                                                                                     | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| SNMP     | Fully supported.                                                                                     | <a href="#">Junos OS Network Management Configuration Guide</a>                                                          |
| VRRP     | Fully supported.                                                                                     | See “Understanding VRRP on EX Series Switches” on page 2501. See also <a href="#">Junos OS High Availability Guide</a> . |

- Related Documentation**
- [Layer 3 Protocols Not Supported on EX Series Switches](#) on page 2940
  - [EX Series Switch Software Features Overview](#)

## Layer 3 Protocols Not Supported on EX Series Switches

EX Series switches do not support the Junos OS Layer 3 protocols and features listed in [Table 315 on page 2940](#):

Table 315: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported

| Feature                                                                                                         | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DVMRP                                                                                                           | <ul style="list-style-type: none"> <li>• <b>dvmrp</b> and subordinate statements</li> </ul>                                                                                                                                                                                                          |
| Flow aggregation (cflowd)                                                                                       | <ul style="list-style-type: none"> <li>• <b>cflow</b> and subordinate statements</li> </ul>                                                                                                                                                                                                          |
| IPsec                                                                                                           | <ul style="list-style-type: none"> <li>• <b>[edit services]</b> statements related to IPsec</li> </ul>                                                                                                                                                                                               |
| IS-IS: <ul style="list-style-type: none"> <li>• ES-IS</li> <li>• IPv6 in multicast routing protocols</li> </ul> | <ul style="list-style-type: none"> <li>• <b>clns-routing</b> statement</li> <li>• <b>ipv6-multicast</b> statement</li> <li>• <b>lsp-interval</b> statement</li> <li>• <b>label-switched-path</b> statement</li> <li>• <b>lsp-lifetime</b> statement</li> <li>• <b>te-metric</b> statement</li> </ul> |
| Logical routers                                                                                                 | <ul style="list-style-type: none"> <li>• <b>logical-routers</b> and subordinate statements</li> </ul>                                                                                                                                                                                                |

Table 315: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)

| Feature                                                                                                                                                                                                                                                                                                                                                                                                                                      | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MPLS: <ul style="list-style-type: none"> <li>Fast Reroute (FRR)</li> <li>Label Distribution Protocol (LDP) (except on EX8200 switches)</li> <li>Layer 3 VPNs (except on EX8200 switches)</li> <li>Multiprotocol BGP (MP-BGP) for VPN-IPv4 family</li> <li>Pseudowire emulation (PWE3)</li> <li>Routing policy statements related to Layer 3 VPNs and MPLS (except on EX8200 switches)</li> <li>Virtual Private LAN Service (VPLS)</li> </ul> | <ul style="list-style-type: none"> <li><b>ldp</b> and all subordinate statements (except on EX8200 switches)</li> </ul>                                                                                                                                                                                                                                                                                           |
| Network Address Translation (NAT)                                                                                                                                                                                                                                                                                                                                                                                                            | <ul style="list-style-type: none"> <li><b>nat</b> and subordinate statements</li> <li>Policy statements related to NAT</li> </ul>                                                                                                                                                                                                                                                                                 |
| OSPF                                                                                                                                                                                                                                                                                                                                                                                                                                         | <ul style="list-style-type: none"> <li><b>demand-circuit</b> statement</li> <li><b>label-switched-path</b> and subordinate statements</li> <li><b>neighbor</b> statement within an OSPF area</li> <li><b>peer-interface</b> and subordinate statements within an OSPF area</li> <li><b>sham-link</b> statement</li> <li><b>te-metric</b> statement</li> </ul>                                                     |
| PPM                                                                                                                                                                                                                                                                                                                                                                                                                                          | <ul style="list-style-type: none"> <li>Not supported on EX2200 and EX3300 switches</li> </ul>                                                                                                                                                                                                                                                                                                                     |
| Routing instances: <ul style="list-style-type: none"> <li>Routing instance forwarding</li> </ul>                                                                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li><b>l2vpn</b> and subordinate statements (except on EX4500, EX4550, and EX8200 switches)</li> <li><b>ldp</b> and subordinate statements (except on EX8200 switches)</li> <li><b>vpls</b> and subordinate statements</li> </ul>                                                                                                                                              |
| Routed VLAN interfaces (RVIs)                                                                                                                                                                                                                                                                                                                                                                                                                | <ul style="list-style-type: none"> <li><b>family mpls</b> statement</li> </ul>                                                                                                                                                                                                                                                                                                                                    |
| SAP and SDP                                                                                                                                                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li><b>sap</b> and all subordinate statements</li> </ul>                                                                                                                                                                                                                                                                                                                       |
| General routing options in the <b>routing-options</b> hierarchy: <ul style="list-style-type: none"> <li>MPLS and label-switched-paths</li> </ul>                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li><b>auto-export</b> and subordinate statements</li> <li><b>dynamic-tunnels</b> and subordinate statements</li> <li><b>lsp-next-hop</b> and subordinate statements</li> <li><b>multicast</b> and subordinate statements</li> <li><b>p2mp-lsp-next-hop</b> and subordinate statements</li> <li><b>route-distinguisher-id</b> statement (except on EX8200 switches)</li> </ul> |

Table 315: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)

| Feature                                                                    | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Traffic sampling and forwarding in the <b>forwarding-options</b> hierarchy | <ul style="list-style-type: none"> <li>• <b>accounting</b> and subordinate statements</li> <li>• <b>family mpls</b> and <b>family multiservice</b> under <b>hash-key</b> hierarchy</li> <li>• Under <b>monitoring group-name</b> family <b>inet output</b> hierarchy: <ul style="list-style-type: none"> <li>• <b>cflowd</b> statement</li> <li>• <b>export-format-cflowd-version-5</b> statement</li> <li>• <b>flow-active-timeout</b> statement</li> <li>• <b>flow-export-destination</b> statement</li> <li>• <b>flow-inactive-timeout</b> statement</li> <li>• <b>interface</b> statement</li> </ul> </li> <li>• <b>port-mirroring</b> statement (On EX Series switches, port mirroring is implemented using the <b>analyzer</b> statement.)</li> <li>• <b>sampling</b> and subordinate statements</li> </ul> |

- Related Documentation**
- [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)
  - [EX Series Switch Software Features Overview](#)

# Configuration

- Configuration Tasks on page 2943
- Configuration Statements on page 2947

## Configuration Tasks

---

- Configuring BGP Sessions (J-Web Procedure) on page 2943

### Configuring BGP Sessions (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You can use the J-Web interface to create BGP peering sessions on a routing device.



**NOTE:** To configure BGP sessions, you must have a license for BGP installed on the EX Series switch.

To configure a BGP peering session:

1. Select **Configure > Routing > BGP**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—Adds a BGP group. Enter information into the configuration page as described in [Table 316 on page 2944](#).
- **Edit**—Modifies an existing BGP group. Enter information into the configuration page as described in [Table 316 on page 2944](#).

- **Delete**—Deletes an existing BGP group.
  - **Disable**—Disables BGP configuration.
3. To modify BGP global settings, click **Edit** in the Global Information section. Enter information as described in [Table 317 on page 2945](#).

**Table 316: BGP Routing Configuration Summary**

| Field                     | Function                                                                                                                  | Your Action                                                                                                                                                                                                                                  |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>General tab</b>        |                                                                                                                           |                                                                                                                                                                                                                                              |
| Group Type                | Specifies whether the group is an internal BGP (IBGP) group or an external BGP (EBGP) group.                              | Select the option: <b>Internal</b> or <b>External</b> .                                                                                                                                                                                      |
| Group Name                | Specifies the name for the group.                                                                                         | Type a new name or select and edit the name.                                                                                                                                                                                                 |
| ASN                       | Sets the unique numeric identifier of the AS in which the routing device is configured.                                   | Type the routing device's 32-bit AS number, in dotted decimal notation.<br><br>If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter <b>3</b> , the value assigned to the AS is <b>0.0.0.3</b> . |
| Preference                | Specifies the degree of preference for an external route. The route with the highest local preference value is preferred. | Type or select and edit the value.                                                                                                                                                                                                           |
| Cluster Id                | Specifies the cluster identifier to be used by the route reflector cluster in an internal BGP group.                      | Type or select and edit the IPv6 or IPv4 address to be used as the identifier.                                                                                                                                                               |
| Description               | Specifies the text description of the global, group, or neighbor configuration.                                           | Type or select and edit the description.                                                                                                                                                                                                     |
| Damping                   | Specifies whether route flap damping is enabled or not.                                                                   | To enable route flap damping, select the check box.<br><br>To disable route flap damping do not select the check box.                                                                                                                        |
| Advertise Inactive Routes | Specifies whether BGP advertises the best route even if the routing table did not select it to be an active route.        | To enable advertising inactive routes, select the check box.<br><br>To disable advertising inactive routes, do not select the check box.                                                                                                     |
| Advertise Peer AS Routes  | Specifies whether to disable the default behavior of suppressing AS routes.                                               | To enable advertising peer AS routes, select the check box.<br><br>To disable advertising peer AS routes, do not select the check box.                                                                                                       |
| <b>Neighbors tab</b>      |                                                                                                                           |                                                                                                                                                                                                                                              |
| Dynamic Neighbors         | Configures a neighbor (peer).                                                                                             | Type the IPv4 address of the peer.                                                                                                                                                                                                           |



Table 316: BGP Routing Configuration Summary (*continued*)

| Field               | Function                                                                                         | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------|--------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Static Neighbors    | Configures the system's peers statically.                                                        | <p>To configure a static neighbor:</p> <ol style="list-style-type: none"> <li>1. Specify the IP address.</li> <li>2. Specify the address of the local end of a BGP session.</li> <li>3. Specify the degree of preference for an external route.</li> <li>4. Enter a description.</li> <li>5. Specify the hold-time value to use when negotiating a connection with the peer.</li> <li>6. Specify how long a route must be present in the routing table before it is exported to BGP. Use this time delay to help bundle routing updates.</li> <li>7. Select <b>Passive</b> if you do not want to send active open messages to the peer.</li> <li>8. Select the option to compare the AS path of an incoming advertised route with the AS number of the BGP peer under the group and replace all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer.</li> <li>9. Specify an import policy and export policy.</li> <li>10. Click <b>OK</b>.</li> </ol> |
| <b>Policies tab</b> |                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Import Policy       | Specifies one or more routing policies to routes being imported into the routing table from BGP. | <p>Click <b>Add</b> to add an import policy. Select the policy and click <b>OK</b>.</p> <p>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.</p> <p>Select the policy and click <b>Remove</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Export Policy       | Specifies one or more policies to routes being exported from the routing table into BGP.         | <p>Click <b>Add</b> to add an export policy. Select the policy and click <b>OK</b>.</p> <p>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.</p> <p>Select the policy and click <b>Remove</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 317: BGP Global Settings

| Field              | Function                                  | Your Action                             |
|--------------------|-------------------------------------------|-----------------------------------------|
| <b>General tab</b> |                                           |                                         |
| Router ASN         | Specifies the routing device's AS number. | Type or select and edit the value.      |
| Router Identifier  | Specify the routing device's IP address.  | Type or select and edit the IP address. |

Table 317: BGP Global Settings (*continued*)

| Field                 | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BGP Status            | Enables or disables BGP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <ul style="list-style-type: none"> <li>To enable BGP, select <b>Enabled</b>.</li> <li>To disable BGP, select <b>Disabled</b>.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Description           | Describes of the global, group, or neighbor configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Type or select and edit the description.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Confederation Number  | Specifies the routing device's confederation AS number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Type or select and edit the value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Confederation Members | Specifies the AS numbers for the confederation members.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <p>To add a member AS number, click <b>Add</b> and enter the number in the <b>Member ASN</b> box. Click <b>OK</b>.</p> <p>To modify a confederation member's AS number, select the member click <b>Edit</b> and, enter the number and click <b>OK</b>.</p> <p>To delete a confederation member, select the member and click <b>Remove</b>.</p>                                                                                                                                                                                                                 |
| Advance Options       | <p>You can configure the following:</p> <ul style="list-style-type: none"> <li><b>Keep routes</b>—Specifies whether routes learned from a BGP peer must be retained in the routing table even if they contain an AS number that was exported from the local AS.</li> <li><b>TCP MSS</b>—Configures the maximum segment size (MSS) for the TCP connection for BGP neighbors.</li> <li><b>MTU Discovery</b>—Select to configure MTU discovery.</li> <li><b>Remove Private ASN</b>—Select to have the local system strip private AS numbers from the AS path when advertising AS paths to remote systems.</li> <li><b>Graceful Restart</b>—Specifies the time period when the restart is expected to be complete. Specify the maximum time that stale routes are kept during restart.</li> <li><b>Multihop</b>—Configures the maximum time-to-live (TTL) value for the TTL in the IP header of BGP packets.</li> <li><b>Authentication Type</b>—Select the authentication algorithm: None, MD5, SHA1, AES.</li> </ul> | <p>Select <b>All</b> or <b>None</b> to configure Keep Routes.</p> <p>Enter a value in the <b>TCP MSS</b> box.</p> <p>Click to enable <b>MTU Discovery</b>.</p> <p>Click to enable <b>Remove Private ASN</b>.</p> <p>Enter the time period for a graceful restart and the maximum time that stale routes must be kept.</p> <p>To configure Multihop, select <b>Nextthop Change</b> to allow unconnected third-party next hops. Enter a TTL value.</p> <p>Select the authentication algorithm. If you select None, specify an authentication key (password).</p> |
| <b>Policies tab</b>   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Import Policy         | Specifies one or more routing policies to routes being imported into the routing table from BGP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <p>Click <b>Add</b> to add an import policy.</p> <p>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.</p> <p>Click <b>Remove</b> to remove an import policy.</p>                                                                                                                                                                                                                                                                                                                                           |

Table 317: BGP Global Settings (*continued*)

| Field                    | Function                                                                                 | Your Action                                                                                                                                                                                                          |
|--------------------------|------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Export Policy            | Specifies one or more policies to routes being exported from the routing table into BGP. | <p>Click <b>Add</b> to add an export policy.</p> <p>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.</p> <p>Click <b>Remove</b> to remove an export policy.</p> |
| <b>Trace Options tab</b> |                                                                                          |                                                                                                                                                                                                                      |
| File Name                | Specifies the name of the file to receive the output of the tracing operation.           | Type or select and edit the name.                                                                                                                                                                                    |
| Number of Files          | Specifies the maximum number of trace files.                                             | Type or select and edit the value.                                                                                                                                                                                   |
| File Size                | Specifies the maximum size for each trace file.                                          | Type or select and edit the value.                                                                                                                                                                                   |
| World Readable           | Specifies whether the trace file can be read by any user or not.                         | <p>Select <b>True</b> to allow any user to read the file.</p> <p>Select <b>False</b> to disallow all users being able to read the file.</p>                                                                          |
| Flags                    | Specifies the tracing operation to perform.                                              | Select a value from the list.                                                                                                                                                                                        |

- Related Documentation**
- [Monitoring BGP Routing Information on page 3035](#)
  - [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)

## Configuration Statements

- [accept-remote-nexthop on page 2949](#)
- [advertise-external on page 2950](#)
- [advertise-inactive on page 2952](#)
- [advertise-peer-as on page 2953](#)
- [aggregate-label on page 2954](#)
- [allow on page 2955](#)
- [as-override on page 2956](#)
- [authentication-algorithm on page 2957](#)
- [authentication-key \(Protocols BGP and BMP\) on page 2958](#)
- [authentication-key-chain \(Protocols BGP and BMP\) on page 2959](#)
- [autonomous-system on page 2960](#)
- [bfd-liveness-detection \(Protocols BGP\) on page 2963](#)
- [bgp on page 2967](#)
- [bgp-orf-cisco-mode on page 2968](#)
- [cluster on page 2970](#)

- [damping \(Protocols BGP\) on page 2972](#)
- [description \(Protocols BGP\) on page 2974](#)
- [disable \(Protocols BGP\) on page 2975](#)
- [explicit-null \(Protocols BGP\) on page 2976](#)
- [export \(Protocols BGP\) on page 2978](#)
- [family \(Protocols BGP\) on page 2979](#)
- [graceful-restart \(Protocols BGP\) on page 2983](#)
- [group \(Protocols BGP\) on page 2984](#)
- [hold-time \(Protocols BGP\) on page 2987](#)
- [idle-after-switch-over on page 2989](#)
- [import \(Protocols BGP\) on page 2990](#)
- [include-mp-next-hop on page 2992](#)
- [keep on page 2993](#)
- [labeled-unicast \(Protocols BGP\) on page 2995](#)
- [local-address \(Protocols BGP\) on page 2997](#)
- [local-as on page 2999](#)
- [local-interface \(IPv6\) on page 3001](#)
- [local-preference on page 3002](#)
- [log-updown \(Protocols BGP\) on page 3003](#)
- [metric-out \(Protocols BGP\) on page 3004](#)
- [mtu-discovery on page 3006](#)
- [multihop on page 3008](#)
- [multipath \(Protocols BGP\) on page 3010](#)
- [neighbor \(Protocols BGP\) on page 3011](#)
- [no-aggregator-id on page 3014](#)
- [no-client-reflect on page 3015](#)
- [no-validate on page 3016](#)
- [out-delay on page 3018](#)
- [outbound-route-filter on page 3020](#)
- [passive \(Protocols BGP\) on page 3021](#)
- [peer-as \(Protocols BGP\) on page 3022](#)
- [preference \(Protocols BGP\) on page 3024](#)
- [prefix-limit on page 3025](#)
- [remove-private on page 3027](#)
- [rib-group \(Protocols BGP\) on page 3029](#)
- [tcp-mss \(Protocols BGP\) on page 3030](#)

- [traceoptions \(Protocols BGP\) on page 3031](#)
- [type \(Protocols BGP\) on page 3034](#)

## accept-remote-nexthop

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>accept-remote-nexthop;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],<br/>         [edit protocols bgp],<br/>         [edit protocols bgp group <i>group-name</i>],<br/>         [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Specify that a single-hop EBGP peer accepts a remote next hop with which it does not share a common subnet. Configure a separate import policy on the EBGP peer to specify the remote next hop. You cannot configure <b>multihop</b> and <b>accept-remote-nexthop</b> statements for the same EBGP peer.</p> <p>For Junos OS Release 13.3 and later releases, specify that a multihop EBGP peer accepts a remote next hop with which it does not share a common subnet. This allows working around current resolver limitations to realize multipath forwarding in recursive next-hop resolution scenarios.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Single-Hop EBGP Peers to Accept Remote Next Hops</a></li> <li>• <a href="#">Understanding Route Advertisement</a></li> <li>• <a href="#">multipath on page 3010</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## advertise-external

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>advertise-external {conditional};</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>     | <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    bgp group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    bgp group <i>group-name</i> neighbor <i>address</i>],</code><br><code>[edit protocols bgp group <i>group-name</i>],</code><br><code>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor</code><br><code>    <i>neighbor-address</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.3.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>         | <p>Specify BGP to advertise the best external route into an IBGP mesh group, a route reflector cluster, or an AS confederation even if the best route is an internal route.</p> <p>In general, deployed BGP implementations do not advertise the external route with the highest local preference value to internal peers unless it is the best route. Although this behavior was required by an earlier version of the BGP version 4 specification, RFC 1771, it was typically not followed in order to minimize the amount of advertised information and to prevent routing loops. However, there are scenarios in which advertising the best external route is beneficial, in particular, situations that can result in IBGP route oscillation.</p> <p>The <b>advertise-external</b> statement is supported at both the group and neighbor level. If you configure the statement at the neighbor level, you must configure it for all neighbors in a group. Otherwise, the group is automatically split into different groups.</p> <p>In a confederation, when advertising a route to a confederation border router, any route from a different confederation sub-AS is considered external. When configuring the <b>advertise-external</b> statement for an AS confederation, it is recommended that EBGp peers belonging to different autonomous systems are configured in a separate EBGp peer group. This ensures consistency while BGP sends the best external route to peers in the configured peer group.</p> <p>To configure the <b>advertise-external</b> statement on a route reflector, you must disable intracluster reflection with the <b>no-client-reflect</b> statement.</p> <p>When a routing device is configured as a route reflector for a cluster, a route advertised by the route reflector is considered internal if it is received from an internal peer with the same cluster identifier or if both peers have no cluster identifier configured. A route received from an internal peer that belongs to another cluster, that is, with a different cluster identifier, is considered external.</p> <p>The <b>conditional</b> option causes BGP to advertise the external route only if the route selection process reaches the point where the multiple exit discriminator (MED) metric</p> |

is evaluated. As a result, an external route with an AS path longer than that of the active path is not advertised.

Junos OS also provides support for configuring a BGP export policy that matches on the state of an advertised route. You can match on either active or inactive routes.

**Default** BGP does not advertise the external route with the highest local preference value to internal peers unless it is the best route.

**Options** **conditional**—(Optional) Advertise the best external path only if the route selection process reaches the point at which the multiple exit discriminator (MED) metric is evaluated. The **conditional** option restricts advertisement to when the best external path and the active path are equal until the MED step of the route selection process. This implies that external routes with a longer AS path length than the active path, for instance, are not advertised. The criteria used for selecting the best external path is the same whether or not the **conditional** option is configured.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring BGP to Advertise the Best External Route to Internal Peers*
- [advertise-inactive on page 2952](#)

## advertise-inactive

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | advertise-inactive;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure the routing table to export to BGP the best route learned by BGP even if Junos OS did not select this route to be an active route.</p> <p>One way to achieve multivendor compatibility is to include the <b>advertise-inactive</b> statement in the external BGP (EBGP) configuration. By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. The <b>advertise-inactive</b> statement causes Junos OS to advertise the best BGP route that is inactive because of IGP preference. When you use the <b>advertise-inactive</b> statement, the Junos OS device uses, for example, the OSPF route for forwarding, and the other vendor's device uses the EBGP route for forwarding. However, from the perspective of an EBGP peer in a neighboring AS, both vendors' devices appear to behave the same way.</p>                                                                                                                                                                                                                                                                         |
| <b>Default</b>                  | By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Setting BGP to Advertise Inactive Routes</i></li> <li>• <i>Example: Configuring the Preference Value for BGP Routes</i></li> <li>• <i>Example: Configuring BGP Route Preference (Administrative Distance)</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



- [advertise-external on page 2950](#)

## advertise-peer-as



|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | advertise-peer-as;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],<br/>         [edit protocols bgp],<br/>         [edit protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Disable the default behavior of suppressing AS routes.</p> <p>If you include the <b>advertise-peer-as</b> statement in the configuration, BGP advertises routes learned from one external BGP (EBGP) peer back to another EBGP peer in the same autonomous system (AS) but not back to the originating peer.</p> <p>Another way to disable the route suppression default behavior is with the <b>as-override</b> statement. If you include both the <b>as-override</b> and <b>no-advertise-peer-as</b> statements in the configuration, the <b>no-advertise-peer-as</b> statement is ignored.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>                  | By default, Junos OS does not advertise the routes learned from one EBGP peer back to the same external BGP (EBGP) peer. In addition, the software does not advertise those routes back to any EBGP peers that are in the same AS as the originating peer, regardless of the routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Disabling Suppression of Route Advertisements</i></li> <li>• <i>Example: Configuring a Layer 3 VPN with Route Reflection and AS Override</i></li> <li>• <i>no-advertise-peer-as</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## aggregate-label


---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>aggregate-label {<br/>    community <i>community-name</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols bgp family inet labeled-unicast],<br>[edit logical-systems <i>logical-system-name</i> protocols bgp family inet6 labeled-unicast],<br>[edit logical-systems <i>logical-system-name</i> protocols bgp family inet-vpn unicast],<br>[edit logical-systems <i>logical-system-name</i> protocols bgp family inet-vpn6 unicast],<br>[edit protocols bgp family inet labeled-unicast],<br>[edit protocols bgp family inet6 labeled-unicast],<br>[edit protocols bgp family inet-vpn unicast],<br>[edit protocols bgp family inet6-vpn unicast] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify matching criteria (in the form of a community) such that all routes which match are assigned the same VPN label, selected from one of the several routes in the set defined by this criteria. This reduces the number of VPN labels that the router must consider, and aggregates the received labels.                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>community <i>community-name</i></b> —Specify the name of the community to which to apply the aggregate label.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Aggregate Labels for VPNs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## allow

|                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                 | <code>allow (all   [ <i>network/mask-length</i> ] );</code>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                        | [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>bgp <b>group</b> <i>group-name</i> ],<br>[edit protocols bgp <b>group</b> <i>group-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> ] |
| <b>Release Information</b>                                                                                                                                                                    | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                         |
| <b>Description</b>                                                                                                                                                                            | Implicitly configure BGP peers, allowing peer connections from any of the specified networks or hosts. To configure multiple BGP peers, configure one or more networks and hosts within a single <b>allow</b> statement or include multiple <b>allow</b> statements.                                                                                                                                                   |
| <div>  <b>NOTE:</b> You cannot define a BGP group with dynamic peers with BGP authentication enabled. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                                                                                                                                                                | <p><b>all</b>—Allow all addresses, which is equivalent to <b>0.0.0.0/0</b> (or <b>::/0</b>).</p> <p><b><i>network/mask-length</i></b>—IPv6 or IPv4 network number of a single address or a range of allowable addresses for BGP peers, followed by the number of significant bits in the subnet mask.</p>                                                                                                              |
| <div>  <b>NOTE:</b> You cannot define a BGP group with dynamic peers with authentication enabled. </div>   |                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                                                                                                                                                               | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>                                                                                                                                                                  | <ul style="list-style-type: none"> <li><a href="#">neighbor on page 3011</a></li> </ul>                                                                                                                                                                                                                                                                                                                                |

## as-override

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | as-override;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Compare the AS path of an incoming advertised route with the AS number of the BGP peer under the group and replace all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <div>  <p><b>NOTE:</b> The <b>as-override</b> statement is specific to a particular BGP group. This statement does not affect peers from the same remote AS configured in different groups.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <p>Enabling the AS override feature allows routes originating from an AS to be accepted by a router residing in the same AS. Without AS override enabled, the routing device refuses the route advertisement once the AS path shows that the route originated from its own AS. This is done by default to prevent route loops. The <b>as-override</b> statement overrides this default behavior.</p> <p>Note that enabling the AS override feature may result in routing loops. Use this feature only for specific applications that require this type of behavior, and in situations with strict network control. One application is the IGP protocol between the provider edge routing device and the customer edge routing device in a virtual private network.</p> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <ul style="list-style-type: none"> <li>• <i>Example: Configuring a Layer 3 VPN with Route Reflection and AS Override</i></li> <li>• <i>Junos OS VPNs Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## authentication-algorithm

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>authentication-algorithm <i>algorithm</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols ldp session <i>session-address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <i>group</i> <i>group-name</i> <b>neighbor</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   ldp session <i>session-address</i>], [edit logical-systems <i>logical-system-name</i> routing-options <b>bmp</b>], [edit logical-systems <i>logical-system-name</i> routing-options bmp station <i>station-name</i>], [edit protocols bgp], [edit protocols bgp <b>group</b> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> <b>neighbor</b> <i>address</i>], [edit protocols ldp session <i>session-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>   <b>neighbor</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols ldp session <i>session-address</i>], [edit routing-options <b>bmp</b>], [edit routing-options bmp station <i>station-name</i>]</pre> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced for BGP in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p>Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced for BMP in Junos OS Release 13.3.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>         | Configure an authentication algorithm type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>             | <p><b><i>algorithm</i></b>—Specify one of the following types of authentication algorithms:</p> <ul style="list-style-type: none"> <li><b>aes-128-cmac-96</b>—Cipher-based message authentication code (AES128, 96 bits).</li> <li><b>hmac-sha-1-96</b>—Hash-based message authentication code (SHA1, 96 bits).</li> <li><b>md5</b>—Message digest 5.</li> </ul> <p><b>Default:</b> <code>hmac-sha-1-96</code></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



**NOTE:** The default is not displayed in the output of the `show bgp bmp` command unless a key or key-chain is also configured.

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Route Authentication for BGP*
- *Configuring BGP Monitoring Protocol Version 3*

---

## authentication-key (Protocols BGP and BMP)

---

**Syntax** authentication-key *key*;

**Hierarchy Level**

```
[edit logical-systems logical-system-name protocols bgp],  
[edit logical-systems logical-system-name protocols bgp group group-name],  
[edit logical-systems logical-system-name protocols bgp group group-name  
  neighbor address],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols  
  bgp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols  
  bgp group group-name],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols  
  bgp group group-name neighbor address],  
[edit logical-systems logical-system-name routing-options bmp],  
[edit logical-systems logical-system-name routing-options bmp station station-name],  
[edit protocols bgp],  
[edit protocols bgp group group-name],  
[edit protocols bgp group group-name neighbor address],  
[edit routing-instances routing-instance-name protocols bgp],  
[edit routing-instances routing-instance-name protocols bgp group group-name],  
[edit routing-instances routing-instance-name protocols bgp group group-name  
  neighbor address],  
[edit routing-options bmp],  
[edit routing-options bmp station station-name]
```

**Release Information**

Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.  
Statement introduced for BMP version 3 in Junos OS Release 13.3.

**Description** Configure an MD5 authentication key (password). Neighboring routing devices use the same password to verify the authenticity of BGP packets sent from this system.

**Options** *key*—Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Route Authentication for BGP*
- *Configuring BGP Monitoring Protocol Version 3*

## authentication-key-chain (Protocols BGP and BMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | authentication-key-chain <i>key-chain</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>bmp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options bmp station <i>station-name</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-options <b>bmp</b>],</p> <p>[edit routing-options bmp station <i>station-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced for BMP in Junos OS Release 13.3.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update feature for BGP, you cannot commit the <b>0.0.0.0/allow</b> statement with authentication keys or key chains. The CLI issues a warning and fails to commit the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>key-chain</b> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Route Authentication for BGP</i></li> <li>• <i>Example: Configuring BFD Authentication for Static Routes</i></li> <li>• <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i></li> <li>• <i>Configuring BGP Monitoring Protocol Version 3</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## autonomous-system

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>autonomous-system <i>autonomous-system</i> &lt;asdot-notation&gt; &lt;loops <i>number</i>&gt; {<br/>    independent-domain &lt;no-attrset&gt;;<br/>}</code>                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options]                                                                                                                                                                                                             |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br><b>asdot-notation</b> option introduced in Junos OS Release 9.3.<br><b>asdot-notation</b> option introduced in Junos OS Release 9.3 for EX Series switches.<br><b>no-attrset</b> option introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers. |
| <b>Description</b>         | Specify the routing device's AS number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

An autonomous system (AS) is a set of routing devices that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it. ASs are identified by a number that is assigned by the Network Information Center (NIC) in the United States (<http://www.isi.edu>).

If you are using BGP on the routing device, you must configure an AS number.

The AS path attribute is modified when a route is advertised to an EBGP peer. Each time a route is advertised to an EBGP peer, the local routing device prepends its AS number to the existing path attribute, and a value of 1 is added to the AS number.

In Junos OS Release 9.1 and later, the numeric range is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. RFC 4893 introduces two new optional transitive BGP attributes, AS4\_PATH and AS4\_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS\_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.

**Options**    ***autonomous-system***—AS number. Use a number assigned to you by the NIC.



**Range:** 1 through 4,294,967,295 ( $2^{32} - 1$ ) in plain-number format for 4-byte AS numbers

In this example, the 4-byte AS number 65,546 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 65546;
}
```

**Range:** 0.0 through 65535.65535 in AS-dot notation format for 4-byte numbers

In this example, 1.10 is the AS-dot notation format for 65,546:

```
[edit]
routing-options {
  autonomous-system 1.10;
}
```

**Range:** 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

In this example, the 2-byte AS number 60,000 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 60000;
}
```

**asdot-notation**—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

**Default:** Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

**loops number**—(Optional) Specify the number of times detection of the AS number in the AS\_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

**Range:** 1 through 10

**Default:** 1



**NOTE:** When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the loops statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.

Use the **independent-domain** option if the loops statement must be enabled only on a subset of routing instances.

The remaining statement is explained separately.

|                              |                                                                                                                                                                 |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege</b>    | routing—To view this statement in the configuration.                                                                                                            |
| <b>Level</b>                 | routing-control—To add this statement to the configuration.                                                                                                     |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Examples: Configuring External BGP Peering</i></li><li>• <i>Examples: Configuring Internal BGP Peering</i></li></ul> |

## bfd-liveness-detection (Protocols BGP)

**Syntax** `bfd-liveness-detection {`  
     `authentication {`  
         `algorithm` *algorithm-name*;  
         `key-chain` *key-chain-name*;  
         `loose-check`;  
     `}`  
     `detection-time {`  
         `threshold` *milliseconds*;  
     `}`  
     `hold-down-interval` *milliseconds*;  
     `minimum-interval` *milliseconds*;  
     `minimum-receive-interval` *milliseconds*;  
     `multiplier` *number*;  
     `no-adaptation`;  
     `session-mode` (automatic | multihop | single-hop);  
     `transmit-interval {`  
         `minimum-interval` *milliseconds*;  
         `threshold` *milliseconds*;  
     `}`  
     `version` (1 | automatic);  
`}`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols bgp],  
 [edit logical-systems *logical-system-name* protocols bgp group *group-name*],  
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],  
 [edit protocols bgp],  
 [edit protocols bgp group *group-name*],  
 [edit protocols bgp group *group-name* neighbor *address*],  
 [edit routing-instances *routing-instance-name* protocols bgp],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*]

**Release Information** Statement introduced in Junos OS Release 8.1.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
**detection-time threshold** and **transmit-interval threshold** options introduced in Junos OS Release 8.2  
 Support for logical routers introduced in Junos OS Release 8.3.  
 Support for IBGP and multihop EBGP sessions introduced in Junos OS Release 8.3.  
**holddown-interval** statement introduced in Junos OS Release 8.5. You can configure this statement only for EBGP peers at the [edit protocols bgp group *group-name* neighbor *address*] hierarchy level.  
**no-adaptation** statement introduced in Junos OS Release 9.0.  
 Support for BFD authentication introduced in Junos OS Release 9.6.

Support for BFD on IPv6 interfaces with BGP introduced in Junos OS Release 11.2.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure bidirectional failure detection (BFD) timers and authentication for BGP.

For IBGP and multihop EBGP support, configure the **bfd-liveness-detection** statement at the global **[edit bgp protocols]** hierarchy level. You can also configure IBGP and multihop support for a routing instance or a logical system.

**Options** **authentication algorithm** *algorithm-name* (Optional)—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**.

**authentication key-chain** *key-chain-name* (Optional)—Associate a security key with the specified BFD session using the name of the security keychain. The keychain name must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

**detection-time threshold** *milliseconds* (Optional)—Configure a threshold. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**holddown-interval** *milliseconds* (Optional)—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. When you configure the hold-down interval for the BFD protocol for EBGp, the BFD session is unaware of the BGP session during this time. In this case, if the BGP session goes down during the configured hold-down interval, BFD already assumes it is down and does not send a state change notification. The **holddown-interval** statement is supported only for EBGp peers at the **[edit protocols bgp group group-name neighbor address]** hierarchy level. If the BFD session goes down and then comes back up during the configured hold-down interval, the timer is restarted. You must configure the hold-down interval on both EBGp peers. If you configure the hold-down interval for a multihop EBGp session, you must also configure a local IP address by including the **local-address** statement at the **[edit protocols bgp group group-name]** hierarchy level.

**Range:** 0 through 255,000

**Default:** 0

**minimum-interval** *milliseconds* (Required)—Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately (using the **minimum-receive-interval** and **transmit-interval** statements).

**Range:** 1 through 255,000

**minimum-receive-interval** *milliseconds* (Optional)—Configure only the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session.

**Range:** 1 through 255,000

**multiplier *number*** (Optional)—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation** (Optional)—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable to not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds*** (Optional)—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval *milliseconds*** (Optional)—Configure only the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.

**Range:** 1 through 255,000

**version** (Optional)—Configure the BFD version to detect.

**Range:** 1 or **automatic** (autodetect the BFD version)

**Default:** **automatic**

The remaining statements are explained separately.

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|


|                              |                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Example: Configuring BFD for Static Routes</i></li><li>• <i>Example: Configuring BFD Authentication for Static Routes</i></li><li>• <i>Example: Configuring BFD on Internal BGP Peer Sessions</i></li><li>• <i>Example: Configuring BFD Authentication for BGP</i></li><li>• <i>Understanding BFD for BGP</i></li></ul> |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## bgp

---

|                                 |                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bgp { ... }</code>                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols bgp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                              |
| <b>Description</b>              | Enable BGP on the routing device or for a routing instance.                                                                                                                                                                                                                 |
| <b>Default</b>                  | BGP is disabled.                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>BGP Feature Guide for Routing Devices</i></li> </ul>                                                                                                                                                                            |

## bgp-orf-cisco-mode

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bgp-orf-cisco-mode;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit routing-options <b>outbound-route-filter</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|                                 | <p> <b>NOTE:</b> To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



**Related Documentation**

- *Example: Configuring BGP Prefix-Based Outbound Route Filtering*

## cluster

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>cluster <i>cluster-identifier</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>         | Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



### CAUTION:

If you configure both route reflection and VPNs on the same routing device, the following modifications to the route reflection configuration cause current BGP sessions to be reset:

- Adding a cluster ID—If a BGP session shares the same AS number with the group where you add the cluster ID, all BGP sessions are reset regardless of whether the BGP sessions are contained in the same group.
- Creating a new route reflector—If you have an IBGP group with an AS number and create a new route reflector group with the same AS number, all BGP sessions in the IBGP group and the new route reflector group are reset.



**NOTE:** If you change the address family specified in the [edit protocols bgp family] hierarchy level, all current BGP sessions on the routing device are dropped and then reestablished.

|                                 |                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <i>cluster-identifier</i> —4-byte identifier (such as an IPv4 address).                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring BGP Route Reflectors</i></li><li>• <i>Understanding External BGP Peering Sessions</i></li><li>• <a href="#">no-client-reflect on page 3015</a></li></ul> |

## damping (Protocols BGP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | damping;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | <pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>   <b>neighbor</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>   family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>   family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <b>group</b> <i>group-name</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <b>group</b> <i>group-name</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> family <i>family</i>], [edit protocols bgp], [edit protocols bgp], [edit protocols bgp <b>group</b> <i>group-name</i>], [edit protocols bgp <b>group</b> <i>group-name</i> family <i>family</i>], [edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>], [edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>   <b>neighbor</b> <i>address</i>] [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>   <b>neighbor</b> <i>address</i> family <i>family</i>] </pre> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for flap damping at the address family level introduced in Junos OS Release 12.2.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>         | <p>Enable route flap damping. BGP route flapping describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability information. Flap damping reduces the number of update messages sent between BGP</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.

You typically apply flap damping to external BGP (EBGP) routes (that is, to routes in different ASs). You can also apply it within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.) The exception to this rule is when flap damping is applied at the address family level. When you apply flap damping at the address family level, it works for both IBGP and EBGP.

**Default** Flap damping is disabled on the routing device.

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- *Examples: Configuring BGP Flap Damping*
- *Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family*

## description (Protocols BGP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description text-description;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems logical-system-name protocols bgp],</code><br><code>[edit logical-systems logical-system-name protocols bgp group group-name],</code><br><code>[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],</code><br><code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],</code><br><code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name],</code><br><code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name neighbor address],</code><br><code>[edit protocols bgp],</code><br><code>[edit protocols bgp group group-name],</code><br><code>[edit protocols bgp group group-name neighbor address],</code><br><code>[edit routing-instances routing-instance-name protocols bgp],</code><br><code>[edit routing-instances routing-instance-name protocols bgp group group-name],</code><br><code>[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Provide a description of the global, group, or neighbor configuration. If the text includes one or more spaces, enclose it in quotation marks (" "). The text is displayed in the output of the <b>show</b> command and has no effect on the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <i>text-description</i> —Text description of the configuration. It is limited to 255 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><i>BGP Feature Guide for Routing Devices</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

---

## disable (Protocols BGP)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols bgp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br>[edit protocols bgp],<br>[edit routing-instances <i>routing-instance-name</i> protocols bgp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                      |
| <b>Description</b>              | Disable BGP on the system.                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |

## explicit-null (Protocols BGP)

**Syntax** explicit-null;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols mpls],  
 [edit logical-systems *logical-system-name* protocols bgp *family* inet labeled-unicast],  
 [edit logical-systems *logical-system-name* protocols bgp *family* inet6 labeled-unicast],  
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* *family* inet labeled-unicast],  
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* *family* inet6 labeled-unicast],  
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* *family* inet labeled-unicast],  
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* *family* inet6 labeled-unicast],  
 [edit logical-systems *logical-system-name* protocols ldp],  
 [edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp *family* inet labeled-unicast],  
 [edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp *family* inet6 labeled-unicast],  
 [edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp group *group-name* *family* inet labeled-unicast],  
 [edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp group *group-name* *family* inet6 labeled-unicast],  
 [edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp group *group-name* neighbor *address* *family* inet labeled-unicast],  
 [edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp group *group-name* neighbor *address* *family* inet6 labeled-unicast],  
 [edit logical-systems *logical-system-name* routing-instances *instance-name* protocols ldp],  
 [edit protocols mpls],  
 [edit protocols bgp *family* inet labeled-unicast],  
 [edit protocols bgp *family* inet6 labeled-unicast],  
 [edit protocols bgp group *group-name* *family* inet labeled-unicast],  
 [edit protocols bgp group *group-name* *family* inet6 labeled-unicast],  
 [edit protocols bgp group *group-name* neighbor *address* *family* inet labeled-unicast],  
 [edit protocols bgp group *group-name* neighbor *address* *family* inet6 labeled-unicast],  
 [edit protocols ldp],  
 [edit routing-instances *instance-name* protocols bgp *family* inet labeled-unicast],  
 [edit routing-instances *instance-name* protocols bgp *family* inet6 labeled-unicast],  
 [edit routing-instances *instance-name* protocols bgp group *group-name* *family* inet labeled-unicast],  
 [edit routing-instances *instance-name* protocols bgp group *group-name* *family* inet6 labeled-unicast],  
 [edit routing-instances *instance-name* protocols bgp group *group-name* neighbor *address* *family* inet labeled-unicast],  
 [edit routing-instances *instance-name* protocols bgp group *group-name* neighbor *address* *family* inet6 labeled-unicast],  
 [edit routing-instances *instance-name* protocols ldp]

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Advertise label 0 to the egress routing device of an LSP.



|                              |                                                                                                                       |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>               | If you do not include the <b>explicit-null</b> statement in the configuration, label 3 (implicit null) is advertised. |
| <b>Required Privilege</b>    | routing—To view this statement in the configuration.                                                                  |
| <b>Level</b>                 | routing-control—To add this statement to the configuration.                                                           |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Advertising Explicit Null Labels to BGP Peers</i></li></ul>                |

## export (Protocols BGP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i></code><br><code>    <i>neighbor</i> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    bgp],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    bgp <i>group</i> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>],</code><br><code>[edit protocols bgp],</code><br><code>[edit protocols bgp <i>group</i> <i>group-name</i>],</code><br><code>[edit protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i></code><br><code>    <i>neighbor</i> <i>address</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Apply one or more policies to routes being exported from the routing table into BGP.</p> <p>If you specify more than one policy, they are evaluated in the order specified, from left to right, and the first matching filter is applied to the route. If no routes match the filters, the routing table exports into BGP only the routes that it learned from BGP. If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during the evaluation of the remaining policies. For example, if the action specified in the first policy of a chain sets a route's metric to 500, this route matches the criterion of <b>metric 500</b> defined in the next policy.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding Route Advertisement</i></li><li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li><li>• <a href="#">import on page 2990</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## family (Protocols BGP)

```
Syntax  family {
    (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) {
        (any | flow | labeled-unicast | multicast | unicast) {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage-threshold> idle-timeout (forever | minutes);
            }
            add-path {
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
                receive;
            }
            algp [disable];
            loops number;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            protection;
            rib-group group-name;
            topology name {
                community {
                    target identifier;
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib (inet.3 | inet6.3);
            rib-group group-name;
            traffic-statistics {
                file filename <world-readable | no-world-readable>;
                interval seconds;
            }
        }
    }
}
```

```
    }
  }
  route-target {
    accepted-prefix-limit {
      maximum number;
      proxy-generate <route-target-policy route-target-policy-name>;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
  }
  (evpn | inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
    signaling {
      accepted-prefix-limit {
        maximum number;
        teardown <percentage-threshold> idle-timeout (forever | minutes);
      }
      add-path {
        send {
          path-count number;
          prefix-policy [ policy-names ];
        }
        receive;
      }
      aigp [disable];
      damping;
      loops number;
      prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      rib-group group-name;
    }
  }
}
```

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>     | <pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>   <b>neighbor</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>], [edit protocols bgp], [edit protocols bgp <b>group</b> <i>group-name</i>], [edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>   <b>neighbor</b> <i>address</i>] </pre> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>inet-mvpn</b> and <b>inet6-mvpn</b> statements introduced in Junos OS Release 8.4.</p> <p><b>inet-mdt</b> statement introduced in Junos OS Release 9.4.</p> <p>Support for the <b>loops</b> statement introduced in Junos OS Release 9.6.</p> <p><b>evpn</b> statement introduced in Junos OS Release 13.2.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | <p>Enable multiprotocol BGP (MP-BGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, to specify MP-BGP to carry NLRI for the IPv6 address family, or to carry NLRI for VPNs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |


- Options**
- any**—Configure the family type to be both unicast and multicast.
  - evpn**—Configure NLRI parameters for Ethernet VPNs (EVPNs).
  - inet**—Configure NLRI parameters for IPv4.
  - inet6**—Configure NLRI parameters for IPv6.
  - inet-mdt**—Configure NLRI parameters for the multicast distribution tree (MDT) subaddress family identifier (SAFI) for IPv4 traffic in Layer 3 VPNs.
  - inet-mvpn**—Configure NLRI parameters for IPv4 for multicast VPNs.
  - inet6-mvpn**—Configure NLRI parameters for IPv6 for multicast VPNs.
  - inet-vpn**—Configure NLRI parameters for IPv4 for Layer 3 VPNs.
  - inet6-vpn**—Configure NLRI parameters for IPv6 for Layer 3 VPNs.
  - iso-vpn**—Configure NLRI parameters for IS-IS for Layer 3 VPNs.
  - l2vpn**—Configure NLRI parameters for IPv4 for MPLS-based Layer 2 VPNs and VPLS.
  - labeled-unicast**—Configure the family type to be labeled-unicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by labeled-unicast for resolving the labeled-unicast routes. This statement is supported only with **inet** and **inet6**.
  - multicast**—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.
  - unicast**—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes. The default family type is **unicast**.

The remaining statements are explained separately.

- Required Privilege Level**
- routing—To view this statement in the configuration.
  - routing-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring IBGP Sessions Between PE Routers in VPNs*
  - *Understanding Multiprotocol BGP*
  - [autonomous-system on page 2960](#)
  - [local-as on page 2999](#)

## graceful-restart (Protocols BGP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>graceful-restart {   disable;   restart-time <i>seconds</i>;   stale-routes-time <i>seconds</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/> [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/> [edit protocols bgp],<br/> [edit protocols bgp <b>group</b> <i>group-name</i>],<br/> [edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/> Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/> Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Configure graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the <b>restart-time</b> statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the <b>stale-routes-time</b> statement at the [edit protocols bgp graceful-restart] hierarchy level.</p> <hr/> <div>  <p><b>NOTE:</b> If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p> </div> <hr/> <p>Configure graceful restart globally at the [edit routing-options] or [edit routing-instances <i>instance-name</i> routing-options] hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/> routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Graceful Restart Options for BGP</i></li> <li>• <i>Configuring Graceful Restart for QFabric Systems</i></li> <li>• <i>Junos OS High Availability Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## group (Protocols BGP)

---

```
Syntax  group group-name {
        advertise-inactive;
        allow [ network/mask-length ];
        authentication-key key;
        cluster cluster-identifier;
        damping;
        description text-description;
        export [ policy-names ];
        family {
            (inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {
                (any | multicast | unicast | signaling) {
                    accepted-prefix-limit {
                        maximum number;
                        teardown <percentage> <idle-timeout (forever | minutes)>;
                    }
                    add-path {
                        send {
                            path-count number;
                            prefix-policy [ policy-names ];
                        }
                        receive;
                    }
                    aigp [disable];
                    damping;
                    prefix-limit {
                        maximum number;
                        teardown <percentage> <idle-timeout (forever | minutes)>;
                    }
                    rib-group group-name;
                    topology name {
                        community {
                            target identifier;
                        }
                    }
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib inet.3;
```



```

        rib-group group-name;
    }
}
route-target {
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-preference local-preference;
log-updown;
metric-out metric;
multihop <ttl-value>;
multipath {
    multiple-as;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
remove-private;
tcp-aggressive-transmission;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
type type;
neighbor address {
    ... peer-specific-options ...
}
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols bgp],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 bgp],  
 [edit protocols bgp],  
 [edit routing-instances *routing-instance-name* protocols bgp]

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Define a BGP peer group. BGP peer groups share a common type, peer autonomous system (AS) number, and cluster ID, if present. To configure multiple BGP groups, include multiple <b>group</b> statements.</p> <p>By default, the group's options are identical to the global BGP options. To override the global options, include group-specific options within the <b>group</b> statement.</p> <p>The <b>group</b> statement is one of the statements you must include in the configuration to run BGP on the routing device.</p> <p>Each group must contain at least one peer.</p> |
| <b>Options</b>                  | <p><b>group-name</b>—Name of the BGP group.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>BGP Feature Guide for Routing Devices</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## hold-time (Protocols BGP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>hold-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>                                                                                                                           |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>         | <p>Specify the hold-time value to use when negotiating a connection with the peer. The hold-time value is advertised in open packets and indicates to the peer the length of time that it should consider the sender valid. If the peer does not receive a keepalive, update, or notification message within the specified hold time, the BGP connection to the peer is closed and routing devices through that peer become unavailable.</p> <p>The hold time is three times the interval at which keepalive messages are sent.</p> <p>BGP on the local routing device uses the smaller of either the local hold-time value or the peer's hold-time value received in the open message as the hold time for the BGP connection between the two peers.</p> <p>Starting in Junos OS Release 12.3, the BGP hold-time value can be zero (0). This implies that the speaker does not expect keepalive messages from its peer to maintain the BGP session. When negotiating between two peers, if one side requests a nonzero hold time and the other requests a zero hold time, the negotiation settles on the nonzero value and keepalive intervals are determined accordingly. Both sides must be set to zero for keepalive messages to stop being sent.</p> |
| <b>Options</b>             | <p><b>seconds</b>—Hold time.</p> <p><b>Range:</b> 10 through 65,535 seconds (or 0 for infinite hold time)</p> <p><b>Default:</b> 90 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |



**TIP:** When you set a hold-time value of 1 through 19 seconds, we recommend that you also configure the BGP `precision-timers` statement. The `precision-timers` statement ensures that if scheduler slip messages occur, the routing device continues to send keepalive messages. When the `precision-timers` statement is included, keepalive message generation is performed in a dedicated kernel thread, which helps to prevent BGP session flaps.

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>BGP Messages Overview</i></li><li>• <i>precision-timers</i></li></ul>    |

## idle-after-switch-over

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>idle-after-switch-over (forever   seconds);</code>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure the routing device so that it does not automatically reestablish BGP peer sessions after a nonstop active routing (NSR) switchover. This feature is particularly useful if you are using dynamic routing policies because the dynamic database is not synchronized with the backup Routing Engine when NSR is enabled.                                                                                                                              |
| <b>Options</b>                  | <p><b>forever</b>—Do not reestablish a BGP peer session after an non-stop routing switchover until the <b>clear bgp neighbor</b> command is issued.</p> <p><b>seconds</b>—Do not reestablish a BGP peer session after an non-stop routing switchover until after the specified period.</p> <p><b>Range:</b> 1 through 4,294,967,295 (<math>2^{32} - 1</math>)</p>                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Preventing Automatic Reestablishment of BGP Peer Sessions After NSR Switchovers</i></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> <li>• <i>Junos OS High Availability Library for Routing Devices</i></li> </ul>                                                                                                                              |

## import (Protocols BGP)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>     | <code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code><br><code>    <i>neighbor address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    bgp],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    bgp <i>group group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    bgp <i>group group-name</i> <i>neighbor address</i>],</code><br><code>[edit protocols bgp],</code><br><code>[edit protocols bgp <i>group group-name</i>],</code><br><code>[edit protocols bgp <i>group group-name</i> <i>neighbor address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code><br><code>    <i>neighbor address</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>         | <p>Apply one or more routing policies to routes being imported into the Junos OS routing table from BGP.</p> <p>If you specify more than one policy, they are evaluated in the order specified, from left to right, and the first matching filter is applied to the route. If no match is found, BGP places into the routing table only those routes that were learned from BGP routing devices. The policy framework software evaluates the routing policies in a chain sequentially. If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during the evaluation of the remaining policies. For example, if the action specified in the first policy of a chain sets a route's metric to 500, this route matches the criterion of <b>metric 500</b> defined in the next policy.</p> <p>It is also important to understand that in Junos OS, although an import policy (inbound route filter) might reject a route, not use it for traffic forwarding, and not include it in an advertisement to other peers, the router retains these routes as hidden routes. These hidden routes are not available for policy or routing purposes. However, they do occupy memory space on the router. A service provider filtering routes to control the amount of information being kept in memory and processed by a router might want the router to entirely drop the routes being rejected by the import policy.</p> <p>Hidden routes can be viewed by using the <b>show route receive-protocol bgp <i>neighbor-address</i> hidden</b> command. The hidden routes can then be retained or dropped from the routing table by configuring the <b>keep all   none</b> statement at the <b>[edit protocols bgp]</b> or <b>[edit protocols bgp <i>group group-name</i>]</b> hierarchy level.</p> |

The rules of BGP route retention are as follows:

- By default, all routes learned from BGP are retained, except those where the AS path is looped. (The AS path includes the local AS.)
- By configuring the **keep all** statement, all routes learned from BGP are retained, even those with the local AS in the AS path.
- By configuring the **keep none** statement, all routes received are discarded. When this statement is configured and the inbound policy changes, Junos OS re-advertises all the routes advertised by the peer.

**Options**    *policy-names*—Name of one or more policies.

**Required Privilege**    routing—To view this statement in the configuration.  
**Level**    routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring BGP Interactions with IGPs*
- *Understanding Route Advertisement*
- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*
- [export on page 2978](#)

## include-mp-next-hop

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | include-mp-next-hop;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>   <i>neighbor</i> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>   <i>neighbor</i> <i>address</i>]</pre> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Enable multiprotocol updates to contain next-hop reachability information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Examples: Configuring Multiprotocol BGP</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |



## keep

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | keep (all   none);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>         | <p>Control whether or not Junos OS keeps in memory and hides certain routes.</p> <p>If the <b>keep none</b> statement is used, Junos OS does not retain in memory and hide routes that are rejected because of a BGP import policy. Nor does BGP keep in memory and hide routes that are declared unfeasible due to BGP sanity checks. The <b>keep none</b> statement causes Junos OS to discard from memory the routes that are rejected due to BGP-specific logic or BGP evaluation. When a route is rejected because of some non-BGP-specific reason, the <b>keep none</b> statement has no effect on this route. This rejected route is retained in memory and hidden even though <b>keep none</b> is configured. An example of this type of hidden route is a route for which the protocol nexthop is unresolved.</p> <p>The routing table can retain the route information learned from BGP in one of the following ways:</p> <ul style="list-style-type: none"> <li>• Default (omit the <b>keep</b> statement)—Keep all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.</li> <li>• <b>keep all</b>—Keep all route information that was learned from BGP.</li> <li>• <b>keep none</b>—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking, such as AS path or next hop. When you configure <b>keep none</b> for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.</li> </ul> |

In an AS path healing situation, routes with looped paths theoretically could become usable during a soft reconfiguration when the AS path loop limit is changed. However, there is a significant memory usage difference between the default and **keep all**.

Consider the following scenarios:

- A peer readvertises routes back to the peer from which it learned them.

This can happen in the following cases:

- Another vendor's routing device advertises the routes back to the sending peer.
- The Junos OS peer's default behavior of not readvertising routes back to the sending peer is overridden by configuring **advertise-peer-as**.
- A provider edge (PE) routing device discards any VPN route that does not have any of the expected route targets.

When **keep all** is configured, the behavior of discarding routes received in the above scenarios is overridden.



**CAUTION:** If you add or remove **keep all** or **keep none** and the peer does not support session restart, the associated BGP sessions are restarted (flapped). To determine if a peer supports refresh, check for **Peer supports Refresh capability** in the output of the **show bgp neighbor** command.

|                                 |                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | By default, BGP retains incoming rejected routes in memory and hides them. If you do not include the <b>keep</b> statement, most routes are retained in the routing table. BGP keeps all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.                        |
| <b>Options</b>                  | <b>all</b> —Retain all routes.<br><br><b>none</b> —Discard routes that were received from a peer and that were rejected by import policy or other sanity checking. When <b>keep none</b> is configured for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">out-delay on page 3018</a></li><li>• <i>Interprovider VPN Example—MP-EBGP Between ISP Peer Routers</i></li><li>• <i>Example: Configuring Conditional Installation of Prefixes in a Routing Table</i></li></ul>                                                                           |

## labeled-unicast (Protocols BGP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> labeled-unicast {     accepted-prefix-limit {         maximum <i>number</i>;         teardown &lt;<i>percentage</i>&gt; &lt;idle-timeout (forever   <i>minutes</i>)&gt;;     }     aggregate-label {         community <i>community-name</i>;     }     explicit-null {         connected-only;     }     prefix-limit {         maximum <i>number</i>;         teardown &lt;<i>percentage</i>&gt; &lt;idle-timeout (forever   <i>minutes</i>)&gt;;     }     protection;     resolve-vpn;     rib (inet.3   inet6.3);     rib-group <i>group-name</i>; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <pre> [edit logical-systems <i>logical-system-name</i> protocols bgp <i>family</i> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <i>family</i> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>family</i> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> (inet   inet6)], [edit protocols bgp <i>family</i> (inet   inet6)], [edit protocols bgp group <i>group-name</i> <i>family</i> (inet   inet6)], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>family</i> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> (inet   inet6)] </pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Configure the family type to be labeled-unicast.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

**Related Documentation** • *Examples: Configuring Multiprotocol BGP*

## local-address (Protocols BGP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>local-address address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | <p>Specify the address of the local end of a BGP session. This address is used to accept incoming connections to the peer and to establish connections to the remote peer. When none of the operational interfaces are configured with the specified local address, a session with a BGP peer is placed in the idle state.</p> <p>You generally configure a local address to explicitly configure the system's IP address from BGP's point of view. This IP address can be either an IPv6 or IPv4 address. Typically, an IP address is assigned to a loopback interface, and that IP address is configured here.</p> <p>For internal BGP (IBGP) peering sessions, generally the loopback interface (lo0) is used to establish connections between the IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address, the IBGP peering session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peering session also goes up and down. Thus, the loopback interface provides fault tolerance in case the physical interface or the link goes down, if the device has link redundancy.</p> <p>When a device peers with a remote device's loopback interface address, the local device expects BGP update messages to come from (be sourced by) the remote device's loopback interface address. The <b>local-address</b> statement enables you to specify the source information in BGP update messages. If you omit the <b>local-address</b> statement, the expected source of BGP update messages is based on the device's source address selection rules, which normally result in the egress interface address being the expected source of update messages. When this happens, the peering session is not established because a mismatch exists between the expected source address (the egress interface</p> |

of the peer) and the actual source (the loopback interface of the peer). To ensure that the expected source address matches the actual source address, specify the loopback interface address in the **local-address** statement.



**NOTE:** Although a BGP session can be established when only one of the paired routing devices has **local-address** configured, we strongly recommend that you configure **local-address** on both paired routing devices for IBGP and multihop EBGP sessions. The **local-address** statement ensures that deterministic fixed addresses are used for the BGP session end-points.

If you include the **default-address-selection** statement in the configuration, the software chooses the system default address as the source for most locally generated IP packets. For protocols in which the local address is unconstrained by the protocol specification, for example IBGP and multihop EBGP, if you do not configure a specific local address when configuring the protocol, the local address is chosen using the same methods as other locally generated IP packets.

**Default** If you do not configure a local address, BGP uses the routing device's source address selection rules to set the local address.

**Options** **address**—IPv6 or IPv4 address of the local end of the connection.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Internal BGP Peering Sessions on Logical Systems*
- *Example: Configuring Internal BGP Peer Sessions*
- *Understanding Internal BGP Peering Sessions*
- [router-id on page 3514](#)

## local-as

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>local-as <i>autonomous-system</i> &lt;loops <i>number</i>&gt; &lt;private   alias&gt; &lt;no-prepend-global-as&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>alias</b> option introduced in Junos OS Release 9.5.</p> <p><b>no-prepend-global-as</b> option introduced in Junos OS Release 9.6.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>         | <p>Specify the local autonomous system (AS) number. An AS is a set of routing devices that are under a single technical administration and generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices.</p> <p>Internet service providers (ISPs) sometimes acquire networks that belong to a different AS. When this occurs, there is no seamless method for moving the BGP peers of the acquired network to the AS of the acquiring ISP. The process of configuring the BGP peers with the new AS number can be time-consuming and cumbersome. In this case, it might not be desirable to modify peer arrangements or configuration. During this kind of transition period, it can be useful to configure BGP-enabled devices in the new AS to use the former AS number in BGP updates. This former AS number is called a <i>local</i> AS.</p>                                                                                                                                                                                                                                                                                                                                                      |



**NOTE:** If you are using BGP on the routing device, you must configure an AS number before you specify the local as number.

In Junos OS Release 9.1 and later, the AS numeric range in plain-number format is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For

example, the 4-byte AS number of 65546 in plain-number format is represented as 1.10 in the AS-dot notation format.

**Options** **alias**—(Optional) Configure the local AS as an alias of the global AS number configured for the router at the **[edit routing-options]** hierarchy level. As a result, a BGP peer considers any local AS to which it is assigned as equivalent to the primary AS number configured for the routing device. When you use the **alias** option, only the AS (global or local) used to establish the BGP session is prepended in the AS path sent to the BGP neighbor.

**autonomous-system**—AS number.

**Range:** 1 through 4,294,967,295 ( $2^{32} - 1$ ) in plain-number format

**Range:** 0.0 through 65535.65535 in AS-dot notation format

**loops number**—(Optional) Specify the number of times detection of the AS number in the AS\_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.



**NOTE:** If you configure the local AS values for any BGP group, the detection of routing loops is performed using both the AS and the local AS values for all BGP groups.

If the local AS for the EBGp or IBGP peer is the same as the current AS, do not use the **local-as** statement to specify the local AS number.

When you configure the local AS within a VRF, this impacts the AS path loop-detection mechanism. All of the **local-as** statements configured on the device are part of a single AS domain. The AS path loop-detection mechanism is based on looking for a matching AS present in the domain.

**Range:** 1 through 10

**Default:** 1

**no-prepend-global-as**—(Optional) Specify to strip the global AS and to prepend only the local AS in AS paths sent to external peers.

**private**—(Optional) Configure to use the local AS only during the establishment of the BGP session with a BGP neighbor but to hide it in the AS path sent to external BGP peers. Only the global AS is included in the AS path sent to external peers.



**NOTE:** The **private** and **alias** options are mutually exclusive. You cannot configure both options with the same **local-as** statement.



**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Examples: Configuring BGP Local AS*
- *Example: Configuring a Local AS for EBGp Sessions*
- [autonomous-system on page 2960](#)
- [family on page 2979](#)

## local-interface (IPv6)

**Syntax** `local-interface interface-name;`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols bgp [group group-name neighbor ipv6-link-local-address](#)],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* [neighbor ipv6-link-local-address](#)],  
[edit protocols bgp group *group-name* [neighbor ipv6-link-local-address](#)],  
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* [neighbor ipv6-link-local-address](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Specify the interface name of the EBGp peer that uses IPv6 link-local addresses. This peer is link-local in scope.

**Options** *interface-name*—Interface name of the EBGp IPv6 peer.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Internal BGP Peering Sessions on Logical Systems*
- *Example: Configuring Internal BGP Peer Sessions*
- *Example: Configuring External BGP on Logical Systems with IPv6 Interfaces*
- *Understanding Internal BGP Peering Sessions*

## local-preference

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>local-preference local-preference;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems logical-system-name protocols bgp],</code><br><code>[edit logical-systems logical-system-name protocols bgp group group-name],</code><br><code>[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],</code><br><code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],</code><br><code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name],</code><br><code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name neighbor address],</code><br><code>[edit protocols bgp],</code><br><code>[edit protocols bgp group group-name],</code><br><code>[edit protocols bgp group group-name neighbor address],</code><br><code>[edit routing-instances routing-instance-name protocols bgp],</code><br><code>[edit routing-instances routing-instance-name protocols bgp group group-name],</code><br><code>[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Modify the value of the <b>LOCAL_PREF</b> path attribute, which is a metric used by BGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.</p> <p>The <b>LOCAL_PREF</b> path attribute always is used in inbound routing policy and is advertised to internal BGP peers and to neighboring confederations. It is never advertised to external BGP peers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | If you omit this statement, the <b>LOCAL_PREF</b> path attribute, if present, is not modified.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>local-preference</b>—Preference to assign to routes learned from BGP or from the group or peer.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> If the <b>LOCAL_PREF</b> path attribute is present, do not modify its value. If a BGP route is received without a <b>LOCAL_PREF</b> attribute, the route is handled locally (it is stored in the routing table and advertised by BGP) as if it were received with a <b>LOCAL_PREF</b> value of 100. By default, non-BGP routes that are advertised by BGP are advertised with a <b>LOCAL_PREF</b> value of 100.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring the Local Preference Value for BGP Routes</i></li><li>• <i>Understanding Internal BGP Peering Sessions</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

- [preference on page 3024](#)

## log-updown (Protocols BGP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | log-updown;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <a href="#">group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <a href="#">group</a> <i>group-name</i> <a href="#">neighbor</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <a href="#">neighbor</a> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <a href="#">group</a> <i>group-name</i>],</p> <p>[edit protocols bgp <i>group</i> <i>group-name</i> <a href="#">neighbor</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">group</a> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <a href="#">neighbor</a> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Specify to generate a log a message whenever a BGP peer makes a state transition. Messages are logged using the system logging mechanism located at the <b>[edit system syslog]</b> hierarchy level.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Preventing BGP Session Resets</i></li> <li>• <i>Junos OS Administration Library for Routing Devices</i></li> <li>• <a href="#">traceoptions on page 3031</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## metric-out (Protocols BGP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>metric-out (<i>metric</i>   minimum-igp <i>offset</i>   igp (delay-med-update   <i>offset</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>]</p>                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Option <b>delay-med-update</b> introduced in Junos OS Release 9.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>         | <p>Specify the metric for all routes sent using the multiple exit discriminator (MED, or <b>MULTI_EXIT_DISC</b>) path attribute in update messages. This path attribute is used to discriminate among multiple exit points to a neighboring AS. If all other factors are equal, the exit point with the lowest metric is preferred.</p> <p>You can specify a constant metric value by including the <b>metric</b> option. For configurations in which a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the <b>multihop</b> command—you can specify a variable metric by including the <b>minimum-igp</b> or <b>igp</b> option.</p> <p>You can increase or decrease the variable metric calculated from the IGP metric (either from the <b>igp</b> or <b>minimum-igp</b> statement) by specifying a value for <b>offset</b>. The metric is increased by specifying a positive value for <b>offset</b>, and decreased by specifying a negative value for <b>offset</b>.</p> <p>In Junos OS Release 9.0 and later, you can specify that a BGP group or peer not advertise updates for the MED path attributes used to calculate IGP costs for BGP next hops unless the MED is lower. You can also configure an interval to delay when MED updates are sent by including the <b>med-igp-update-interval minutes</b> statement at the <b>[edit routing-options]</b> hierarchy level.</p> |
| <b>Options</b>             | <p><b>delay-med-update</b>—Specify that a BGP group or peer configured with the <b>metric-out igp</b> statement not advertise MED updates unless the current MED value is lower than</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

the previously advertised MED value, or another attribute associated with the route has changed, or the BGP peer is responding to a refresh route request.



**NOTE:** You cannot configure the `delay-med-update` statement at the global BGP level.

**igp**—Set the metric to the most recent metric value calculated in the IGP to get to the BGP next hop. Routes learned from an EBGP peer usually have a next hop on a directly connected interface and thus the IGP value is equal to zero. This is the value advertised.

**metric**—Primary metric on all routes sent to peers.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**Default:** No metric is sent.

**minimum-igp**—Set the metric to the minimum metric value calculated in the IGP to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value. When you change a neighbor's export policy from any configuration to a configuration that sets the minimum IGP offset on an exported route, the advertised MED is not updated if the value would increase as a result, even if the previous configuration does not use a minimum IGP-based MED value. This behavior helps to prevent unnecessary route flapping when an IGP cost changes, by not forcing a route update if the metric value increases past the previous lowest known value.

**offset**—Increases or decreases the metric by this value.

**Range:**  $-2^{31}$  through  $2^{31} - 1$

**Default:** None

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates</i></li> <li>• <i>Examples: Configuring BGP MED</i></li> <li>• <i>Example: Configuring the MED Attribute Directly</i></li> <li>• <i>Understanding the MED Attribute</i></li> <li>• <a href="#">med-igp-update-interval on page 3481</a></li> </ul> |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## mtu-discovery

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mtu-discovery;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>                                                                                                                                                                                                                            |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure TCP path maximum transmission unit (MTU) discovery.</p> <p>TCP path MTU discovery enables BGP to automatically discover the best TCP path MTU for each BGP session. In Junos OS, TCP path MTU discovery is disabled by default for all BGP neighbor sessions.</p> <p>When MTU discovery is disabled, TCP sessions that are not directly connected transmit packets of 512-byte maximum segment size (MSS). These small packets minimize the chances of packet fragmentation at a device along the path to the destination. However, because most links use an MTU of at least 1500 bytes, 512-byte packets do not result in the most efficient use of link bandwidth. For directly connected EBGP sessions, MTU mismatches prevent the BGP session from being established. As a workaround, enable path MTU discovery within the EBGP group.</p> <p>Path MTU discovery dynamically determines the MTU size on the network path between the source and the destination, with the goal of avoiding IP fragmentation. Path MTU discovery works by setting the Don't Fragment (DF) bit in the IP headers of outgoing packets. When a device along the path has an MTU that is smaller than the packet, the device drops the packet. The device also sends back an ICMP Fragmentation Needed (Type 3, Code 4) message that contains the device's MTU, thus allowing the source to reduce its path MTU appropriately. The process repeats until the MTU is small enough to traverse the entire path without fragmentation.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

- Related Documentation**
- *Example: Limiting TCP Segment Size for BGP*
  - *Configuring Junos OS for IPv6 Path MTU Discovery*
  - *Configuring Junos OS for Path MTU Discovery on Outgoing GRE Tunnel Connections*

## multihop

---

**Syntax**    multihop {  
              no-nexthop-change;  
              ttl *ttl-value*;  
              }

**Hierarchy Level**    [edit logical-systems *logical-system-name* protocols bgp],  
                          [edit logical-systems *logical-system-name* protocols bgp *group group-name*],  
                          [edit logical-systems *logical-system-name* protocols bgp *group group-name*  
                              *neighbor address*],  
                          [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
                              bgp],  
                          [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
                              bgp *group group-name*],  
                          [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
                              bgp group *group-name neighbor address*],  
                          [edit protocols bgp],  
                          [edit protocols bgp *group group-name*],  
                          [edit protocols bgp group *group-name neighbor address*],  
                          [edit routing-instances *routing-instance-name* protocols bgp],  
                          [edit routing-instances *routing-instance-name* protocols bgp *group group-name*],  
                          [edit routing-instances *routing-instance-name* protocols bgp group *group-name*  
                              *neighbor address*]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description**    Configure an EBGp multihop session.

For Layer 3 VPNs, you configure the EBGp multihop session between the PE and CE routing devices. This allows you to configure one or more routing devices between the PE and CE routing devices.

An external confederation peer is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case because multihop behavior is implied.

If you have external BGP confederation peer-to-loopback addresses, you still need the multihop configuration.



**NOTE:** You cannot configure the `accept-remote-nexthop` statement at the same time.

---

**Default**    If you omit this statement, all EBGp peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or “regular,” BGP session), and the default time-to-live (TTL) value is 1.



The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring EBGp Multihop Sessions</i></li><li>• <i>Configuring EBGp Multihop Sessions Between PE and CE Routers in Layer 3 VPNs</i></li><li>• <a href="#">accept-remote-nextthop on page 2949</a></li><li>• <i>no-nextthop-change</i></li><li>• <i>tth</i></li></ul> |

## multipath (Protocols BGP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>multipath {<br/>  multiple-as;<br/>  vpn-unequal-cost equal-external-internal;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Allow load sharing among multiple EBGP paths and multiple IBGP paths. A path is considered a BGP equal-cost path (and will be used for forwarding) if a tie-break is performed. The tie-break is performed after the BGP route path selection step that chooses the next-hop path that is resolved through the IGP route with the lowest metric. All paths with the same neighboring AS, learned by a multipath-enabled BGP neighbor, are considered.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>multiple-as</b>—Disable the default check requiring that paths accepted by BGP multipath must have the same neighboring AS.</p> <p><b>vpn-unequal-cost equal-external-internal</b>—Enable load-balancing in a Layer 3 VPN with unequal cost paths.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding BGP Path Selection</i></li><li>• <i>Example: Load Balancing BGP Traffic</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## neighbor (Protocols BGP)

```
Syntax neighbor address {
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    as-override;
    authentication-algorithm algorithm;
    authentication-key key;
    authentication-key-chain key-chain;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family {
        (inet | inet6 | inet-mvpn | inet6-mpvn | inet-vpn | inet6-vpn | iso-vpn | l2-vpn) {
            (any | flow | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                damping;
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                rib-group group-name;
                topology name {
                    community {
                        target identifier;
                    }
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib inet.3;
        }
    }
}
```

```
    rib-group group-name;
    topology name {
        community {
            target identifier;
        }
    }
}
}
route-target {
    advertise-default;
    external-paths number;
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
signaling {
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-interface interface-name;
local-preference preference;
log-updown;
metric-out (metric | minimum-igp <offset> | igp <offset>);
mtu-discovery;
multihop <ttl-value>;
multipath {
    multiple-as;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
tcp-aggressive-transmission;
tcp-mss segment-size;
traceoptions {
```

```

    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  vpn-apply-export;
}

```

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>bgp <b>group</b> <i>group-name</i> ],<br>[edit protocols bgp <b>group</b> <i>group-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> ]                                                                                                                                                             |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                     |
| Description              | <p>Explicitly configure a neighbor (peer). To configure multiple BGP peers, include multiple <b>neighbor</b> statements.</p> <p>By default, the peer's options are identical to those of the group. You can override these options by including peer-specific option statements within the <b>neighbor</b> statement.</p> <p>The <b>neighbor</b> statement is one of the statements you can include in the configuration to define a minimal BGP configuration on the routing device. (You can include an <b>allow all</b> statement in place of a <b>neighbor</b> statement.)</p> |
| Options                  | <p><b>address</b>—IPv6 or IPv4 address of a single peer.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Related Documentation    | <ul style="list-style-type: none"> <li>• <i>BGP Feature Guide for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## no-aggregator-id

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-aggregator-id;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit protocols bgp],<br/>[edit protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Prevent different routing devices within an AS from creating aggregate routes that contain different AS paths.</p> <p>Junos OS performs route aggregation, which is the process of combining the characteristics of different routes so that only a single route is advertised. Aggregation reduces the amount of information that BGP must store and exchange with other BGP systems. When aggregation occurs, the local routing device adds the local AS number and the router ID to the aggregator path attribute. The <b>no-aggregator-id</b> statement causes Junos OS to place a 0 in the router ID field and thus eliminate the possibility of having multiple aggregate advertisements in the network, each with different path information.</p>                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>                  | If you omit this statement, the router ID is included in the BGP aggregator path attribute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>BGP Messages Overview</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## no-client-reflect

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-client-reflect;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Disable intracluster route redistribution by the system acting as the route reflector. Include this statement when the client cluster is fully meshed to prevent the sending of redundant route advertisements. Route reflection provides a way to decrease BGP control traffic and minimizing the number of update messages sent within the AS.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring BGP Route Reflectors</i></li> <li>• <a href="#">cluster on page 2970</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## no-validate

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>no-validate <i>policy-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | <code>[edit protocols bgp group <i>group-name</i> <b>family</b> (inet   inet flow)],</code><br><code>[edit protocols bgp group <i>group-name</i> neighbor address <b>family</b> (inet   inet flow)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>family</b> (inet   inet flow)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor address <b>family</b> (inet   inet flow)]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>         | <p>When BGP is carrying flow-specification network layer reachability information (NLRI) messages, the <b>no-validate</b> statement omits the flow route validation procedure after packets are accepted by a policy.</p> <p>The receiving BGP-enabled device accepts a flow route if it passes the following criteria:</p> <ul style="list-style-type: none"><li>• The originator of a flow route matches the originator of the best match unicast route for the destination address that is embedded in the route.</li><li>• There are no more specific unicast routes, when compared to the destination address of the flow route, for which the active route has been received from a different next-hop autonomous system.</li></ul> <p>The first criterion ensures that the filter is being advertised by the next-hop used by unicast forwarding for the destination address embedded in the flow route. For example, if a flow route is given as 10.1.1.1, proto=6, port=80, the receiving BGP-enabled device selects the more specific unicast route in the unicast routing table that matches the destination prefix 10.1.1.1/32. On a unicast routing table containing 10.1/16 and 10.1.1/24, the latter is chosen as the unicast route to compare against. Only the active unicast route entry is considered. This follows the concept that a flow route is valid if advertised by the originator of the best unicast route.</p> <p>The second criterion addresses situations in which a given address block is allocated to different entities. Flows that resolve to a best-match unicast route that is an aggregate route are only accepted if they do not cover more specific routes that are being routed to different next-hop autonomous systems.</p> <p>You can bypass the validation process and use your own specific import policy. To disable the validation procedure and use an import policy instead, include the <b>no-validate</b> statement in the configuration.</p> <p>Flow routes configured for VPNs with family <b>inet-vpn</b> are not automatically validated, so the <b>no-validate</b> statement is not supported at the <code>[edit protocols bgp group <i>group-name</i> <b>family</b> <i>inet-vpn</i>]</code> hierarchy level. No validation is needed if the flow routes are configured locally between devices in a single AS.</p> |



**Options**    *policy-name*—Import policy to match NLRI messages.

**Required Privilege**    routing—To view this statement in the configuration.  
**Level**    routing-control—To add this statement to the configuration.

**Related**    • *Example: Configuring Flow Routes*  
**Documentation**

## out-delay

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>out-delay seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>     | <code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code><br><code>  <i>neighbor address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  bgp],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  bgp <i>group group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  bgp <i>group group-name neighbor address</i>],</code><br><code>[edit protocols bgp],</code><br><code>[edit protocols bgp <i>group group-name</i>],</code><br><code>[edit protocols bgp <i>group group-name neighbor address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code><br><code>  <i>neighbor address</i>]</code>                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>         | <p>Control how often BGP and the routing table exchange route information by specifying how long a route must be present in the Junos OS routing table before it is exported to BGP. Use this time delay to help bundle routing updates and to avoid sending updates too often.</p> <p>Alternatively or in addition, external BGP (EBGP) sessions can also use the route-flap damping mechanism upon the reception of BGP messages coming from an external neighbor.</p> <p>BGP stores the route information it receives from update messages in the routing table, and the routing table exports active routes from the routing table into BGP. BGP then advertises the exported routes to its peers. The <b>out-delay</b> statement enables a form of rate limiting. The delay is added to each update for each prefix individually. When a routing device changes its best path to a destination prefix, the device does not inform its peer about the change unless the route has been present in its routing table for the specified <b>out-delay</b>. If you use <b>out-delay</b> to perform rate-limiting, you can expect a less bursty pattern of updates. You will see a pattern in which updates arrive in a steady flow, and two updates for the same prefix are always spaced by at least the <b>out-delay</b> timer value (for example, 30 seconds). Thus, the <b>out-delay</b> setting is useful for limiting oscillation (sometimes called <i>churn</i>) in a network. Keep in mind that, regardless of the <b>out-delay</b> setting, BGP peers exchange routes immediately after neighbor establishment. The <b>out-delay</b> setting is only designed to delay the exchange of routes between BGP and the local routing table.</p> |

Caution is warranted because an **out-delay** can delay convergence. If your network is configured in a way that avoids oscillation, setting an **out-delay** is not necessary.

When configured, the **out-delay** value displays as **Outbound Timer** when using **show bgp group** or **show bgp group neighbor** commands.


**Default** By default, the exchange of route information between BGP and the routing table occurs immediately after the routes are received. This immediate exchange of route information might cause instabilities in the network reachability information. If you omit this statement, routes are exported to BGP immediately after they have been added to the routing table.

**Options** *seconds*—Output delay time.  
**Range:** 0 through 65,535 seconds  
**Default:** 0 seconds

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • [keep on page 2993](#)

## outbound-route-filter

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> outbound-route-filter {     <b>bgp-orf-cisco-mode</b>;     prefix-based {         accept {             (inet   inet6);         }     } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor   <i>address</i>] </pre> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Configure a BGP peer to accept outbound route filters from a remote peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>accept</b>—Specify that outbound route filters from a BGP peer be accepted.</p> <p><b>inet</b>—Specify that IPv4 prefix-based outbound route filters be accepted.</p> <p><b>inet6</b>—Specify that IPv6 prefix-based outbound route filters be accepted.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|                                 | <p> <b>NOTE:</b> You can specify that both IPv4 and IPv6 outbound route filters be accepted.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|                                 | <p><b>prefix-based</b>—Specify that prefix-based filters be accepted.</p> <p>The <b>bgp-orf-cisco-mode</b> statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

**Related Documentation** • *Example: Configuring BGP Prefix-Based Outbound Route Filtering*

## passive (Protocols BGP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | passive;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure the routing device so that active open messages are not sent to the peer. Once you configure the routing device to be passive, the routing device will wait for the peer to issue an open request before a message is sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>                  | If you omit this statement, all explicitly configured peers are active, and each peer periodically sends open requests until its peer responds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | • <i>Example: Preventing BGP Session Flaps When VPN Families Are Configured</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## peer-as (Protocols BGP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>peer-as <i>autonomous-system</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>         | <p>Specify the neighbor (peer) autonomous system (AS) number.</p> <p>For EBGP, the peer is in another AS, so the AS number you specify in the <b>peer-as</b> statement must be different from the local router's AS number, which you specify in the <b>autonomous-system</b> statement. For IBGP, the peer is in the same AS, so the two AS numbers that you specify in the <b>autonomous-system</b> and <b>peer-as</b> statements must be the same.</p> <p>The AS numeric range in plain-number format has been extended in Junos OS Release 9.1 to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of the Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i>&lt;16-bit high-order value in decimal&gt;.&lt;16-bit low-order value in decimal&gt;</i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p> <p>With the introduction of 4-byte AS numbers, you might have a combination of routers that support 4-byte AS numbers and 2-byte AS numbers. For more information about what happens when establishing BGP peer relationships between 4-byte and 2-byte capable routers, see the following topics:</p> |

- *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview.*

**Options** *autonomous-system*—AS number.

**Range:** 1 through 4,294,967,295 ( $2^{32} - 1$ ) in plain-number format for 4-byte AS numbers

**Range:** 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)


**Range:** 0.0 through 65535.65535 in AS-dot notation format for 4-byte AS numbers

**Required Privilege** routing—To view this statement in the configuration.

**Level** routing-control—To add this statement to the configuration.

**Related  
Documentation**

## preference (Protocols BGP)

|                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                      | <code>preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>                                                                                                                                             | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>                                                                                                                                         | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                                                                                                                                 | <p>Specify the preference for routes learned from BGP.</p> <p>At the BGP global level, the preference statement sets the preference for routes learned from BGP. You can override this preference in a BGP group or peer preference statement.</p> <p>At the group or peer level, the preference statement sets the preference for routes learned from the group or peer. Use this statement to override the preference set in the BGP global preference statement when you want to favor routes from one group or peer over those of another.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <div>  <b>NOTE:</b> Do not set preference2 for BGP route-policy.         </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                                                                                                                                     | <p><b>preference</b>—Preference to assign to routes learned from BGP or from the group or peer.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 170 for the primary preference</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b>                                                                                                                                    | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                       | <ul style="list-style-type: none"> <li>• <a href="#">local-preference on page 3002</a></li> <li>• <i>Example: Configuring the Preference Value for BGP Routes</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |



## prefix-limit

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>prefix-limit {     maximum <i>number</i>;     teardown &lt;<i>percentage</i>&gt; &lt;idle-timeout (forever   <i>minutes</i>)&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>family</b> (inet   inet6) (any   flow   labeled-unicast   multicast   unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <b>family</b> (inet   inet6) (any   flow   labeled-unicast   multicast   unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>family</b> (inet   inet6) (any   flow   labeled-unicast   multicast   unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>family</b> (inet   inet6) (any   flow   labeled-unicast   multicast   unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>family</b> (inet   inet6) (any   flow   labeled-unicast   multicast   unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>family</b> (inet   inet6) (any   flow   labeled-unicast   multicast   unicast)],</p> <p>[edit protocols bgp <b>family</b> (inet   inet6) (any   flow   labeled-unicast   multicast   unicast)],</p> <p>[edit protocols bgp group <i>group-name</i> <b>family</b> (inet   inet6) (any   labeled-unicast   multicast   unicast)],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>family</b> (inet   inet6) (any   flow   labeled-unicast   multicast   unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>family</b> (inet   inet6) (any   flow   labeled-unicast   multicast   unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>family</b> (inet   inet6) (any   flow   labeled-unicast   multicast   unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>family</b> (inet   inet6) (any   flow   labeled-unicast   multicast   unicast)]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | <p>Limit the number of prefixes received on a BGP peer session and a rate-limit logging when injected prefixes exceed a set limit.</p> <p>This functionality is identical to the <b>accepted-prefix-limit</b> functionality except that it operates against received prefixes rather than accepted prefixes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>             | <p><b>maximum <i>number</i></b>—When you set the maximum number of prefixes, a message is logged when that number is exceeded.</p> <p><b>Range:</b> 1 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>teardown &lt;<i>percentage</i>&gt;</b>—If you include the <b>teardown</b> statement, the session is torn down when the maximum number of prefixes is reached. If you specify a percentage, messages are logged when the number of prefixes exceeds that percentage. After the session is torn down, it is reestablished in a short time unless you include the <b>idle-timeout</b> statement. Then the session can be kept down for a specified amount</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

of time, or forever. If you specify **forever**, the session is reestablished only after you issue a **clear bgp neighbor** command.

**Range:** 1 through 100


**idle-timeout (*forever* | *timeout-in-minutes*)**—(Optional) If you include the **idle-timeout** statement, the session is torn down for a specified amount of time, or forever. If you specify a period of time, the session is allowed to reestablish after this timeout period. If you specify **forever**, the session is reestablished only after you intervene with a **clear bgp neighbor** command.

**Range:** 1 through 2400

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                      |                                                                                          |
|----------------------|------------------------------------------------------------------------------------------|
| <b>Related</b>       | <ul style="list-style-type: none"><li>• <i>accepted-prefix-limit</i></li></ul>           |
| <b>Documentation</b> | <ul style="list-style-type: none"><li>• <i>Understanding Multiprotocol BGP</i></li></ul> |

## remove-private

|                                                                                                                                                                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                          | remove-private all replace nearest;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                 | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                             | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                     | <p>When advertising AS paths to remote systems, have the local system strip private AS numbers from the AS path. The numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. If the AS path contains the AS number of the external BGP (EBGP) neighbor, BGP does not remove the private AS number.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <div>  <p><b>NOTE:</b> As of Junos OS 10.0R2 and higher, if there is a need to send prefixes to an EBGP peer that has an AS number that matches an AS number in the AS path, consider using the <code>as-override</code> statement instead of the <code>remove-private</code> statement.</p> </div>                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <p>The operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.</p> <p>The Junos OS recognizes the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document.</p> <p>The set of reserved AS numbers is in the range from 64,512 through 65,535.</p> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                         | <p><b>all</b>—Remove all private AS numbers from the original path. Do not stop the process of removing private AS numbers, even if a public AS number is encountered.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**nearest**—When you use the **all** and **replace** options, choose the last (right-most) public AS number encountered in the original AS path for the replacement value, as the AS path is processed from left to right. If no public AS number is encountered, the default replacement value is used. (See the **replace** option for information about the default replacement value.)

**replace**—When you use the **all** option, instead of removing private AS numbers, perform a replace operation. The default replacement value for the private AS number is the local AS number at the BGP group level for the BGP peer. If you are unsure about the replacement value, check the local AS value displayed in the output of the **show bgp group group-name** command.

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |                                                                                                             |
|------------------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Example: Removing Private AS Numbers from AS Paths</i></li></ul> |
|------------------------------|-------------------------------------------------------------------------------------------------------------|

## rib-group (Protocols BGP)


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rib-group group-name;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>family</b> inet (labeled-unicast   unicast   multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <b>family</b> inet (labeled-unicast   unicast   multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>family</b> inet (labeled-unicast   unicast   multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>family</b> inet (labeled-unicast   unicast   multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>family</b> inet (labeled-unicast   unicast   multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>family</b> inet (labeled-unicast   unicast   multicast)],</p> <p>[edit protocols bgp <b>family</b> inet (labeled-unicast   unicast   multicast)],</p> <p>[edit protocols bgp group <i>group-name</i> <b>family</b> inet (labeled-unicast   unicast   multicast)],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>family</b> inet (labeled-unicast   unicast   multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>family</b> inet (labeled-unicast   unicast   multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>family</b> inet (labeled-unicast   unicast   multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>family</b> inet (labeled-unicast   unicast   multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Add unicast prefixes to unicast and multicast tables.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>group-name</b> —Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. You generally specify only one routing table group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li> <li>• <i>Example: Importing Direct and Static Routes Into a Routing Instance</i></li> <li>• <i>Understanding Multiprotocol BGP</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## tcp-mss (Protocols BGP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>tcp-mss <i>segment-size</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor</code><br><code>    <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    bgp],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    bgp group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    bgp group <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit protocols bgp],</code><br><code>[edit protocol bgp group <i>group-name</i>],</code><br><code>[edit protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor</code><br><code>    <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure the maximum segment size (MSS) for the TCP connection for BGP neighbors.<br><br>The MSS is only valid in increments of 2 KB. The value used is based on the value set, but is rounded down to the nearest multiple of 2048.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b><i>segment-size</i></b> —MSS for the TCP connection.<br><b>Range:</b> 1 through 4096                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Limiting TCP Segment Size for BGP</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## traceoptions (Protocols BGP)

|                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                        | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>                                                                                                                                                               | <pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor   <i>address</i>] </pre> |
| <b>Release Information</b>                                                                                                                                                           | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>4byte-as</b> statement introduced in Junos OS Release 9.2.</p> <p><b>4byte-as</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                                                                                                                                                   | Configure BGP protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <div>  <b>NOTE:</b> The <b>traceoptions</b> statement is not supported on QFabric systems. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>                                                                                                                                                                       | <p>The default BGP protocol-level tracing options are inherited from the routing protocols <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level. The default group-level trace options are inherited from the BGP protocol-level <b>traceoptions</b> statement. The default peer-level trace options are inherited from the group-level <b>traceoptions</b> statement.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                                                                                                                                                                       | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>name</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place BGP tracing output in the file <b>bgp-log</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file.0*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

#### BGP Tracing Flags

- **4byte-as**—4-byte AS events.
- **bfd**—BFD protocol events.
- **damping**—Damping operations.
- **graceful-restart**—Graceful restart events.
- **keepalive**—BGP keepalive messages. If you enable the the BGP **update** flag only, received keepalive messages do not generate a trace message.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **open**—Open packets. These packets are sent between peers when they are establishing a connection.
- **packets**—All BGP protocol packets.
- **refresh**—BGP refresh packets.
- **update**—Update packets. These packets provide routing updates to BGP systems. If you enable only this flag, received keepalive messages do not generate a trace message. Use the **keepalive** flag to generate a trace message for keepalive messages.

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing



**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **filter**—Provide filter trace information. Applies only to **route**, **damping**, and **update** tracing flags.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.<br>routing-control and trace-control—To add this statement to the configuration.                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">log-updown on page 3003</a> statement</li> <li>• <a href="#">Tracing Nonstop Active Routing Synchronization Events on page 2573</a></li> <li>• <i>Understanding Trace Operations for BGP Protocol Traffic</i></li> <li>• <i>Configuring OSPF Refresh and Flooding Reduction in Stable Topologies</i></li> </ul> |

## type (Protocols BGP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>type type;</code>                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>bgp <b>group</b> <i>group-name</i> ],<br>[edit protocols bgp <b>group</b> <i>group-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Specify the type of BGP peer group.</p> <p>When configuring a BGP group, you can indicate whether the group is an IBGP group or an EBGP group. All peers in an IBGP group are in the same AS, while peers in an EBGP group are in different ASs and normally share a subnet.</p>                                                                                                                                    |
| <b>Options</b>                  | <p><b>type</b>—Type of group:</p> <ul style="list-style-type: none"><li>• <b>external</b>—External group, which allows inter-AS BGP routing</li><li>• <b>internal</b>—Internal group, which allows intra-AS BGP routing</li></ul>                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>BGP Feature Guide for Routing Devices</i></li></ul>                                                                                                                                                                                                                                                                                                                         |

# Administration

- [Routine Monitoring on page 3035](#)
- [Operational Commands on page 3037](#)

## Routine Monitoring

- [Monitoring BGP Routing Information on page 3035](#)

### Monitoring BGP Routing Information

Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to monitor BGP routing information on the routing device.

Action

To view BGP routing information in the J-Web interface, select **Monitor > Routing > BGP Information**.

To view BGP routing information in the CLI, enter the following commands:

- `show bgp summary`
- `show bgp neighbor`

Meaning

[Table 318 on page 3035](#) summarizes key output fields in the BGP routing display in the J-Web interface.

Table 318: Summary of Key BGP Routing Output Fields

| Field            | Values                           | Additional Information |
|------------------|----------------------------------|------------------------|
| BGP Peer Summary |                                  |                        |
| Total Groups     | Number of BGP groups.            |                        |
| Total Peers      | Number of BGP peers.             |                        |
| Down Peers       | Number of unavailable BGP peers. |                        |

Table 318: Summary of Key BGP Routing Output Fields (*continued*)

| Field                  | Values                                                                                                                                                                 | Additional Information |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| Unconfigured Peers     | Address of each BGP peer.                                                                                                                                              |                        |
| <b>RIB Summary tab</b> |                                                                                                                                                                        |                        |
| RIB Name               | Name of the RIB group.                                                                                                                                                 |                        |
| Total Prefixes         | Total number of prefixes from the peer, both active and inactive, that are in the routing table.                                                                       |                        |
| Active Prefixes        | Number of prefixes received from the EBGp peers that are active in the routing table.                                                                                  |                        |
| Suppressed Prefixes    | Number of routes received from EBGp peers currently inactive because of damping or other reasons.                                                                      |                        |
| History Prefixes       | History of the routes received or suppressed.                                                                                                                          |                        |
| Dumped Prefixes        | Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. |                        |
| Pending Prefixes       | Number of pending routes.                                                                                                                                              |                        |
| State                  | Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete.  |                        |
| <b>BGP Neighbors</b>   |                                                                                                                                                                        |                        |
| Details                | Click this button to view the selected BGP neighbor details.                                                                                                           |                        |
| Peer Address           | Address of the BGP neighbor.                                                                                                                                           |                        |
| Autonomous System      | AS number of the peer.                                                                                                                                                 |                        |

Table 318: Summary of Key BGP Routing Output Fields (*continued*)

| Field        | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Additional Information                                                                                                                                                                                                                                                                                                 |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Peer State   | <p>Current state of the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—BGP is initiating a TCP connection in an attempt to connect to a peer. If the connection is successful, BGP sends an open message.</li> <li>• <b>Connect</b>—BGP is waiting for the TCP connection to become complete.</li> <li>• <b>Established</b>—The BGP session has been established, and the peers are exchanging BGP update messages.</li> <li>• <b>Idle</b>—This is the first stage of a connection. BGP is waiting for a Start event.</li> <li>• <b>OpenConfirm</b>—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.</li> <li>• <b>OpenSent</b>—BGP has sent an open message and is waiting to receive an open message from the peer.</li> </ul> | <p>Generally, the most common states are <b>Active</b>, which indicates a problem establishing the BGP connection, and <b>Established</b>, which indicates a successful session setup. The other states are transition states, and BGP sessions normally do not stay in those states for extended periods of time.</p> |
| Elapsed Time | Elapsed time since the peering session was last reset.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                                                                                                                                                                                                        |
| Description  | Description of the BGP session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                        |

- Related Documentation**
- [Configuring BGP Sessions \(J-Web Procedure\) on page 2943](#)
  - [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)

## Operational Commands

- `clear bgp damping`
- `clear bgp neighbor`
- `clear bgp table`
- `show bgp bmp`
- `show bgp group`
- `show bgp neighbor`
- `show bgp summary`
- `show policy damping`

## clear bgp damping

---

|                                          |                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                           | <a href="#">Syntax on page 3038</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3038</a>                                                                                                                                                                                                                             |
| Syntax                                   | <code>clear bgp damping</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;prefix&gt;</code>                                                                                                                                                                                                  |
| Syntax (EX Series Switch and QFX Series) | <code>clear bgp damping</code><br><code>&lt;prefix&gt;</code>                                                                                                                                                                                                                                                                            |
| Release Information                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                 |
| Description                              | Clear BGP route flap damping information.                                                                                                                                                                                                                                                                                                |
| Options                                  | <b>none</b> —Clear all BGP route flap damping information.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>prefix</b> —(Optional) Clear route flap damping information for only the specified destination prefix. |
| Required Privilege Level                 | clear                                                                                                                                                                                                                                                                                                                                    |
| Related Documentation                    | <ul style="list-style-type: none"><li>• <a href="#">show policy damping on page 3072</a></li><li>• <a href="#">show route damping on page 3577</a></li></ul>                                                                                                                                                                             |
| List of Sample Output                    | <a href="#">clear bgp damping on page 3038</a>                                                                                                                                                                                                                                                                                           |
| Output Fields                            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                    |

## Sample Output

### clear bgp damping

```
user@host> clear bgp damping
```

## clear bgp neighbor

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3039</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3039</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax</b>                                   | <pre>clear bgp neighbor &lt;as <i>as-number</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;malformed-route&gt; &lt;neighbor&gt; &lt;soft   soft-inbound&gt; &lt;soft-minimum-igp&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>clear bgp neighbor &lt;as <i>as-number</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;malformed-route&gt; &lt;neighbor&gt; &lt;soft   soft-inbound&gt; &lt;soft-minimum-igp&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>malformed-route</b> option introduced in Junos OS Release 13.2.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                              | <p>Perform one of the following tasks:</p> <ul style="list-style-type: none"> <li>• Change the state of one or more BGP neighbors to <b>IDLE</b>. For neighbors in the <b>ESTABLISHED</b> state, this command drops the TCP connection to the neighbors and then reestablishes the connection.</li> <li>• (<b>soft</b> keyword only) Reapply export policies or import policies, respectively, to one or more BGP neighbors without changing their state.</li> <li>• (<b>soft-inbound</b> keyword only) Reapply export policies or import policies, respectively, and send refresh updates to one or more BGP neighbors without changing their state.</li> </ul>                                                                                                   |
| <b>Options</b>                                  | <p><b>none</b>—Change the state of all BGP neighbors to <b>IDLE</b>.</p> <p><b>as <i>as-number</i></b>—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).</p> <p><b>instance <i>instance-name</i></b>—(Optional) Apply this command only to neighbors for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>malformed-route</b>—(Optional) Remove malformed routes. If a specific neighbor is provided, Junos OS removes malformed routes for that particular neighbor. Otherwise, Junos OS removes malformed routes for all BGP neighbors. To find routes that have</p> |

malformed attributes, run the **show route hidden** command, and look for routes marked with **MalformedAttr** in the AS path field.

**neighbor**—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.

**soft**—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.

**soft-inbound**—(Optional) Reapply any import policies and send refresh updates to neighbors without clearing the state.

**soft-minimum-igp**—(Optional) Provides soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.

**Required Privilege  
Level**

clear

**Related  
Documentation**

- [show bgp neighbor on page 3052](#)

**List of Sample Output**

[clear bgp neighbor on page 3040](#)

**Output Fields**

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

**clear bgp neighbor**

```
user@host> clear bgp neighbor
```



## clear bgp table

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                   | <code>clear bgp table <i>table-name</i></code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <code>clear bgp table <i>table-name</i></code>                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>                      | Command introduced in Junos OS Release 9.0.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                     |
| <b>Description</b>                              | Request that BGP refresh routes in a specified routing table.                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                                  | <b><code>logical-system (all   <i>logical-system-name</i>)</code></b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><code>table-name</code></b> —Request that BGP refresh routes in the specified table.                                                                                                                                        |
| <b>Additional Information</b>                   | In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the <b>clear bgp table</b> command to request that BGP refresh routes in a VPN instance table. |
| <b>Required Privilege Level</b>                 | clear                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>                    | <a href="#">clear bgp table private.inet.0 on page 3041</a><br><a href="#">clear bgp table inet.6 logical-system all on page 3041</a><br><a href="#">clear bgp table private.inet.6 logical-system ls1 on page 3041</a><br><a href="#">clear bgp table logical-system all inet.0 on page 3041</a><br><a href="#">clear bgp table logical-system ls2 private.inet.0 on page 3042</a>                      |
| <b>Output Fields</b>                            | This command produces no output.                                                                                                                                                                                                                                                                                                                                                                         |

## Sample Output

[clear bgp table private.inet.0](#)

```
user@host> clear bgp table private.inet.0
```

[clear bgp table inet.6 logical-system all](#)

```
user@host> clear bgp table inet.6 logical-system all
```

[clear bgp table private.inet.6 logical-system ls1](#)

```
user@host> clear bgp table private.inet.6 logical-system ls1
```

[clear bgp table logical-system all inet.0](#)

```
user@host> clear bgp table logical-system all inet.0
```

### clear bgp table logical-system ls2 private.inet.0

```
user@host> clear bgp table logical-system ls2 private.inet.0
```

## show bgp bmp

|                                 |                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show bgp bmp</b>                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.5.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series. |
| <b>Description</b>              | Display information about the BGP Monitoring Protocol (BMP).                                                                                                                                |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                        |
| <b>List of Sample Output</b>    | <a href="#">show bgp bmp on page 3043</a>                                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 319 on page 3043</a> lists the output fields for the <b>show bgp bmp</b> command. Output fields are listed in the approximate order in which they appear.                 |

**Table 319: show bgp bmp Output Fields**

| Field Name                          | Field Description                                                                                                                                              |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>BMP station address/port</b>     | IP address and port number of the monitoring station to which BGP Monitoring Protocol (BMP) statistics are sent.                                               |
| <b>BMP session state</b>            | Status of the BMP session: <b>UP</b> or <b>DOWN</b> .                                                                                                          |
| <b>Memory consumed by BMP</b>       | Memory used by the active BMP session.                                                                                                                         |
| <b>Statistics timeout</b>           | Amount of time, in seconds, between transmissions of BMP data to the monitoring station.                                                                       |
| <b>Memory limit</b>                 | Threshold, in bytes, at which the routing device stops collecting BMP data.                                                                                    |
| <b>Memory-connect retry timeout</b> | Amount of time, in seconds, after which the routing device attempts to resume a BMP session that was ended after the configured memory threshold was exceeded. |

## Sample Output

### show bgp bmp

```

user@host> show bgp bmp
  BMP station address/port: 172.24.24.157+5454
  BMP session state: DOWN
  Memory consumed by BMP: 0
  Statistics timeout: 15
  Memory limit: 10485760
  Memory connect retry timeout: 600

```



## show bgp group

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3045</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3045</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax</b>                                   | <pre>show bgp group &lt;brief   detail   summary&gt; &lt;group-name&gt; &lt;exact-instance instance-name&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt; &lt;rtf&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show bgp group &lt;brief   detail   summary&gt; &lt;group-name&gt; &lt;exact-instance instance-name&gt; &lt;instance instance-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>exact-instance</b> option introduced in Junos OS Release 11.4.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                              | Display information about the configured BGP groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                  | <p><b>none</b>—Display group information about all BGP groups.</p> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output.</p> <p><b>group-name</b>—(Optional) Display group information for the specified group.</p> <p><b>exact-instance instance-name</b>—(Optional) Display information for the specified instance only.</p> <p><b>instance instance-name</b>—(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, <b>cust1</b>, <b>cust11</b>, and <b>cust111</b> are all displayed when you run the <b>show bgp group instance cust1</b> command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>rtf</b>—(Optional) Display BGP group route targeting information.</p> |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>                    | <a href="#">show bgp group on page 3049</a><br><a href="#">show bgp group brief on page 3049</a><br><a href="#">show bgp group detail on page 3050</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

[show bgp group rtf detail on page 3051](#)  
[show bgp group summary on page 3051](#)

**Output Fields** [Table 320 on page 3046](#) describes the output fields for the **show bgp group** command. Output fields are listed in the approximate order in which they appear.

**Table 320: show bgp group Output Fields**

| Field Name                                | Field Description                                                                                                                                                                                                                                      | Level of Output             |
|-------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|
| <b>Group Type or Group</b>                | Type of BGP group: <b>Internal</b> or <b>External</b> .                                                                                                                                                                                                | All levels                  |
| <b>group-index</b>                        | Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.                                                          | <b>rtf detail</b>           |
| <b>AS</b>                                 | AS number of the peer. For internal BGP (IBGP), this number is the same as <b>Local AS</b> .                                                                                                                                                           | <b>brief detail</b><br>none |
| <b>Local AS</b>                           | AS number of the local routing device.                                                                                                                                                                                                                 | <b>brief detail</b><br>none |
| <b>Name</b>                               | Name of a specific BGP group.                                                                                                                                                                                                                          | <b>brief detail</b><br>none |
| <b>Index</b>                              | Unique index number of a BGP group.                                                                                                                                                                                                                    | <b>brief detail</b><br>none |
| <b>Flags</b>                              | Flags associated with the BGP group. This field is used by Juniper Networks customer support.                                                                                                                                                          | <b>brief detail</b><br>none |
| <b>Remove-private options</b>             | Options associated with the <a href="#">remove-private</a> statement.                                                                                                                                                                                  | <b>brief detail</b><br>none |
| <b>Holdtime</b>                           | Maximum number of seconds allowed to elapse between successive keepalive or update messages that BGP receives from a peer in the BGP group, after which the connection to the peer is closed and routing devices through that peer become unavailable. | <b>brief detail</b><br>none |
| <b>Export</b>                             | Export policies configured for the BGP group with the <b>export</b> statement.                                                                                                                                                                         | <b>brief detail</b><br>none |
| <b>MED tracks IGP metric update delay</b> | Time, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire                                                                                                 | All levels                  |
| <b>Traffic Statistics Interval</b>        | Time between sample periods for labeled-unicast traffic statistics, in seconds.                                                                                                                                                                        | <b>brief detail</b><br>none |
| <b>Total peers</b>                        | Total number of peers in the group.                                                                                                                                                                                                                    | <b>brief detail</b><br>none |

Table 320: show bgp group Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Established</b>                     | Number of peers in the group that are in the established state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels      |
| <b>Active/Received/Accepted/Damped</b> | <p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established in the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> <li>If a peer is not established, the field shows the state of the peer session: <b>Active</b>, <b>Connect</b>, or <b>Idle</b>.</li> <li>If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the <b>inet.0</b> (main) and <b>inet.2</b> (multicast) routing tables. For example, <b>8/10/10/2</b> and <b>2/4/4/0</b> indicate the following: <ul style="list-style-type: none"> <li>8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the <b>inet.0</b> routing table.</li> <li>2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the <b>inet.2</b> routing table.</li> </ul> </li> </ul> | <b>summary</b>  |
| <b>ip-addresses</b>                    | List of peers who are members of the group. The address is followed by the peer's port number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels      |
| <b>Route Queue Timer</b>               | Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail</b>   |
| <b>Route Queue</b>                     | Number of prefixes that are queued up for sending to the peers in the group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail</b>   |
| <b>inet.number</b>                     | <p>Number of active, received, accepted, and damped routes in the routing table. For example, <b>inet.0: 7/10/9/0</b> indicates the following:</p> <ul style="list-style-type: none"> <li>7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the <b>inet.0</b> routing table.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | none            |

Table 320: show bgp group Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output     |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| <b>Table inet.number</b> | Information about the routing table. <ul style="list-style-type: none"> <li>• <b>Received prefixes</b>—Total number of prefixes from the peer, both active and inactive, that are in the routing table.</li> <li>• <b>Active prefixes</b>—Number of prefixes received from the peer that are active in the routing table.</li> <li>• <b>Suppressed due to damping</b>—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.</li> <li>• <b>Advertised prefixes</b>—Number of prefixes advertised to a peer.</li> <li>• <b>Received external prefixes</b>—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table.</li> <li>• <b>Active external prefixes</b>—Number of prefixes received from the EBGP peers that are active in the routing table.</li> <li>• <b>Externals suppressed</b>—Number of routes received from EBGP peers currently inactive because of damping or other reasons.</li> <li>• <b>Received internal prefixes</b>—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table.</li> <li>• <b>Active internal prefixes</b>—Number of prefixes received from the IBGP peers that are active in the routing table.</li> <li>• <b>Internals suppressed</b>—Number of routes received from IBGP peers currently inactive because of damping or other reasons.</li> <li>• <b>RIB State</b>—Status of the graceful restart process for this routing table: <b>BGP restart is complete</b>, <b>BGP restart in progress</b>, <b>VPN restart in progress</b>, or <b>VPN restart is complete</b>.</li> </ul> | <b>detail</b>       |
| <b>Groups</b>            | Total number of groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels          |
| <b>Peers</b>             | Total number of peers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels          |
| <b>External</b>          | Total number of external peers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels          |
| <b>Internal</b>          | Total number of internal peers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels          |
| <b>Down peers</b>        | Total number of unavailable peers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels          |
| <b>Flaps</b>             | Total number of flaps that occurred.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | All levels          |
| <b>Table</b>             | Name of a routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>brief</b> , none |
| <b>Tot Paths</b>         | Total number of routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>brief</b> , none |
| <b>Act Paths</b>         | Number of active routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>brief</b> , none |
| <b>Suppressed</b>        | Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>brief</b> , none |



Table 320: show bgp group Output Fields (*continued*)

| Field Name   | Field Description                                                                                                                                                                                           | Level of Output |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| History      | Number of withdrawn routes stored locally to keep track of damping history.                                                                                                                                 | brief, none     |
| Damp State   | Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs.                                                                                 | brief, none     |
| Pending      | Routes being processed by the BGP import policy.                                                                                                                                                            | brief, none     |
| Group        | Group the peer belongs to in the BGP configuration.                                                                                                                                                         | detail          |
| Receive mask | Mask of the received target included in the advertised route.                                                                                                                                               | detail          |
| Entries      | Number of route entries received.                                                                                                                                                                           | detail          |
| Target       | Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer. | detail          |
| Mask         | Mask which specifies that the peer receive routes with the given route target.                                                                                                                              | detail          |

## Sample Output

### show bgp group

```

user@host> show bgp group
Groups: 2  Peers: 2  External: 0  Internal: 2  Down peers: 1  Flaps: 0
Table      Tot Paths  Act Paths  Suppressed  History  Damp State  Pending

inet.0
          0         0         0         0         0         0

bgp.13vpn.0
          0         0         0         0         0         0

bgp.rtarget.0
          2         0         0         0         0         0

```

### show bgp group brief

```

user@host> show bgp group brief
Groups: 2  Peers: 2  External: 0  Internal: 2  Down peers: 1  Flaps: 0
Table      Tot Paths  Act Paths  Suppressed  History  Damp State  Pending

inet.0
          0         0         0         0         0         0

bgp.13vpn.0
          0         0         0         0         0         0

bgp.rtarget.0
          2         0         0         0         0         0

```

## show bgp group detail

```

user@host> show bgp group detail
Group Type: Internal   AS: 1                      Local AS: 1
Name: ibgp             Index: 0                    Flags: <Export Eval>
Holdtime: 0
Total peers: 3         Established: 0
22.0.0.2
22.0.0.8
22.0.0.5

Groups: 1 Peers: 3   External: 0   Internal: 3   Down peers: 3   Flaps: 3
Table bgp.l3vpn.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table bgp.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.inet.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0

```

```

Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete

```

### show bgp group rtf detail

```

user@host> show bgp group rtf detail
Group: internal (group-index: 0)
  Receive mask: 00000002
  Table: bgp.rtarget.0
    Target
    100:100/64
    200:201/64
    Mask
    00000002
    (Group)
    Entries: 2
Group: internal (group-index: 1)
  Table: bgp.rtarget.0
    Target
    200:201/64
    Mask
    (Group)
    Entries: 1

```

### show bgp group summary

```

user@host> show bgp group summary
Group      Type      Peers      Established      Active/Received/Accepted/Damped
ibgp       Internal  3           0
Groups: 1  Peers: 3      External: 0      Internal: 3      Down peers: 3      Flaps: 3
  bgp.l3vpn.0      : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  bgp.mdt.0        : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  VPN-A.inet.0     : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  VPN-A.mdt.0      : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0

```

## show bgp neighbor

---

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3052</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3052</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax</b>                                   | <pre>show bgp neighbor &lt;exact-instance <i>instance-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>neighbor-address</i>&gt; &lt;orf (detail   <i>neighbor-address</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show bgp neighbor &lt;instance <i>instance-name</i>&gt; &lt;exact-instance <i>instance-name</i>&gt; &lt;<i>neighbor-address</i>&gt; &lt;orf (<i>neighbor-address</i>   detail)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>orf</b> option introduced in Junos OS Release 9.2.</p> <p><b>exact-instance</b> option introduced in Junos OS Release 11.4.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                              | Display information about BGP peers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                                  | <p><b>none</b>—Display information about all BGP peers.</p> <p><b>exact-instance <i>instance-name</i></b>—(Optional) Display information for the specified instance only.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, <b>cust1</b>, <b>cust11</b>, and <b>cust111</b> are all displayed when you run the <b>show bgp neighbor instance cust1</b> command).</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>neighbor-address</i></b>—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p><b>orf (detail   <i>neighbor-address</i>)</b>—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the <b>detail</b> option to display detailed output.</p> |
| <b>Additional Information</b>                   | For information about the <b>local-address</b> , <b>nlri</b> , <b>hold-time</b> , and <b>preference</b> statements, see the <i>Junos OS Routing Protocols Library for Routing Devices</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Related Documentation** • [clear bgp neighbor on page 3039](#)

**List of Sample Output** [show bgp neighbor on page 3059](#)  
[show bgp neighbor \(CLNS\) on page 3060](#)  
[show bgp neighbor \(Layer 2 VPN\) on page 3061](#)  
[show bgp neighbor \(Layer 3 VPN\) on page 3063](#)  
[show bgp neighbor neighbor-address on page 3063](#)  
[show bgp neighbor neighbor-address on page 3064](#)  
[show bgp neighbor orf neighbor-address detail on page 3065](#)

**Output Fields** [Table 321 on page 3053](#) describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

**Table 321: show bgp neighbor Output Fields**

| Field Name   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Peer</b>  | Address of the BGP neighbor. The address is followed by the neighbor port number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>AS</b>    | AS number of the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Local</b> | Address of the local routing device. The address is followed by the peer port number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Type</b>  | Type of peer: <b>Internal</b> or <b>External</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>State</b> | Current state of the BGP session: <ul style="list-style-type: none"> <li>• <b>Active</b>—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message.</li> <li>• <b>Connect</b>—BGP is waiting for the transport protocol connection to be completed.</li> <li>• <b>Established</b>—The BGP session has been established, and the peers are exchanging update messages.</li> <li>• <b>Idle</b>—This is the first stage of a connection. BGP is waiting for a Start event.</li> <li>• <b>OpenConfirm</b>—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.</li> <li>• <b>OpenSent</b>—BGP has sent an open message and is waiting to receive an open message from the peer.</li> </ul> |

Table 321: show bgp neighbor Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Flags</b>      | <p>Internal BGP flags:</p> <ul style="list-style-type: none"> <li>• <b>Aggregate Label</b>—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label.</li> <li>• <b>CleanUp</b>—The peer session is being shut down.</li> <li>• <b>Delete</b>—This peer has been deleted.</li> <li>• <b>Idled</b>—This peer has been permanently idled.</li> <li>• <b>ImportEval</b>—At the last commit operation, this peer was identified as needing to reevaluate all received routes.</li> <li>• <b>Initializing</b>—The peer session is initializing.</li> <li>• <b>SendRtn</b>—Messages are being sent to the peer.</li> <li>• <b>Sync</b>—This peer is synchronized with the rest of the peer group.</li> <li>• <b>TryConnect</b>—Another attempt is being made to connect to the peer.</li> <li>• <b>Unconfigured</b>—This peer is not configured.</li> <li>• <b>WriteFailed</b>—An attempt to write to this peer failed.</li> </ul>                                                                                                                                                                                                                                                                 |
| <b>Last state</b> | <p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message.</li> <li>• <b>Connect</b>—BGP is waiting for the transport protocol connection to be completed.</li> <li>• <b>Established</b>—The BGP session has been established, and the peers are exchanging update messages.</li> <li>• <b>Idle</b>—This is the first stage of a connection. BGP is waiting for a Start event.</li> <li>• <b>OpenConfirm</b>—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.</li> <li>• <b>OpenSent</b>—BGP has sent an open message and is waiting to receive an open message from the peer.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Last event</b> | <p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Closed</b>—The BGP session closed.</li> <li>• <b>ConnectRetry</b>—The transport protocol connection failed, and BGP is trying again to connect.</li> <li>• <b>HoldTime</b>—The session ended because the hold timer expired.</li> <li>• <b>KeepAlive</b>—The local routing device sent a BGP keepalive message to the peer.</li> <li>• <b>Open</b>—The local routing device sent a BGP open message to the peer.</li> <li>• <b>OpenFail</b>—The local routing device did not receive an acknowledgment of a BGP open message from the peer.</li> <li>• <b>RecvKeepAlive</b>—The local routing device received a BGP keepalive message from the peer.</li> <li>• <b>RecvNotify</b>—The local routing device received a BGP notification message from the peer.</li> <li>• <b>RecvOpen</b>—The local routing device received a BGP open message from the peer.</li> <li>• <b>RecvUpdate</b>—The local routing device received a BGP update message from the peer.</li> <li>• <b>Start</b>—The peering session started.</li> <li>• <b>Stop</b>—The peering session stopped.</li> <li>• <b>TransportError</b>—A TCP error occurred.</li> </ul> |

Table 321: show bgp neighbor Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Last error                  | <p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Cease</b>—An error occurred, such as a version mismatch, that caused the session to close.</li> <li>• <b>Finite State Machine Error</b>—In setting up the session, BGP received a message that it did not understand.</li> <li>• <b>Hold Time Expired</b>—The session's hold time expired.</li> <li>• <b>Message Header Error</b>—The header of a BGP message was malformed.</li> <li>• <b>Open Message Error</b>—A BGP open message contained an error.</li> <li>• <b>None</b>—No errors occurred in the BGP session.</li> <li>• <b>Update Message Error</b>—A BGP update message contained an error.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Export                      | Name of the export policy that is configured on the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Import                      | Name of the import policy that is configured on the peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Options                     | <p>Configured BGP options:</p> <ul style="list-style-type: none"> <li>• <b>AddressFamily</b>—Configured address family: <b>inet</b> or <b>inet-vpn</b>.</li> <li>• <b>AuthKeyChain</b>—Authentication key change is enabled.</li> <li>• <b>DropPathAttributes</b>—Certain path attributes are configured to be dropped from neighbor updates during inbound processing.</li> <li>• <b>GracefulRestart</b>—Graceful restart is configured.</li> <li>• <b>HoldTime</b>—Hold time configured with the <b>hold-time</b> statement. The hold time is three times the interval at which keepalive messages are sent.</li> <li>• <b>IgnorePathAttributes</b>—Certain path attributes are configured to be ignored in neighbor updates during inbound processing.</li> <li>• <b>Local Address</b>—Address configured with the <b>local-address</b> statement.</li> <li>• <b>Multihop</b>—Allow BGP connections to external peers that are not on a directly connected network.</li> <li>• <b>NLRI</b>—Configured MBGP state for the BGP group: <b>multicast</b>, <b>unicast</b>, or both if you have configured <b>nlri any</b>.</li> <li>• <b>Peer AS</b>—Configured peer autonomous system (AS).</li> <li>• <b>Preference</b>—Preference value configured with the <b>preference</b> statement.</li> <li>• <b>Refresh</b>—Configured to refresh automatically when the policy changes.</li> <li>• <b>Rib-group</b>—Configured routing table group.</li> </ul> |
| Path-attributes dropped     | Path attribute codes that are dropped from neighbor updates.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Path-attributes ignored     | Path attribute codes that are ignored during neighbor updates.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Authentication key change   | (appears only if the <b>authentication-keychain</b> statement has been configured) Name of the authentication keychain enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Authentication algorithm    | (appears only if the <b>authentication-algorithm</b> statement has been configured) Type of authentication algorithm enabled: <b>hmac</b> or <b>md5</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Address families configured | Names of configured address families for the VPN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 321: show bgp neighbor Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local Address                          | Address of the local routing device.                                                                                                                                                                                                                                                                                                                                            |
| Remove-private options                 | Options associated with the <code>remove-private</code> statement.                                                                                                                                                                                                                                                                                                              |
| Holdtime                               | Hold time configured with the <code>hold-time</code> statement. The hold time is three times the interval at which keepalive messages are sent.                                                                                                                                                                                                                                 |
| Flags for NLRI<br>inet-label-unicast   | Flags related to labeled-unicast: <ul style="list-style-type: none"> <li>• <b>TrafficStatistics</b>—Collection of statistics for labeled-unicast traffic is enabled.</li> </ul>                                                                                                                                                                                                 |
| Traffic statistics                     | Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> <li>• <b>Options</b>—Options configured for collecting statistics about labeled-unicast traffic.</li> <li>• <b>File</b>—Name and location of statistics log files.</li> <li>• <b>size</b>—Size of all the log files, in bytes.</li> <li>• <b>files</b>—Number of log files.</li> </ul> |
| Traffic Statistics<br>Interval         | Time between sample periods for labeled-unicast traffic statistics, in seconds.                                                                                                                                                                                                                                                                                                 |
| Preference                             | Preference value configured with the <code>preference</code> statement.                                                                                                                                                                                                                                                                                                         |
| Outbound Timer                         | Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the <code>out-delay</code> parameter is configured to a non-zero value.                                                                                                                                                             |
| Number of flaps                        | Number of times the BGP session has gone down and then come back up.                                                                                                                                                                                                                                                                                                            |
| Peer ID                                | Router identifier of the peer.                                                                                                                                                                                                                                                                                                                                                  |
| Group index                            | Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.                                                                                                                                                                                   |
| Peer index                             | Index that is unique within the BGP group to which the peer belongs.                                                                                                                                                                                                                                                                                                            |
| Local ID                               | Router identifier of the local routing device.                                                                                                                                                                                                                                                                                                                                  |
| Local Interface                        | Name of the interface on the local routing device.                                                                                                                                                                                                                                                                                                                              |
| Active holdtime                        | Hold time that the local routing device negotiated with the peer.                                                                                                                                                                                                                                                                                                               |
| Keepalive Interval                     | Keepalive interval, in seconds.                                                                                                                                                                                                                                                                                                                                                 |
| BFD                                    | Status of BFD failure detection.                                                                                                                                                                                                                                                                                                                                                |
| Local Address                          | Name of directly connected interface over which direct EBGP peering is established.                                                                                                                                                                                                                                                                                             |
| NLRI for restart<br>configured on peer | Names of address families configured for restart.                                                                                                                                                                                                                                                                                                                               |



Table 321: show bgp neighbor Output Fields (*continued*)

| Field Name                                      | Field Description                                                                                                                                                                                                                         |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NLRI advertised by peer                         | Address families supported by the peer: <b>unicast</b> or <b>multicast</b> .                                                                                                                                                              |
| NLRI for this session                           | Address families being used for this session.                                                                                                                                                                                             |
| Peer supports Refresh capability                | Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .                                                    |
| Restart time configured on peer                 | Configured time allowed for restart on the neighbor.                                                                                                                                                                                      |
| Stale routes from peer are kept for             | When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.                                                                                                          |
| Peer does not support Restarter functionality   | Graceful restart restarter-mode is disabled on the peer.                                                                                                                                                                                  |
| Peer does not support Receiver functionality    | Graceful restart helper-mode is disabled on the peer.                                                                                                                                                                                     |
| Restart time requested by this peer             | Restart time requested by this neighbor during capability negotiation.                                                                                                                                                                    |
| Restart flag received from the peer             | When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the <b>end-of-rib</b> marker from the speaker before advertising routing information to the speaker.                               |
| NLRI that peer supports restart for             | Neighbor supports graceful restart for this address family.                                                                                                                                                                               |
| NLRI peer can save forwarding state             | Neighbor supporting this address family saves all forwarding states.                                                                                                                                                                      |
| NLRI that peer saved forwarding for             | Neighbor saves all forwarding states for this address family.                                                                                                                                                                             |
| NLRI that restart is negotiated for             | Router supports graceful restart for this address family.                                                                                                                                                                                 |
| NLRI of received end-of-rib markers             | Address families for which end-of-routing-table markers are received from the neighbor.                                                                                                                                                   |
| NLRI of all end-of-rib markers sent             | Address families for which end-of-routing-table markers are sent to the neighbor.                                                                                                                                                         |
| Peer supports 4 byte AS extension (peer-as1)    | Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.                                                                                                                                    |
| NLRIs for which peer can receive multiple paths | Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route.<br><br>Possible value is <b>inet-unicast</b> . |

Table 321: show bgp neighbor Output Fields (*continued*)

| Field Name                                                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NLRIs for which peer can send multiple paths: inet-unicast | Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route.<br><br>Possible value is <b>inet-unicast</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Table inet.number                                          | Information about the routing table: <ul style="list-style-type: none"> <li>• <b>RIB State</b>—BGP is in the graceful restart process for this routing table: <b>restart is complete</b> or <b>restart in progress</b>.</li> <li>• <b>Bit</b>—Number that represents the entry in the routing table for this peer.</li> <li>• <b>Send state</b>—State of the BGP group: <b>in sync</b>, <b>not in sync</b>, or <b>not advertising</b>.</li> <li>• <b>Active prefixes</b>—Number of prefixes received from the peer that are active in the routing table.</li> <li>• <b>Received prefixes</b>—Total number of prefixes from the peer, both active and inactive, that are in the routing table.</li> <li>• <b>Accepted prefixes</b>—Total number of prefixes from the peer that have been accepted by a routing policy.</li> <li>• <b>Suppressed due to damping</b>—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.</li> </ul> |
| Last traffic (seconds)                                     | Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Input messages                                             | Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Output messages                                            | Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Input dropped path attributes                              | Information about dropped path attributes: <ul style="list-style-type: none"> <li>• <b>Code</b>—Path attribute code.</li> <li>• <b>Count</b>—Path attribute count.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Input ignored path attributes                              | Information about ignored path attributes: <ul style="list-style-type: none"> <li>• <b>Code</b>—Path attribute code.</li> <li>• <b>Count</b>—Path attribute count.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Output queue                                               | Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.<br><br>It also specifies the routing table name and the NLRI they represent in the format ( <b>routing table name, NLRI</b> ).<br><br><b>NOTE:</b> The output queues of routing tables that are not advertised, will only show up at <b>extensive</b> output level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Trace options                                              | Configured tracing of BGP protocol packets and operations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

Table 321: show bgp neighbor Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                   |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Trace file              | Name of the file to receive the output of the tracing operation.                                                                                                                                                                                                                                                                                    |
| Filter Updates rcv      | (orf option only) Number of outbound-route filters received for each configured address family.<br><br><b>NOTE:</b> The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.                                                                       |
| Immediate               | (orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes.<br><br><b>NOTE:</b> The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list. |
| Filter                  | (orf option only) Type of prefix filter received: <b>prefix-based</b> or <b>extended-community</b> .                                                                                                                                                                                                                                                |
| Received filter entries | (orf option only) List of received filters displayed.                                                                                                                                                                                                                                                                                               |
| seq                     | (orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.                                                                                                                                                                                                                        |
| prefix                  | (orf option only) Address for the prefix entry that matches the filter.                                                                                                                                                                                                                                                                             |
| minlength               | (orf option only) Minimum prefix length, in bits, required to match this prefix.                                                                                                                                                                                                                                                                    |
| maxlength               | (orf option only) Maximum prefix length, in bits, required to match this prefix.                                                                                                                                                                                                                                                                    |
| match                   | (orf option only) For this prefix match, whether to <b>permit</b> or <b>deny</b> route updates.                                                                                                                                                                                                                                                     |

## Sample Output

### show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250   Local ID: 10.255.7.248   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast

```

```

NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 10)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 9    Sent 5    Checked 5
Input messages: Total 36    Updates 2    Refreshes 0    Octets 718
Output messages: Total 37    Updates 1    Refreshes 0    Octets 796
Output Queue[0]: 0 (inet.0, inet-unicast)

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-vpn-unicast route-target
  Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.162.214 Local ID: 10.255.167.205 Active Holdtime: 90
  Keepalive Interval: 30 Group index: 0 Peer index: 1

```

### show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
  Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1 Local ID: 10.245.245.3 Active Holdtime: 90
  Keepalive Interval: 30 Peer index: 0
  NLRI advertised by peer: iso-vpn-unicast
  NLRI for this session: iso-vpn-unicast
  Peer supports Refresh capability (2)
  Table bgp.isovpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          3
    Received prefixes:        3
    Suppressed due to damping: 0
    Advertised prefixes:      3
  Table aaaa.iso.0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: not advertising
    Active prefixes:          3
    Received prefixes:        3
    Suppressed due to damping: 0
  Last traffic (seconds): Received 6    Sent 5    Checked 5
  Input messages: Total 1736    Updates 4    Refreshes 0    Octets 33385
  Output messages: Total 1738    Updates 3    Refreshes 0    Octets 33305

```

```
Output Queue[0]: 0 (bgp.isovpn.0, iso-vpn-unicast)
Output Queue[1]: 0 (aaaa.iso.0, iso-vpn-unicast)
```

### show bgp neighbor (Layer 2 VPN)

```
user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65100 Local: 10.69.103.1      AS 65103
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle     Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.69.104.2      AS 65100 Local: 10.69.104.1      AS 65104
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle     Last Event: Start
  Last Error: None
  Export: [ BGP-L-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-labeled-unicast
  Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.255.14.182+179 AS 69      Local: 10.255.14.176+2131 AS 69
  Type: Internal      State: Established  Flags: <ImportEval>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer ID: 10.255.14.182      Local ID: 10.255.14.176      Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
```

```

Received prefixes:          1
Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            2
Received prefixes:          2
Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            2
Received prefixes:          2
Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            1
Received prefixes:          1
Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            2
Received prefixes:          2
Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            2
Received prefixes:          2
Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            1
Received prefixes:          1
Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            1
Received prefixes:          1
Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
Output Queue[1]: 0 (bgp.l2vpn.0, inet-vpn-unicast)
Output Queue[2]: 0 (BGP-INET.inet.0, inet-vpn-unicast)
Output Queue[3]: 0 (BGP-L.inet.0, inet-vpn-unicast)
Output Queue[4]: 0 (LDP.inet.0, inet-vpn-unicast)
Output Queue[5]: 0 (OSPF.inet.0, inet-vpn-unicast)
Output Queue[6]: 0 (RIP.inet.0, inet-vpn-unicast)

```

```
Output Queue[7]: 0 (STATIC.inet.0, inet-vpn-unicast)
Output Queue[8]: 0 (L2VPN.l2vpn.0, inet-vpn-unicast)
```

### show bgp neighbor (Layer 3 VPN)

```
user@host> show bgp neighbor
Peer: 4.4.4.4+179      AS 10045 Local: 5.5.5.5+1214      AS 10045
Type: Internal      State: Established      Flags: <ImportEval>
Last State: OpenConfirm      Last Event: RecvKeepAlive
Last Error: None
Export: [ match-all ] Import: [ match-all ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
      Rib-group Refresh>
Address families configured: inet-vpn-unicast
Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
Flags for NLRI inet-labeled-unicast: TrafficStatistics
Traffic Statistics: Options: all File: /var/log/bstat.log
                        size 131072 files 10

Traffic Statistics Interval: 60
Number of flaps: 0
Peer ID: 192.168.1.110      Local ID: 192.168.1.111      Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast
NLRI advertised by peer: inet-vpn-unicast
NLRI for this session: inet-vpn-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast
NLRI peer can save forwarding state: inet-vpn-unicast
NLRI that peer saved forwarding for: inet-vpn-unicast
NLRI that restart is negotiated for: inet-vpn-unicast
NLRI of received end-of-rib markers: inet-vpn-unicast
NLRI of all end-of-rib markers sent: inet-vpn-unicast
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table vpn-green.inet.0 Bit: 20001
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Last traffic (seconds): Received 15      Sent 20      Checked 20
Input messages: Total 40      Updates 2      Refreshes 0      Octets 856
Output messages: Total 44      Updates 2      Refreshes 0      Octets 1066
Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
Output Queue[1]: 0 (vpn-green.inet.0, inet-vpn-unicast)
Trace options: detail packets
Trace file: /var/log/bgpgr.log size 131072 files 10
```

### show bgp neighbor neighbor-address

```
user@host> show bgp neighbor 192.168.1.111
```

```

Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
  Refresh>
  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90
  Keepalive Interval: 30
BFD: disabled
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 4
    Received prefixes: 6
    Suppressed due to damping: 0
  Table inet6.0 Bit: 20000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 3 Sent 3 Checked 3
  Input messages: Total 9 Updates 6 Refreshes 0 Octets 403
  Output messages: Total 7 Updates 3 Refreshes 0 Octets 365
  Output Queue[0]: 0 (inet.0, inet-unicast)
  Output Queue[1]: 0 (inet6.0, inet6-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpr size 131072 files 10

```

### show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
  Type: External State: Established Flags: <Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: Cease
  Export: [ export-policy ] Import: [ import-policy ]
  Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
  Address families configured: inet-unicast inet-multicast
  Holdtime: 60000 Preference: 170
  Number of flaps: 4
  Last flap event: RecvUpdate
  Error: 'Cease' Sent: 5 Recv: 0
  Peer ID: 10.255.245.6 Local ID: 10.255.245.5 Active Holdtime: 60000
  Keepalive Interval: 20000 Peer index: 0
  BFD: disabled, down

```



```

Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:           8
  Received prefixes:        10
  Accepted prefixes:        10
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:           0
  Received prefixes:         0
  Accepted prefixes:         0
  Suppressed due to damping: 0
  Advertised prefixes:       0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0 (inet.0, inet-unicast)
Output Queue[1]: 0 (inet.2, inet-multicast)
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

#### show bgp neighbor orf neighbor-address detail

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
  Filter updates rcv:           1 Immediate:           1
  Filter: prefix-based receive
  Received filter entries:
    seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast
  Filter updates rcv:           0 Immediate:           1
  Filter: prefix-based receive
  Received filter entries:
    *.*

```

## show bgp summary

---

|                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                           | <a href="#">Syntax on page 3066</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3066</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Syntax                                   | <pre>show bgp summary &lt;exact-instance <i>instance-name</i>&gt; &lt;group <i>group-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Syntax (EX Series Switch and QFX Series) | <pre>show bgp summary &lt;exact-instance <i>instance-name</i>&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Release Information                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.<br><b>exact-instance</b> option introduced in Junos OS Release 11.4.<br><b>group</b> option introduced in Junos OS Release 13.3.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Description                              | Display BGP summary information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Options                                  | <b>none</b> —Display BGP summary information for all routing instances.<br><br><b>exact-instance <i>instance-name</i></b> —(Optional) Display information for the specified instance only.<br><br><b>group</b> —Display overview of bgp information for a particular group<br><br><b>instance <i>instance-name</i></b> —(Optional) Display information for all routing instances whose name begins with this string (for example, <b>cust1</b> , <b>cust11</b> , and <b>cust111</b> are all displayed when you run the <b>show bgp summary instance cust1</b> command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| List of Sample Output                    | <a href="#">show bgp summary (When a Peer Is Not Established) on page 3069</a><br><a href="#">show bgp summary (When a Peer Is Established) on page 3069</a><br><a href="#">show bgp summary (CLNS) on page 3069</a><br><a href="#">show bgp summary (Layer 2 VPN) on page 3070</a><br><a href="#">show bgp summary (Layer 3 VPN) on page 3070</a><br><a href="#">show bgp summary group on page 3070</a>                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Output Fields                            | <a href="#">Table 322 on page 3067</a> describes the output fields for the <b>show bgp summary</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 322: show bgp summary Output Fields

| Field Name          | Field Description                                                                                                                                                      |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Groups</b>       | Number of BGP groups.                                                                                                                                                  |
| <b>Peers</b>        | Number of BGP peers.                                                                                                                                                   |
| <b>Down peers</b>   | Number of down BGP peers.                                                                                                                                              |
| <b>Table</b>        | Name of routing table.                                                                                                                                                 |
| <b>Tot Paths</b>    | Total number of paths.                                                                                                                                                 |
| <b>Act Paths</b>    | Number of active routes.                                                                                                                                               |
| <b>Suppressed</b>   | Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. |
| <b>History</b>      | Number of withdrawn routes stored locally to keep track of damping history.                                                                                            |
| <b>Damp State</b>   | Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs.               |
| <b>Pending</b>      | Routes in process by BGP import policy.                                                                                                                                |
| <b>Peer</b>         | Address of each BGP peer. Each peer has one line of output.                                                                                                            |
| <b>AS</b>           | Peer's AS number.                                                                                                                                                      |
| <b>InPkt</b>        | Number of packets received from the peer.                                                                                                                              |
| <b>OutPkt</b>       | Number of packets sent to the peer.                                                                                                                                    |
| <b>OutQ</b>         | Number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.                       |
| <b>Flaps</b>        | Number of times the BGP session has gone down and then come back up.                                                                                                   |
| <b>Last Up/Down</b> | Last time since the neighbor transitioned to or from the established state.                                                                                            |

Table 322: show bgp summary Output Fields (*continued*)

| Field Name                                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>State #Active<br/>/Received/Accepted<br/>/Damped</b> | <p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established on the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> <li>If a peer is not established, the field shows the state of the peer session: <b>Active</b>, <b>Connect</b>, or <b>Idle</b>. In general, the Idle state is the first stage of a connection. BGP is waiting for a Start event. A session can be idle for other reasons as well. The reason that a session is idle is sometimes displayed. For example: <b>Idle (Removal in progress)</b> or <b>Idle (LicenseFailure)</b>.</li> <li>If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the <b>inet.0</b> (main) and <b>inet.2</b> (multicast) routing tables. For example, <b>8/10/10/2</b> and <b>2/4/4/0</b> indicate the following: <ul style="list-style-type: none"> <li>8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the <b>inet.0</b> routing table.</li> <li>2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the <b>inet.2</b> routing table.</li> </ul> </li> <li>If a BGP session is established in a routing instance, the field indicates the established (<b>Establ</b>) state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, <b>Establ VPN-AB.inet.0: 2/4/0</b> indicates the following: <ul style="list-style-type: none"> <li>The BGP session is established.</li> <li>Routes are received in the <b>VPN-AB.inet.0</b> routing table.</li> <li>The local routing device has two active routes, four received routes, and no damped routes from a BGP peer.</li> </ul> </li> </ul> <p>When a BGP session is established, the peers are exchanging update messages.</p> |

## Sample Output

### show bgp summary (When a Peer Is Not Established)

```

user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0          6          4          0          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.3        65002      86       90       0        2      42:54 0/0/0

0/0/0
10.0.0.4        65002      90       91       0        1      42:54 0/2/0

0/0/0
10.0.0.6        65002      87       90       0        3          3 Active
10.1.12.1       65001      89       89       0        1      42:54 4/4/0

0/0/0

```

### show bgp summary (When a Peer Is Established)

```

user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0          6          4          0          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2        65002    88675    88652      0        2      42:38 2/4/0

0/0/0
10.0.0.3        65002    54528    54532      0        1     2w4d22h 0/0/0

0/0/0
10.0.0.4        65002    51597    51584      0        0     2w3d22h 2/2/0

0/0/0

user@host> show bgp summary logical-system R3
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
bgp.13vpn.0      2          2          0          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
1.1.1.2          2       204     206       0        0      1:30:59
Establ
  bgp.13vpn.0: 2/2/2/0
  red.inet.0: 2/2/2/0
10.1.1.10        3       206     207       0        0      1:31:36
Establ
  red.inet.0: 2/2/2/0

```

### show bgp summary (CLNS)

```

user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.245.245.1    200     1735     1737      0        0     14:26:12 Establ

```

```

bgp.isovpn.0: 3/3/0
aaaa.iso.0: 3/3/0

```

### show bgp summary (Layer 2 VPN)

```

user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l2vpn.0 1 1 0 0 0 0
inet.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.255.245.35 65299 72 74 0 1 19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0
10.255.245.36 65299 2164 2423 0 4 19:50 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.37 65299 36 37 0 4 17:07 Establ
  inet.0: 0/0/0
10.255.245.39 65299 138 168 0 6 53:48 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.69 65299 134 140 0 6 53:42 Establ
  inet.0: 0/0/0

```

### show bgp summary (Layer 3 VPN)

```

user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5 2 21 22 0 0 6:26 Establ
  VPN-AB.inet.0: 1/1/0
10.255.71.15 1 19 21 0 0 6:17 Establ
  bgp.l3vpn.0: 2/2/0
  VPN-A.inet.0: 1/1/0
  VPN-AB.inet.0: 2/2/0
  VPN-B.inet.0: 1/1/0

```

### show bgp summary group

```

user@host> show bgp summary group Group2
Groups: 3 Peers: 3 Down peers: 3
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.1 56 0 0 0 0 51
Idle

user@host> show bgp summary logical-system R3 group toR4
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.1.1.10 3 207 207 0 0 1:31:40

```

```
Estab1  
red.inet.0: 2/2/2/0
```

## Sample Output

## show policy damping

---

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3072</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3072</a>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax</b>                                   | show policy damping<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax (EX Series Switch and QFX Series)</b> | show policy damping                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                             |
| <b>Description</b>                              | Display information about BGP route flap damping parameters.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                  | <b>none</b> —Display information about BGP route flap damping parameters.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.                                                                                                                                                                                                                               |
| <b>Additional Information</b>                   | In the output from this command, figure-of-merit values correlate with the probability of future instability of a routing device. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time. A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes. |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"><li>• “Configuring BGP Flap Damping Parameters” in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li><li>• <a href="#">clear bgp damping on page 3038</a></li><li>• <a href="#">show route damping on page 3577</a></li></ul>                                                                                                                                              |
| <b>List of Sample Output</b>                    | <a href="#">show policy damping on page 3073</a>                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>                            | <a href="#">Table 323 on page 3073</a> describes the output fields for the <b>show policy damping</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                               |



Table 323: show policy damping Output Fields

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Halflife</b>              | Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes. |
| <b>Reuse merit</b>           | Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.                                                                                                                                                  |
| <b>Suppress/cutoff merit</b> | Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols.                                                           |
| <b>Maximum suppress time</b> | Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability.                                                                                                                                                                                                                                                                                           |
| <b>Computed values</b>       | <ul style="list-style-type: none"> <li>• <b>Merit ceiling</b>—Maximum merit that a flapping route can collect.</li> <li>• <b>Maximum decay</b>—Maximum decay half-life, in minutes.</li> </ul>                                                                                                                                                                                                                                                                         |

## Sample Output

### show policy damping

```

user@host> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "standard-damping":
  Halflife: 10 minutes
  Reuse merit: 4000 Suppress/cutoff merit: 8000
  Maximum suppress time: 30 minutes
  Computed values:
    Merit ceiling: 32120
    Maximum decay: 12453

```



## PART 16

# IS-IS

- [Overview on page 3077](#)
- [Configuration on page 3081](#)
- [Administration on page 3141](#)



## CHAPTER 51

# Overview

- [Layer 3 Protocols on page 3077](#)

## Layer 3 Protocols

---

- [Layer 3 Protocols Supported on EX Series Switches on page 3077](#)
- [Layer 3 Protocols Not Supported on EX Series Switches on page 3078](#)

## Layer 3 Protocols Supported on EX Series Switches

EX Series switches support the Junos OS Layer 3 features and configuration statements listed in [Table 314 on page 2939](#):

**Table 324: Supported Junos OS Layer 3 Protocol Statements and Features**

| Protocol           | Notes                                                                                                                                            | For More Information                                             |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| BGP                | Fully supported.                                                                                                                                 | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| BFD                | Fully supported.                                                                                                                                 | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| ICMP               | Fully supported.                                                                                                                                 | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| IGMPv1, v2, and v3 | Fully supported.                                                                                                                                 | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |
| IS-IS              | Supported, with the exceptions noted in " <a href="#">Layer 3 Protocols Not Supported on EX Series Switches</a> " on <a href="#">page 2940</a> . | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| MLD                | Fully supported (MLD versions 1 and 2).                                                                                                          | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |
| MPLS               | Supported, with the exceptions noted in " <a href="#">Layer 3 Protocols Not Supported on EX Series Switches</a> " on <a href="#">page 2940</a> . | <a href="#">Junos OS MPLS Applications Configuration Guide</a>   |
| OSPFv1, v2 and v3  | Supported, with the exceptions noted in " <a href="#">Layer 3 Protocols Not Supported on EX Series Switches</a> " on <a href="#">page 2940</a> . | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| PIM                | Fully supported.                                                                                                                                 | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |

Table 324: Supported Junos OS Layer 3 Protocol Statements and Features (*continued*)

| Protocol | Notes                                                                                                | For More Information                                                                                                     |
|----------|------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| PPM      | Supported. See <i>EX Series Switch Software Features Overview</i> for specific platform information. | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| RIP      | Fully supported.                                                                                     | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| RIPng    | Fully supported.                                                                                     | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| SNMP     | Fully supported.                                                                                     | <a href="#">Junos OS Network Management Configuration Guide</a>                                                          |
| VRRP     | Fully supported.                                                                                     | See “Understanding VRRP on EX Series Switches” on page 2501. See also <a href="#">Junos OS High Availability Guide</a> . |

- Related Documentation**
- [Layer 3 Protocols Not Supported on EX Series Switches](#) on page 2940
  - [EX Series Switch Software Features Overview](#)

## Layer 3 Protocols Not Supported on EX Series Switches

EX Series switches do not support the Junos OS Layer 3 protocols and features listed in [Table 315 on page 2940](#):

Table 325: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported

| Feature                                                                                                         | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DVMRP                                                                                                           | <ul style="list-style-type: none"> <li>• <b>dvmrp</b> and subordinate statements</li> </ul>                                                                                                                                                                                                          |
| Flow aggregation (cflowd)                                                                                       | <ul style="list-style-type: none"> <li>• <b>cflow</b> and subordinate statements</li> </ul>                                                                                                                                                                                                          |
| IPsec                                                                                                           | <ul style="list-style-type: none"> <li>• <b>[edit services]</b> statements related to IPsec</li> </ul>                                                                                                                                                                                               |
| IS-IS: <ul style="list-style-type: none"> <li>• ES-IS</li> <li>• IPv6 in multicast routing protocols</li> </ul> | <ul style="list-style-type: none"> <li>• <b>clns-routing</b> statement</li> <li>• <b>ipv6-multicast</b> statement</li> <li>• <b>lsp-interval</b> statement</li> <li>• <b>label-switched-path</b> statement</li> <li>• <b>lsp-lifetime</b> statement</li> <li>• <b>te-metric</b> statement</li> </ul> |
| Logical routers                                                                                                 | <ul style="list-style-type: none"> <li>• <b>logical-routers</b> and subordinate statements</li> </ul>                                                                                                                                                                                                |

Table 325: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)

| Feature                                                                                                                                                                                                                                                                                                                                                                                                                                      | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MPLS: <ul style="list-style-type: none"> <li>Fast Reroute (FRR)</li> <li>Label Distribution Protocol (LDP) (except on EX8200 switches)</li> <li>Layer 3 VPNs (except on EX8200 switches)</li> <li>Multiprotocol BGP (MP-BGP) for VPN-IPv4 family</li> <li>Pseudowire emulation (PWE3)</li> <li>Routing policy statements related to Layer 3 VPNs and MPLS (except on EX8200 switches)</li> <li>Virtual Private LAN Service (VPLS)</li> </ul> | <ul style="list-style-type: none"> <li><b>ldp</b> and all subordinate statements (except on EX8200 switches)</li> </ul>                                                                                                                                                                                                                                                                                           |
| Network Address Translation (NAT)                                                                                                                                                                                                                                                                                                                                                                                                            | <ul style="list-style-type: none"> <li><b>nat</b> and subordinate statements</li> <li>Policy statements related to NAT</li> </ul>                                                                                                                                                                                                                                                                                 |
| OSPF                                                                                                                                                                                                                                                                                                                                                                                                                                         | <ul style="list-style-type: none"> <li><b>demand-circuit</b> statement</li> <li><b>label-switched-path</b> and subordinate statements</li> <li><b>neighbor</b> statement within an OSPF area</li> <li><b>peer-interface</b> and subordinate statements within an OSPF area</li> <li><b>sham-link</b> statement</li> <li><b>te-metric</b> statement</li> </ul>                                                     |
| PPM                                                                                                                                                                                                                                                                                                                                                                                                                                          | <ul style="list-style-type: none"> <li>Not supported on EX2200 and EX3300 switches</li> </ul>                                                                                                                                                                                                                                                                                                                     |
| Routing instances: <ul style="list-style-type: none"> <li>Routing instance forwarding</li> </ul>                                                                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li><b>l2vpn</b> and subordinate statements (except on EX4500, EX4550, and EX8200 switches)</li> <li><b>ldp</b> and subordinate statements (except on EX8200 switches)</li> <li><b>vpls</b> and subordinate statements</li> </ul>                                                                                                                                              |
| Routed VLAN interfaces (RVIs)                                                                                                                                                                                                                                                                                                                                                                                                                | <ul style="list-style-type: none"> <li><b>family mpls</b> statement</li> </ul>                                                                                                                                                                                                                                                                                                                                    |
| SAP and SDP                                                                                                                                                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li><b>sap</b> and all subordinate statements</li> </ul>                                                                                                                                                                                                                                                                                                                       |
| General routing options in the <b>routing-options</b> hierarchy: <ul style="list-style-type: none"> <li>MPLS and label-switched-paths</li> </ul>                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li><b>auto-export</b> and subordinate statements</li> <li><b>dynamic-tunnels</b> and subordinate statements</li> <li><b>lsp-next-hop</b> and subordinate statements</li> <li><b>multicast</b> and subordinate statements</li> <li><b>p2mp-lsp-next-hop</b> and subordinate statements</li> <li><b>route-distinguisher-id</b> statement (except on EX8200 switches)</li> </ul> |

Table 325: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)

| Feature                                                                    | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Traffic sampling and forwarding in the <b>forwarding-options</b> hierarchy | <ul style="list-style-type: none"> <li>• <b>accounting</b> and subordinate statements</li> <li>• <b>family mpls</b> and <b>family multiservice</b> under <b>hash-key</b> hierarchy</li> <li>• Under <b>monitoring group-name</b> family inet output hierarchy: <ul style="list-style-type: none"> <li>• <b>cflowd</b> statement</li> <li>• <b>export-format-cflowd-version-5</b> statement</li> <li>• <b>flow-active-timeout</b> statement</li> <li>• <b>flow-export-destination</b> statement</li> <li>• <b>flow-inactive-timeout</b> statement</li> <li>• <b>interface</b> statement</li> </ul> </li> <li>• <b>port-mirroring</b> statement (On EX Series switches, port mirroring is implemented using the <b>analyzer</b> statement.)</li> <li>• <b>sampling</b> and subordinate statements</li> </ul> |

- Related Documentation**
- [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)
  - [EX Series Switch Software Features Overview](#)



## CHAPTER 52

# Configuration

- [Configuration Statements on page 3081](#)

### Configuration Statements

---

- [authentication-key \(Protocols IS-IS\) on page 3083](#)
- [authentication-type \(Protocols IS-IS\) on page 3084](#)
- [bfd-liveness-detection \(Protocols IS-IS\) on page 3085](#)
- [checksum \(Protocols IS-IS\) on page 3087](#)
- [csnp-interval on page 3088](#)
- [disable \(Protocols IS-IS\) on page 3089](#)
- [export \(Protocols IS-IS\) on page 3090](#)
- [external-preference \(Protocols IS-IS\) on page 3091](#)
- [graceful-restart \(Protocols IS-IS\) on page 3092](#)
- [hello-authentication-key on page 3093](#)
- [hello-authentication-type on page 3094](#)
- [hello-interval \(Protocols IS-IS\) on page 3095](#)
- [hello-padding on page 3096](#)
- [hold-time \(Protocols IS-IS\) on page 3098](#)
- [ignore-attached-bit on page 3099](#)
- [interface \(Protocols IS-IS\) on page 3100](#)
- [ipv4-multicast on page 3102](#)
- [ipv4-multicast-metric on page 3103](#)
- [ipv6-multicast on page 3103](#)
- [ipv6-multicast-metric on page 3104](#)
- [ipv6-unicast on page 3105](#)
- [ipv6-unicast-metric on page 3106](#)
- [isis on page 3107](#)
- [level \(Global IS-IS\) on page 3108](#)
- [link-protection \(Protocols IS-IS\) on page 3109](#)

- [loose-authentication-check](#) on page 3109
- [lsp-interval](#) on page 3110
- [lsp-lifetime](#) on page 3111
- [max-areas](#) on page 3112
- [mesh-group \(Protocols IS-IS\)](#) on page 3113
- [metric \(Protocols IS-IS\)](#) on page 3114
- [no-adjacency-holddown](#) on page 3115
- [no-authentication-check](#) on page 3116
- [no-csnp-authentication](#) on page 3116
- [no-eligible-backup \(Protocols IS-IS\)](#) on page 3117
- [no-hello-authentication](#) on page 3117
- [no-ipv4-multicast](#) on page 3118
- [no-ipv4-routing](#) on page 3119
- [no-ipv6-multicast](#) on page 3120
- [no-ipv6-routing](#) on page 3121
- [no-ipv6-unicast](#) on page 3122
- [no-psnp-authentication](#) on page 3122
- [no-unicast-topology](#) on page 3123
- [node-link-protection \(Protocols IS-IS\)](#) on page 3123
- [overload \(Protocols IS-IS\)](#) on page 3124
- [passive \(Protocols IS-IS\)](#) on page 3127
- [point-to-point](#) on page 3128
- [preference \(Protocols IS-IS\)](#) on page 3129
- [prefix-export-limit \(Protocols IS-IS\)](#) on page 3130
- [priority \(Protocols IS-IS\)](#) on page 3131
- [reference-bandwidth \(Protocols IS-IS\)](#) on page 3132
- [rib-group \(Protocols IS-IS\)](#) on page 3133
- [spf-options \(Protocols IS-IS\)](#) on page 3134
- [topologies \(Protocols IS-IS\)](#) on page 3135
- [traceoptions \(Protocols IS-IS\)](#) on page 3136
- [wide-metrics-only](#) on page 3139

## authentication-key (Protocols IS-IS)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | authentication-key <i>key</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols isis <b>level</b> <i>level-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <b>level</b> <i>level-number</i> ],<br>[edit protocols isis <b>level</b> <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>level</b> <i>level-number</i> ]                                                                                                                   |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>         | <p>Authentication key (password). Neighboring routing devices use the password to verify the authenticity of packets sent from this interface. For the key to work, you also must include the <b>authentication-type</b> statement.</p> <p>All routing devices must use the same password. If you are using the Junos OS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces adjacent to the Juniper Networks routing device.</p> |
| <b>Default</b>             | If you do not include this statement and the <b>authentication-type</b> statement, IS-IS authentication is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>             | <b>key</b> —Authentication password. The password can be up to 1024 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").                                                                                                                                                                                                                                                                                                                                               |



**CAUTION:** A simple password for authentication is truncated if it exceeds 254 characters.

|                                 |                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Example: Configuring Hitless Authentication Key Rollover for IS-IS</li> </ul> |

## authentication-type (Protocols IS-IS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-type <i>authentication</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ]                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Enable authentication and specify the authentication scheme for IS-IS. If you enable authentication, you must specify a password by including the <b>authentication-key</b> statement.                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | If you do not include this statement and the <b>authentication-key</b> statement, IS-IS authentication is disabled.                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b><i>authentication</i></b> —Authentication scheme: <ul style="list-style-type: none"><li>• <b>md5</b>—Use HMAC authentication in combination with MD5. HMAC-MD5 authentication is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>.</li><li>• <b>simple</b>—Use a simple password for authentication. The password is included in the transmitted packet, making this method of authentication relatively insecure. We recommend that you <i>not</i> use this authentication method.</li></ul> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS</a></li><li>• <a href="#">authentication-key on page 3083</a></li><li>• <a href="#">no-authentication-check on page 3116</a></li></ul>                                                                                                                                                                                                                                                              |

## bfd-liveness-detection (Protocols IS-IS)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> bfd-liveness-detection {     authentication {         algorithm <i>algorithm-name</i>;         key-chain <i>key-chain-name</i>;         loose-check;     }     detection-time {         threshold <i>milliseconds</i>;     }     minimum-interval <i>milliseconds</i>;     minimum-receive-interval <i>milliseconds</i>;     multiplier <i>number</i>;     no-adaptation;     transmit-interval {         minimum-interval <i>milliseconds</i>;         threshold <i>milliseconds</i>;     }     version (1   automatic); } </pre>                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/> isis <a href="#">interface interface-name</a>],<br/> [edit protocols isis <a href="#">interface interface-name</a>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a>]</p>                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.<br/> Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/> <b>detection-time threshold</b> and <b>transmit-interval threshold</b> options added in Junos OS Release 8.2.<br/> Support for logical systems introduced in Junos OS Release 8.3.<br/> <b>no-adaptation</b> statement introduced in Junos OS Release 9.0.<br/> <b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> options introduced in Junos OS Release 9.6.<br/> Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                     |
| <b>Description</b>         | Configure bidirectional failure detection timers and authentication.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>             | <p><b>authentication algorithm <i>algorithm-name</i></b>—Configure the algorithm used to authenticate the specified BFD session: <b>simple-password</b>, <b>keyed-md5</b>, <b>keyed-sha-1</b>, <b>meticulous-keyed-md5</b>, <b>meticulous-keyed-sha-1</b>.</p> <p><b>authentication key-chain <i>key-chain-name</i></b>—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the <b>authentication-key-chains key-chain</b> statement at the <b>[edit security]</b> hierarchy level.</p> <p><b>authentication loose-check</b>—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.</p> |

**detection-time threshold *milliseconds***—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**minimum-interval *milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.

**Range:** 1 through 255,000

**minimum-receive-interval *milliseconds***—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**multiplier *number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation**—Specify that BFD sessions not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval *milliseconds***—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version)

**Default:** automatic

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                                                                                |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Example: Configuring BFD for IS-IS</i></li><li>• <i>Example: Configuring BFD Authentication for IS-IS</i></li></ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|

## checksum (Protocols IS-IS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | checksum;                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a>],</p> <p>[edit protocols isis <a href="#">interface interface-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Enable checksums for packets on this interface.</p> <p>Junos OS supports IS-IS checksums as documented in RFC 3358, <i>Optional Checksums in Intermediate System to Intermediate System (ISIS)</i>.</p> <p>The checksum cannot be enabled with MD5 hello authentication on the same interface.</p>                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Enabling Packet Checksums on IS-IS Interfaces</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                              |

## csnp-interval

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | csnp-interval ( <i>seconds</i>   disable);                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <b>interface</b> <i>interface-name</i> ],<br>[edit protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i> ]                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Configure the interval between complete sequence number PDUs (CSNPs) on a LAN interface.</p> <p>If the routing device is the designated router on a LAN, IS-IS sends CSN packets every 10 seconds. If the routing device is on a point-to-point interface, it sends CSN packets every 5 seconds multiplied by the number of IS-IS adjacencies over point-to-point links, which are in UP state.</p> <p>To configure the interface not to send any CSNPs, specify the <b>disable</b> option.</p> |
| <b>Default</b>                  | By default, IS-IS sends CSNPs periodically. If the routing device is the designated router on a LAN, IS-IS sends CSNPs every 10 seconds. If the routing device is on a point-to-point interface, it sends CSNPs every 5 seconds multiplied by the number of IS-IS adjacencies over point-to-point links, which are in UP state.                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>disable</b>—Do not send CSNPs on this interface.</p> <p><b>seconds</b>—Number of seconds between the sending of CSNPs.</p> <p><b>Range:</b> 1 through 65,535 seconds</p> <p><b>Default:</b> 10 seconds on LAN broadcast links. 5 seconds on point-to-point links.</p>                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces</i></li></ul>                                                                                                                                                                                                                                                                                                                                                      |




## disable (Protocols IS-IS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>isis</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>isis</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering],</p> <p>[edit protocols <b>isis</b>],</p> <p>[edit protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis traffic-engineering],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>isis</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Disable IS-IS on the routing device, on an interface, or on a level.</p> <p>At the <b>[edit protocols isis traffic-engineering]</b> hierarchy level, disable IS-IS support for traffic engineering.</p> <p>Enabling IS-IS on an interface (by including the <b>interface</b> statement at the <b>[edit protocols isis]</b> or the <b>[edit routing-instances routing-instance-name protocols isis]</b> hierarchy level), disabling it (by including the <b>disable</b> statement), and not actually having IS-IS run on an interface (by including the <b>passive</b> statement) are mutually exclusive states.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>                  | <p>IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which <b>family iso</b> is enabled.</p> <p>IS-IS support for traffic engineering is enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Multi-Level IS-IS</i></li> <li>• <i>IS-IS Overview</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## export (Protocols IS-IS)

---

|                                                                                                                                                                                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                  | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                         | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                     | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                                                                                                                             | <p>Apply one or more policies to routes being exported from the routing table into IS-IS.</p> <p>All routing protocols store the routes that they learn in the routing table. The routing table uses this collected route information to determine the active routes to destinations. The routing table then installs the active routes into its forwarding table and exports them into the routing protocols. It is these exported routes that the protocols advertise.</p> <p>For each protocol, you control which routes the protocol stores in the routing table and which routes the routing table exports into the protocol from the routing table by defining a <i>routing policy</i> for that protocol.</p> |
| <div> <b>NOTE:</b> For IS-IS, you cannot apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol can easily lead to an inconsistent topology database.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                                                                                                                                                                                                                                                                                                 | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                   | <ul style="list-style-type: none"><li>• <i>Example: Redistributing OSPF Routes into IS-IS</i></li><li>• <i>Example: Configuring an IS-IS Default Route Policy on Logical Systems</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## external-preference (Protocols IS-IS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>external-preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>],</p> <p>[edit protocols isis <a href="#">level level-number</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure the preference of external routes.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><i>preference</i>—Preference value.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)</p>                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Route Preferences Overview</i></li> <li>• <i>Example: Redistributing OSPF Routes into IS-IS</i></li> <li>• <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i></li> <li>• <a href="#">preference on page 3129</a></li> </ul>                                                                                                                                                |

## graceful-restart (Protocols IS-IS)

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>graceful-restart {<br/>  disable;<br/>  helper-disable;<br/>  restart-duration <i>seconds</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <i>isis</i> ],<br>[edit protocols <i>isis</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Configure graceful restart parameters for IS-IS.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the <b>[edit routing-options]</b> hierarchy level. When graceful restart is enabled, the restarting routing device is not removed from the network topology during the restart period. The adjacencies are reestablished after restart is complete.</p> <p>On LAN interfaces where IS-IS is configured on a transit router that serves as the designated router (DR), a graceful restart causes:</p> <ul style="list-style-type: none"><li>• The ingress router of the label-switched path (LSP), which passes through the DR, to break the LSP.</li><li>• The ingress router to re-signal the LSP.</li></ul> |
| <b>Options</b>                  | <p><b>disable</b>—Disable graceful restart for IS-IS.</p> <p><b>helper-disable</b>—Disable graceful restart helper capability. Helper mode is enabled by default.</p> <p><b>restart-duration <i>seconds</i></b>—Time period for the restart to last, in seconds.<br/><b>Range:</b> 30 through 300 seconds<br/><b>Default:</b> 30 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Routing Protocols Graceful Restart</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## hello-authentication-key

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hello-authentication-key password;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure an authentication key (password) for hello packets. Neighboring routing devices use the password to verify the authenticity of packets sent from an interface. For the key to work, you also must include the <b>hello-authentication-type</b> statement.                                                                                                                                                                                                                                         |
| <b>Default</b>                  | By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>password</b> —Authentication password. The password can be up to 255 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">authentication-key on page 3083</a></li> <li>• <a href="#">authentication-type on page 3084</a></li> <li>• <a href="#">hello-authentication-type on page 3094</a></li> </ul>                                                                                                                                                                                                                                                                           |

## hello-authentication-type

---

|                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                              | hello-authentication-type (md5   simple);                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                     | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i> ] |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                 | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                                                                                                                                                                                                                                                                                                         | <p>Enable authentication on an interface for hello packets. If you enable authentication on hello packets, you must specify a password by including the <b>hello-authentication-key</b> statement.</p> <p>You can configure authentication for a given IS-IS level on an interface. On a point-to-point link, if you enable hello authentication for both IS-IS levels, the password configured for Level 1 is used for both levels.</p>                                                                    |
| <div> <b>CAUTION:</b> If no authentication is configured for Level 1 on a point-to-point link with both levels enabled, the hello packets are sent without any password, regardless of the Level 2 authentication configurations.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>                                                                                                                                                                                                                                                                                                             | By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                                                                                                                                                                                                                                                                                                             | <b>md5</b> —Specifies Message Digest 5 as the packet verification type.<br><br><b>simple</b> —Specifies simple authentication as the packet verification type.                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                            | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                               | <ul style="list-style-type: none"><li>• <a href="#">authentication-key on page 3083</a></li><li>• <a href="#">authentication-type on page 3084</a></li><li>• <a href="#">hello-authentication-key on page 3093</a></li></ul>                                                                                                                                                                                                                                                                                |

## hello-interval (Protocols IS-IS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hello-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Modify the frequency with which the routing device sends hello packets out of an interface, in seconds.</p> <p>Routing devices send hello packets at a fixed interval on all interfaces to establish and maintain neighbor relationships. This interval is advertised in the hello interval field in the hello packet.</p> <p>You can send out hello packets in subsecond intervals. To send out hello packets every 333 milliseconds, set the <b>hello-interval</b> value to 1.</p>                                                            |
| <b>Options</b>                  | <p><b><i>seconds</i></b>—Frequency of transmission for hello packets.</p> <p><b>Range:</b> 1 through 20,000 seconds</p> <p><b>Default:</b> 3 seconds (for designated intermediate system [DIS] routers), 9 seconds (for non-DIS routers)</p>                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>hold-time</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## hello-padding

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | hello-padding (adaptive   disable   loose   strict);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>         | <p>Configure padding on hello packets to accommodate asymmetrical maximum transfer units (MTUs) from different hosts.</p> <p>This helps to prevent a premature adjacency Up state when one routing device's MTU does not meet the requirements to establish the adjacency.</p> <p>As an OSI Layer 2 protocol, IS-IS does not support data fragmentation. Therefore, maximum packet sizes must be established and supported between two routers. During adjacency establishment, the IS-IS protocol makes sure that the link supports a packet size of 1492 bytes by padding outgoing hello packets up to the maximum packet size of 1492 bytes.</p> <p>This is the default behavior of the Junos OS IS-IS implementation. However, Junos OS provides an option to disable hello padding that can override this behavior.</p> <p>There are four types of hello padding:</p> <ul style="list-style-type: none"><li>• Adaptive padding—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state type, length, and value (TLV) tuple. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. Adaptive padding has more overhead than loose padding and is able to detect MTU asymmetry from one side of the connection. This one-sided detection can result in generation of extra link-state PDUs that are flooded throughout the network. Specify the <b>adaptive</b> option to configure enough padding to establish an adjacency to neighbors.</li><li>• Disabled padding—Padding is disabled on all types of interfaces for all adjacency states. Specify the <b>disable</b> option to accommodate interfaces that support less than the default packet size of 1492 bytes.</li><li>• Loose padding (the default)—The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. Loose padding might not be able to detect certain situations such as asymmetrical MTUs between the routing devices. Specify the <b>loose</b> option to configure enough padding to initialize an adjacency to neighbors.</li></ul> |



- **Strict padding**—Padding is done on all interface types and for all adjacency states, and is continuous. Strict padding has the most overhead. The advantage is that strict padding detects MTU issues on both sides of a link. Specify the **strict** option to configure padding to allow all adjacency states with neighbors.

**Options**    **adaptive**—Configure padding until the neighbor adjacency is established and active.

**disable**—Disable padding on all types of interfaces for all adjacency states.

**loose**—Configure padding until the state of the adjacency is initialized.

**strict**—Configure padding for all adjacency states.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                         routing-control—To add this statement to the configuration.

**Related Documentation**    • *Example: Configuring IS-IS*

## hold-time (Protocols IS-IS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hold-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code><br><code>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Set the length of time a neighbor considers this router to be operative (up) after receiving a hello packet. If the neighbor does not receive another hello packet within the specified time, it marks this routing device as inoperative (down). The hold time itself is advertised in the hello packets.</p> <p>The hold time specifies how long a neighbor should consider this routing device to be operative without receiving another hello packet. If the neighbor does not receive a hello packet from this routing device within the hold time, it marks the routing device as being unavailable.</p> <p>For systems configured with graceful routing switchover (GRES) with Graceful Restart, the hold time for Master and Backup Routing Engines should be set to a value higher than 40 seconds. This ensures that adjacencies between the Routing Engine and the neighboring peer 'helper' routers do not time out, stopping graceful restart, and all traffic.</p> |
| <b>Options</b>                  | <p><b>seconds</b>—Hold-time value, in seconds.</p> <p><b>Range:</b> 3 through 65,535 seconds, or 1 to send out hello packets every 333 milliseconds</p> <p><b>Default:</b> 9 seconds (for designated intermediate system [DIS] routers), 27 seconds (for non-DIS routers; three times the default hello interval)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Graceful Routing Engine Switchover</i></li><li>• <i>Example: Configuring IS-IS</i></li><li>• <i>Example: Configuring IS-IS for GRES with Graceful Restart</i></li><li>• <a href="#">hello-interval on page 3095</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## ignore-attached-bit

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ignore-attached-bit;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Ignore the attached bit on IS-IS Level 1 routers. Configuring this statement enables the routing device to ignore the attached bit on incoming Level 1 link-state PDUs. If the attached bit is ignored, no default route, which points to the routing device which has set the attached bit, is installed.</p> <p>There might be times, such as during a denial-of-service (DoS) attack, that you do not want a Level 1 router to be able to forward traffic based on a default route.</p> <p>To prevent a routing device from being able to reach interarea destinations, you can prevent the routing device from installing the default route without affecting the status of its IS-IS adjacencies. The <b>ignore-attached-bit</b> statement is used to tell the routing device to ignore the presence of the attached bit in Level 1 link-state PDUs, which blocks the installation of the IS-IS default route.</p> |
| <b>Default</b>                  | The <b>ignore-attached-bit</b> statement is disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | •                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## interface (Protocols IS-IS)

---

```
Syntax interface (all | interface-name) {
    disable;
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        transmit-interval {
            threshold milliseconds;
            minimum-interval milliseconds;
        }
        multiplier number;
    }
    checksum;
    csnp-interval (seconds | disable);
    hello-padding (adaptive | loose | strict);
    ldp-synchronization {
        disable;
        hold-time seconds;
    }
    lsp-interval milliseconds;
    mesh-group (value | blocked);
    no-adjacency-holddown;
    no-ipv4-multicast;
    no-ipv6-multicast;
    no-ipv6-unicast;
    no-unicast-topology;
    passive;
    point-to-point;
    level level-number {
        disable;
        hello-authentication-key key;
        hello-authentication-key-chain key-chain-name;
        hello-authentication-type authentication;
        hello-interval seconds;
        hold-time seconds;
        ipv4-multicast-metric metric;
        ipv6-multicast-metric metric;
        ipv6-unicast-metric metric;
        metric metric;
        passive;
        priority number;
        te-metric metric;
    }
}
```

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>isis</b>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>isis</b>],<br/> [edit protocols <b>isis</b>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols <b>isis</b>]</p>                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/> Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/> Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure interface-specific IS-IS properties. To configure more than one interface, include the <b>interface</b> statement multiple times.</p> <p>Enabling IS-IS on an interface (by including the <b>interface</b> statement at the [edit protocols <b>isis</b>] or the [edit routing-instances <i>routing-instance-name</i> protocols <b>isis</b>] hierarchy level), disabling it (by including the <b>disable</b> statement), and not actually having IS-IS run on an interface (by including the <b>passive</b> statement) are mutually exclusive states.</p> |
| <b>Options</b>                  | <p><b>all</b>—Have Junos OS create IS-IS interfaces automatically. If you include this option, disable IS-IS on the management interface (fxp0).</p> <p><b>interface-name</b>—Name of an interface. Specify the full interface name, including the physical and logical address components.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/> routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IS-IS</i></li> <li>• <i>Example: Configuring Multi-Level IS-IS</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                        |

## ipv4-multicast

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | ipv4-multicast;                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit protocols isis <a href="#">topologies</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                          |
| <b>Description</b>         | Configure alternate IPv4 multicast topologies.                                                                                                                                                                                                                                                                                                                                                          |



**NOTE:** The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.

---

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | Multicast topologies are disabled.                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IS-IS Multicast Topology</i></li></ul>              |

## ipv4-multicast-metric

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ipv4-multicast-metric <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Specify the multicast topology metric value for the level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>metric</i> —Metric value.<br><b>Range:</b> 0 through 16,777,215                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IS-IS Multicast Topology</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                            |

## ipv6-multicast

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ipv6-multicast;</code>                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <b>topologies</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>topologies</b> ],<br>[edit protocols isis <b>topologies</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>topologies</b> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure alternate IPv6 multicast topologies.                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                  | Multicast topologies are disabled.                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IS-IS Multicast Topology</i></li> </ul>                                                                                                                                                                                                                                                            |


## ipv6-multicast-metric

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ipv6-multicast-metric <i>metric</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify the IPv6 alternate multicast topology metric value for the level.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>metric</i> —Metric value.<br><b>Range:</b> 0 through 16,777,215                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IS-IS Multicast Topology</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                              |



## ipv6-unicast

|                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                | ipv6-unicast;                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                       | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">topologies</a> ],<br>[edit protocols isis <a href="#">topologies</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ] |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                   | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                           | Configure alternate IPv6 unicast topologies.<br><br>This statement causes IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to inet6.0.                                                                                                                                                                             |
| <div>  <p><b>NOTE:</b> The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                               | IPv6 unicast topologies are disabled.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                              | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i></li> </ul>                                                                                                                                                                                                                                                                                     |

## ipv6-unicast-metric

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ipv6-unicast-metric <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify the IPv6 unicast topology metric value for the level. The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics.                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <i>metric</i> —Metric value.<br><b>Range:</b> 0 through 16,777,215                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                |

## isis

|                                 |                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | isis { ... }                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                      |
| <b>Description</b>              | Enable IS-IS routing on the routing device or for a routing instance.<br><br>The <b>isis</b> statement is the one statement you must include in the configuration to run IS-IS on the routing device or in a routing instance.                                      |
| <b>Default</b>                  | IS-IS is disabled on the routing device.                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IS-IS</i></li> <li>• <i>Example: Configuring Multi-Level IS-IS</i></li> </ul>                                                                                                                      |

## level (Global IS-IS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>level <i>level-number</i> {<br/>    authentication-key <i>key</i>;<br/>    authentication-key-chain (Protocols IS-IS) <i>key-chain-name</i>;<br/>    authentication-type <i>type</i>;<br/>    disable;<br/>    external-preference <i>preference</i>;<br/>    no-csnp-authentication;<br/>    no-hello-authentication;<br/>    no-psnp-authentication;<br/>    preference <i>preference</i>;<br/>    wide-metrics-only;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols isis],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>  isis],<br/>[edit protocols isis],<br/>[edit routing-instances <i>routing-instance-name</i> protocols isis]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure the global-level properties.</p> <p>You can administratively divide a single AS into smaller groups called areas. You configure each routing device interface to be in an area. Any interface can be in any area. The area address applies to the entire routing device. You cannot specify one interface to be in one area and another interface in a different area. To route between areas, you must have two adjacent Level 2 routers that communicate with each other.</p> <p>Level 1 routers can only route within their IS-IS area. To send traffic outside their area, Level 1 routers must send packets to the nearest intra-area Level 2 router. A routing device can be a Level 1 router, a Level 2 router, or both. You specify the router level on a per-interface basis, and a routing device becomes adjacent to other routing devices on the same level on that link only.</p> <p>You can configure one Level 1 routing process and one Level 2 routing process on each interface, and you can configure the two levels differently.</p> |
| <b>Options</b>                  | <p><i>level-number</i>—IS-IS level number.</p> <p><b>Values:</b> 1 or 2</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

- Related Documentation**
- *Example: Configuring IS-IS*
  - *Example: Configuring Multi-Level IS-IS*

## link-protection (Protocols IS-IS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-protection;                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit protocols isis interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                  |
| <b>Description</b>              | Enable link protection on the specified IS-IS interface. Junos OS creates a backup loop-free alternate path to the primary next hop for all destination routes that traverse the protected interface.                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Link and Node Protection for IS-IS Routes</i></li> <li>• <a href="#">node-link-protection on page 3123</a></li> </ul>                                                                                                                                                                                                                                      |

## loose-authentication-check

|                                 |                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | loose-authentication-check;                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                              |
| <b>Description</b>              | Allow the use of MD5 authentication without requiring network-wide deployment.                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i></li> </ul>                                                                                                                                                                                                                               |

## **lsp-interval**

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>lsp-interval milliseconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure the link-state PDU interval time.</p> <p>By default, the routing device sends one link-state PDU packet out an interface every 100 milliseconds. To disable the transmission of all link-state PDUs, set the interval to 0.</p> <p>Link-state PDU throttling by use of the <b>lsp-interval</b> statement controls the flooding pace to neighboring routing devices in order to not overload them.</p> <p>Also, consider that control traffic (such as link-state PDUs and related packets) might delay user traffic (information packets) because control traffic always has precedence in terms of scheduling on the routing device interface cards. Unfortunately, the control traffic transmission rate is not decreased on low-bandwidth interfaces, such as DS-0 or fractional T1 and E1 interface. Line control traffic stays the same. On a low-bandwidth circuit that is transmitting 30 full-MTU-sized packets, there is not much bandwidth left over for other types of packets.</p> |
| <b>Default</b>                  | By default, the routing device sends one link-state PDU out an interface every 100 milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <b>milliseconds</b> —Number of milliseconds between the sending of link-state PDUs. Specifying a value of 0 blocks all link-state PDU transmission.<br><b>Range:</b> 0 through 1000 milliseconds<br><b>Default:</b> 100 milliseconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## lsp-lifetime

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>lsp-lifetime <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],</p> <p>[edit protocols <i>isis</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Specify how long a link-state PDU originating from the routing device should persist in the network. The routing device sends link-state PDUs often enough so that the link-state PDU lifetime never expires.</p> <p>Because link-state PDUs have a maximum lifetime, they need to be refreshed. Refreshing means that a routing device needs to re-originate its link-state PDUs periodically. The re-origination interval must be less than the link-state PDU's lifetime. For example, if the link-state PDU is valid for 1200 seconds, the routing device needs to refresh the link-state PDU in less than 1200 seconds to avoid removal of the link-state PDU from the link-state database by other routing devices. The recommended maximum link-state PDU origination interval is the lifetime minus 300 seconds. So, in a default environment this would be 900 seconds. In Junos OS, the refresh interval is derived from the lifetime and is equal to the lifetime minus 317 seconds. You can change the lifetime to a higher value to reduce the number of refreshes in the network. (You would rarely want to increase the number of refreshes.) Often these periodic link-state PDU refreshes are referred to as refresh noise, and network administrators want to reduce this noise as much as possible.</p> <p>The <code>show isis overview</code> command displays the link-state PDU lifetime.</p> |
| <b>Default</b>                  | By default, link-state PDUs are maintained in network databases for 1200 seconds (20 minutes) before being considered invalid. This length of time, called the <i>LSP lifetime</i> , normally is sufficient to guarantee that link-state PDUs never expire.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b><i>seconds</i></b>—link-state PDU lifetime, in seconds.</p> <p><b>Range:</b> 350 through 65,535 seconds</p> <p><b>Default:</b> 1200 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i></li> <li>• <a href="http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf">http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## max-areas

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>max-areas <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <i>isis</i>]</code><br><code>[edit protocols <i>isis</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</code>                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Modify the maximum number of IS-IS areas advertised.</p> <p>This value is included in the Maximum Address Area field of the IS-IS common PDU header included in all outgoing PDUs.</p> <p>The maximum number of areas you can advertise is restricted to 36 to ensure that the IIH PDUs have enough space to include other type, length, and value (TLV) fields, such as the Authentication and IPv4 and IPv6 Interface Address TLVs.</p> |
| <b>Options</b>                  | <p><b><i>number</i></b>—Maximum number of areas to include in the IS-IS hello (IIH) PDUs and link-state PDUs.</p> <p><b>Range:</b> 3 through 36</p> <p><b>Default:</b> 3</p>                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multi-Level IS-IS</i></li></ul>                                                                                                                                                                                                                                                                                                                                              |



## mesh-group (Protocols IS-IS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mesh-group (blocked   <i>value</i> );                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ]                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure an interface to be part of a mesh group, which is a set of fully connected nodes.</p> <p>A <i>mesh group</i> is a set of routing devices that are fully connected. That is, they have a fully meshed topology. When link-state PDUs are being flooded throughout an area, each router within a mesh group receives only a single copy of a link-state PDU instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of link-state PDUs.</p> <p>To create a mesh group and designate that an interface be part of the group, assign a mesh-group number to all the routing device interfaces in the group. To prevent an interface in the mesh group from flooding link-state PDUs, configure blocking on that interface.</p> |
| <b>Options</b>                  | <p><b>blocked</b>—Configure the interface so that it does not flood link-state PDUs.</p> <p><b>value</b>—Number that identifies the mesh group.</p> <p><b>Range:</b> 1 through 4,294,967,295 (<math>2^{32} - 1</math>; 32 bits are allocated to identify a mesh group)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Mesh Groups of IS-IS Interfaces</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## metric (Protocols IS-IS)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>metric <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>                                                                                                              |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>         | <p>Specify the metric value for the level.</p> <p>All IS-IS routes have a cost, which is a routing metric that is used in the IS-IS link-state calculation. The cost is an arbitrary, dimensionless integer that can be from 1 through 63, or from 1 through 16,777,215 (<math>2^{24} - 1</math>) if you are using wide metrics.</p> <p>Similar to other routing protocols, IS-IS provides a way of exporting routes from the routing table into the IS-IS network. When a route is exported into the IS-IS network without a specified metric, IS-IS uses default metric values for the route, depending on the protocol that was used to learn the route.</p> |

Table 326 on page 3114 depicts IS-IS route export metric default values.

**Table 326: Default Metric Values for Routes Exported into IS-IS**

| Protocol Used for Learning the Route | Default Metric Value                                          |
|--------------------------------------|---------------------------------------------------------------|
| Direct                               | 10                                                            |
| Static                               | Same as reported by the protocol used for exporting the route |
| Aggregate                            | 10                                                            |
| Generate                             | 10                                                            |
| RIP                                  | Same as reported by the protocol used for exporting the route |
| OSPF                                 | Same as reported by the protocol used for exporting the route |
| BGP                                  | 10                                                            |

The default metric values behavior can be customized by using routing policies.

- Options** *metric*—Metric value.
- Range:** 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)

**Default:** 10 (for all interfaces except lo0), 0 (for the lo0 interface)

|                                 |                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i></li> <li>• <i>te-metric</i></li> <li>• <a href="#">wide-metrics-only on page 3139</a></li> </ul> |

## no-adjacency-holddown

|                                 |                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-adjacency-holddown;                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                   |
| <b>Description</b>              | Disable the hold-down timer for IS-IS adjacencies.<br><br>A hold-down timer delays the advertising of adjacencies by waiting until a time period has elapsed before labeling adjacencies in the up state. You can disable this hold-down timer, which labels adjacencies up faster. However, disabling the hold-down timer creates more frequent link-state PDU updates and SPF computation. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">hold-time on page 3098</a></li> </ul>                                                                                                                                                                                                                                                                                                   |

## no-authentication-check

---

|                                 |                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-authentication-check;                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                              |
| <b>Description</b>              | Generate authenticated packets and check the authentication on received packets, but do not reject packets that cannot be authenticated.                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">hello-authentication-type on page 3094</a></li></ul>                                                                                                                                                                                                                                                    |

## no-csnp-authentication

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-csnp-authentication;                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                          |
| <b>Description</b>              | Suppress authentication check on complete sequence number PDU (CSNP) packets.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">csnp-interval on page 3088</a></li></ul>                                                                                                                                                                                                                                                                                                                                            |

## no-eligible-backup (Protocols IS-IS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-eligible-backup;                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis interface <i>interface-name</i> ],<br>[edit protocols isis interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                     |
| <b>Description</b>              | Exclude the specified interface as a backup interface for IS-IS interfaces on which link protection or node-link protection is enabled.                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Link and Node Protection for IS-IS Routes</i></li> <li>• <a href="#">link-protection on page 3109</a></li> <li>• <a href="#">node-link-protection on page 3123</a></li> </ul>                                                                                                                                                                                 |

## no-hello-authentication


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-hello-authentication;                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                             |
| <b>Description</b>              | Suppress authentication check on complete sequence number hello packets.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">hello-authentication-type on page 3094</a></li> </ul>                                                                                                                                                                                                                                                                                                                                 |

## no-ipv4-multicast

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-ipv4-multicast;                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Exclude an interface from IPv4 multicast topologies.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>                  | Multicast topologies are disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IS-IS Multicast Topology</i></li></ul>                                                                                                                                                                                                                                                                                                                                                             |

## no-ipv4-routing

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-ipv4-routing;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>],<br/> [edit protocols <a href="#">isis</a>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>]</p>                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Disable IP version 4 (IPv4) routing.</p> <p>Disabling IPv4 routing has the following results:</p> <ul style="list-style-type: none"> <li>• The routing device does not advertise the network layer protocol identifier (NLPID) for IPv4 in the Junos OS link-state PDU fragment zero.</li> <li>• The routing device does not advertise any IPv4 prefixes in Junos OS link-state PDUs.</li> <li>• The routing device does not advertise the NLPID for IPv4 in Junos OS hello packets.</li> <li>• The routing device does not advertise any IPv4 addresses in Junos OS hello packets.</li> <li>• The routing device does not calculate any IPv4 routes.</li> </ul> |
|                                 | <p> <b>NOTE:</b> Note: Even when <code>no-ipv4-routing</code> is configured, an IS-IS traceoptions log can list rejected IPv4 addresses. When a configuration is committed, IS-IS schedules a scan of the routing table to determine whether any routes need to be exported into the IS-IS link state database. The implicit default export policy action is to reject everything. IPv4 addresses from the routing table are examined for export, rejected by the default policy, and the rejections are logged.</p>                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## no-ipv6-multicast

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-ipv6-multicast;                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Exclude an interface from the IPv6 multicast topologies.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | Multicast topologies are disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IS-IS Multicast Topology</i></li></ul>                                                                                                                                                                                                                                                                                                                                                             |



## no-ipv6-routing

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-ipv6-routing;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Disable IP version 6 (IPv6) routing.<br><br>Disabling IPv6 routing has the following results: <ul style="list-style-type: none"> <li>• The routing device does not advertise the network layer protocol identifier (NLPID) for IPv6 in the Junos OS link-state PDU fragment zero.</li> <li>• The routing device does not advertise any IPv6 prefixes in Junos OS link-state PDUs.</li> <li>• The routing device does not advertise the NLPID for IPv6 in Junos OS hello packets.</li> <li>• The routing device does not advertise any IPv6 addresses in Junos OS hello packets.</li> <li>• The routing device does not calculate any IPv6 routes.</li> </ul> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## no-ipv6-unicast

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-ipv6-unicast;                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Exclude an interface from the IPv6 unicast topologies. This enables you to exercise control over the paths that unicast data takes through a network.                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>                  | IPv6 unicast topologies are disabled.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i></li></ul>                                                                                                                                                                                                                                                                                                                                               |

## no-psnp-authentication

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-psnp-authentication;                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                             |
| <b>Description</b>              | Suppress authentication check on partial sequence number PDU (PSNP) packets.                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring IS-IS Authentication</i></li></ul>                                                                                                                                                                                                                                                                                                                                                  |

## no-unicast-topology

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-unicast-topology;                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Exclude an interface from the IPv4 unicast topologies.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>                  | IPv4 unicast topologies are disabled.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IS-IS Multicast Topology</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                        |

## node-link-protection (Protocols IS-IS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | node-link-protection;                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit protocols isis interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                  |
| <b>Description</b>              | Enable node-link protection on the specified IS-IS interface. Junos OS creates an alternate loop-free path to the primary next hop for all destination routes that traverse a protected interface. This alternate path avoids the primary next-hop routing device altogether and establishes a path through a different routing device.                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Link and Node Protection for IS-IS Routes</i></li> <li>• <a href="#">link-protection on page 3109</a></li> </ul>                                                                                                                                                                                                                                           |

## overload (Protocols IS-IS)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>overload {<br/>    advertise-high-metrics;<br/>    allow-route-leaking;<br/>    timeout <i>seconds</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <i>isis</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i> ],<br>[edit protocols <i>isis</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>         | <p>Configure the local routing device so that it appears to be overloaded. This statement causes the routing device to continue participating in IS-IS routing, but prevents it from being used for transit traffic. Traffic destined to immediately attached subnets continues to transit the routing device.</p> <p>You can also advertise maximum link metrics in network layer reachability information (NLRI) instead of setting the overload bit.</p> <p>You configure or disable overload mode in IS-IS with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the IS-IS instance started is less than the specified timeout.</p> <p>A timer is started for the difference between the timeout and the time elapsed since the instance started. If the time elapsed after the IS-IS instance is enabled is less than the specified timeout, overload mode is set. When the timer expires, overload mode is cleared. In overload mode, the routing device IS-IS advertisements are originated with the overload bit set. This causes the transit traffic to take paths around the routing device. However, the overloaded routing device's own links are still accessible.</p> <p>The value of the overload bit depends on these three scenarios:</p> <ol style="list-style-type: none"><li>1. When the overload bit has already been set to a given value and the routing process is restarted: Link-state PDUs are regenerated with the overload bit cleared.</li><li>2. When the overload bit is reset to a lesser value while the routing process is running: Link-state PDUs are regenerated with the overload bit cleared.</li><li>3. When the overload bit is reset to a greater value while the routing process is running: Link-state PDUs are regenerated with the overload bit set to the difference between the old and new value.</li></ol> <p>In overload mode, the routing device advertisement is originated with all the transit routing device links (except stub) set to a metric of 0xFFFF. The stub routing device links are</p> |

advertised with the actual cost of the interfaces corresponding to the stub. This causes the transit traffic to avoid the overloaded routing device and take paths around the routing device.

To understand the reason for setting the overload bit, consider that BGP converges slowly. It is not very good at detecting that a neighbor is down because it has slow-paced keepalive timers. Once the BGP neighbor is determined to be down, it can take up to 2 minutes for a BGP router to declare the neighbor down. IS-IS is much quicker. IS-IS only takes 10-30 seconds to detect absent peers. It is the slowness of BGP, more precisely the slowness of internal BGP (IBGP), that necessitates the use of the overload bit. IS-IS and BGP routing are mutually dependent on each other. If both do not converge at the same time, traffic is dropped without notification (black holed).

You might want to configure the routing device so that it appears to be overloaded when you are restarting routing on the device. Setting the overload bit for a fixed amount of time right after a restart of the routing protocol process (rpd) ensures that the router does not receive transit traffic while the routing protocols (especially IBGP) are still converging.

Setting the overload bit is useful when performing hardware or software maintenance work on a routing device. After the maintenance work, clear the overload bit to carry on forwarding transit traffic. Manual clearing of the overload bit is not always possible. What is needed is an automated way of clearing the overload bit after some amount of time. Most networks use a time value of 300 seconds. This 5-minute value provides a good balance, allowing time to bring up even large internal IBGP meshes, while still relatively quick.

Another appropriate application for setting for the overload bit is on dedicated devices such as BGP route reflectors, which are intentionally not meant to carry any transit traffic. In this case, you would not use the timer.

You can verify that the overload bit is set by running the **show isis database** command.

**Options**    **advertise-high-metrics**—Advertise maximum link metrics in NLRIs instead of setting the overload bit.

The **advertise-high-metric** setting is only valid while the routing device is in overload mode.

When **advertise-high-metric** is configured, IS-IS does not set the overload bit. Rather, it sets the metric to 63 or 16,777,214, depending whether wide metrics are enabled. This allows the overloaded routing device to be used for transit as a last resort.

An L1-L2 router in overload mode stops leaking route information between L1 and L2 levels and clears its attached bit. This is also true when **advertise-high-metrics** is configured.

**allow-route-leaking**—Enable leaking of route information into the network even if the overload bit is set.



**NOTE:** The **allow-route-leaking** option does not work if the routing device is in dynamic overload mode. Dynamic overload can occur if the device has exceeded its resource limits, such as the prefix limit.

---

**timeout seconds**—Number of seconds at which the overloading is reset.


**Range:** 60 through 1800 seconds

**Default:** 0 seconds

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                         routing-control—To add this statement to the configuration.

**Related Documentation**    • *Example: Configuring IS-IS*

## passive (Protocols IS-IS)

|                                                                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                     | <code>passive;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                            | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],<br/>         [edit protocols isis <b>interface</b> <i>interface-name</i>],<br/>         [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                        | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                | <p>Advertise the direct interface addresses on an interface or into a level on the interface without actually running IS-IS on that interface or level.</p> <p>This statement effectively prevents IS-IS from running on the interface. To enable IS-IS on an interface, include the <b>interface</b> statement at the [edit protocols isis] or the [edit routing-instances <i>routing-instance-name</i> protocols isis] hierarchy level. To disable it, include the <b>disable</b> statement at those hierarchy levels. The three states—enabling, disabling, or not running IS-IS on an interface—are mutually exclusive.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <div>  <p><b>NOTE:</b> Configuring IS-IS on a loopback interface automatically renders it as a passive interface, irrespective of whether the <b>passive</b> statement was used in the configuration of the interface.</p> </div>                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <p>If neither passive mode nor the <b>family iso</b> option is configured on the IS-IS interface, then the routing device treats the interface as not being operational, and no direct IPv4/IPv6 routes are exported into IS-IS. (You configure the <b>family iso</b> option at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.)</p> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                    | By default, IS-IS must be configured on an interface or a level for direct interface addresses to be advertised into that level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                   | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                      | <ul style="list-style-type: none"> <li>Example: Configuring Multi-Level IS-IS</li> <li>disable</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## point-to-point

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | point-to-point;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <b>interface</b> <i>interface-name</i> ],<br>[edit protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i> ]                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Configure an IS-IS interface to behave like a point-to-point connection.</p> <p>You can use the <b>point-to-point</b> statement to configure a LAN interface to act like a point-to-point interface for IS-IS. You do not need an unnumbered LAN interface, and it has no effect if configured on an interface that is already point-to-point.</p> <p>The <b>point-to-point</b> statement affects only IS-IS protocol procedures on that interface. All other protocols continue to treat the interface as a LAN interface. Only two IS-IS routing devices can be connected to the LAN interface, and both must be configured as point-to-point.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>IS-IS Overview</i></li><li>• <i>Understanding IS-IS Designated Routers</i></li><li>• <i>Example: Configuring Synchronization Between IS-IS and LDP</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                             |



## preference (Protocols IS-IS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>],</p> <p>[edit protocols isis <a href="#">level level-number</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>]</p>                |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure the preference of internal routes.</p> <p>Route preferences (also known as administrative distances) are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.</p> <p>To change the preference values, include the <b>preference</b> statement (for internal routes) or the <b>external-preference</b> statement.</p> |
| <b>Options</b>                  | <p><b><i>preference</i></b>—Preference value.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)</p>                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Route Preferences Overview</i></li> <li>• <i>Example: Redistributing OSPF Routes into IS-IS</i></li> <li>• <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i></li> <li>• <a href="#">external-preference on page 3091</a></li> </ul>                                                                                                                                                      |

## prefix-export-limit (Protocols IS-IS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>prefix-export-limit <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>isis <a href="#">level level-number</a>],</code><br><code>[edit protocols isis <a href="#">level level-number</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Configure a limit to the number of prefixes exported into IS-IS.</p> <p>By default, there is no limit to the number of prefixes that can be exported into IS-IS. To configure a limit to the number of prefixes that can be exported into IS-IS, include the <b>prefix-export-limit</b> statement. The <b>prefix-export-limit</b> statement protects the rest of the network from a malicious policy by applying a threshold filter for exported routes.</p> <p>The number of prefixes depends on the size of your network. Good design advice is to set it to double the total number of IS-IS Level 1 and Level 2 routing devices in your network.</p> <p>If the number of prefixes exported into IS-IS exceeds the configured limit, the overload bit is set and the overload state is reached. When other routers detect that this bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes. The overload state can be cleared by using the <a href="#">clear isis overload</a> command.</p> <p>The <a href="#">show isis overview</a> command displays the prefix export limit when it is configured.</p> |
| <b>Options</b>                  | <p><b><i>number</i></b>—Prefix limit.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> None</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i></li><li>• <i>Example: Redistributing OSPF Routes into IS-IS</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## priority (Protocols IS-IS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Configure the interface's priority for becoming the designated router. The interface with the highest priority value becomes that level's designated router.</p> <p>The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.</p> <p>A routing device advertises its priority to become a designated router in its hello packets. On all multiaccess networks, IS-IS uses the advertised priorities to elect a designated router for the network. This routing device is responsible for sending network link-state advertisements, which describe all the routing devices attached to the network. These advertisements are flooded throughout a single area.</p> <p>A routing device's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127. Routing devices with a higher value are more likely to become the designated router.</p> |
| <b>Options</b>                  | <p><i>number</i>—Priority value.</p> <p><b>Range:</b> 0 through 127</p> <p><b>Default:</b> 64</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IS-IS Designated Routers</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## reference-bandwidth (Protocols IS-IS)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>reference-bandwidth <i>reference-bandwidth</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],</code><br><code>[edit protocols <i>isis</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Optimize routing based on bandwidth by setting the reference bandwidth used in calculating the default interface cost.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. When there are several equal-cost routes to a destination, traffic is distributed equally among them.</p> <p>The cost of a route is described by a single dimensionless metric that is determined using the following formula:</p> $\text{cost} = \text{reference-bandwidth} / \text{bandwidth}$ <p>For example, if you set the reference bandwidth to 1 Gbps (that is, <i>reference-bandwidth</i> is set to 1,000,000,000), a 100-Mbps interface has a routing metric of 10.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics.</p> |
| <b>Options</b>                  | <p><i>reference-bandwidth</i>—Reference bandwidth value in bits per second.</p> <p><b>Range:</b> 9600 through 1,000,000,000,000 bps</p> <p><b>Default:</b> None</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IS-IS</i></li><li>• <a href="http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf">http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## rib-group (Protocols IS-IS)

|                                 |                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> rib-group {     inet <i>group-name</i>;     inet6 <i>group-name</i>; } </pre>                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>],</p> <p>[edit protocols <a href="#">isis</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                              |
| <b>Description</b>              | <p>Install routes learned from IS-IS routing instances into routing tables in the IS-IS routing table group. You can install IPv4 routes or IPv6 routes.</p> <p>Support for IPv6 routing table groups in IS-IS enables IPv6 routes that are learned from IS-IS routing instances to be installed into other routing tables defined in an IS-IS routing table group.</p>    |
| <b>Options</b>                  | <p><b><i>group-name</i></b>—Name of the routing table group.</p> <p><b>inet</b>—Install IPv4 IS-IS routes.</p> <p><b>inet6</b>—Install IPv6 IS-IS routes.</p>                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p><b>routing</b>—To view this statement in the configuration.</p> <p><b>routing-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li> <li>• <i>Example: Importing Direct and Static Routes Into a Routing Instance</i></li> <li>• <i>Understanding Multiprotocol BGP</i></li> </ul>                                                                                    |

## spf-options (Protocols IS-IS)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>spf-options {<br/>    delay <i>milliseconds</i>;<br/>    holddown <i>milliseconds</i>;<br/>    rapid-runs <i>number</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <i>isis</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i> ],<br>[edit protocols <i>isis</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>         | <p>Configure options for running the shortest-path-first (SPF) algorithm.</p> <p>Running the SPF algorithm is usually the beginning of a series of larger system-wide events. For example, the SPF algorithm can lead to interior gateway protocol (IGP) prefix changes, which then lead to BGP nexthop resolution changes. Consider what happens if there are rapid link changes in the network. The local routing device can become overwhelmed. This is why it sometimes makes sense to throttle the scheduling of the SPF algorithm.</p> <p>You can configure the following SPF options:</p> <ul style="list-style-type: none"><li>• The delay in the time between the detection of a topology change and when the SPF algorithm actually runs.</li><li>• The maximum number of times that the SPF algorithm can run in succession before the hold-down timer begins.</li><li>• The time to hold down, or wait, before running another SPF calculation after the SPF algorithm has run in succession the configured maximum number of times.</li></ul> <p>If the network stabilizes during the hold-down period and the SPF algorithm does not need to run again, the system reverts to the configured values for the <b>delay</b> and <b>rapid-runs</b> statements.</p> |
| <b>Options</b>             | <p><b>delay <i>milliseconds</i></b>—Time interval between the detection of a topology change and when the SPF algorithm runs.</p> <p><b>Range:</b> 50 through 1000 milliseconds</p> <p><b>Default:</b> 200 milliseconds</p> <p><b>holddown <i>milliseconds</i></b>—Time interval to hold down, or wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession.</p> <p><b>Range:</b> 2000 through 10,000 milliseconds</p> <p><b>Default:</b> 5000 milliseconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

**rapid-runs *number***—Maximum number of times the SPF algorithm can run in succession.  
After the maximum is reached, the holddown interval begins.

**Range:** 1 through 5

**Default:** 3

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Link and Node Protection for IS-IS Routes*

## topologies (Protocols IS-IS)

**Syntax**

```
topologies {
  ipv4-multicast;
  ipv6-multicast;
  ipv6-unicast;
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols *isis*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols *isis*],  
[edit protocols *isis*],  
[edit routing-instances *routing-instance-name* protocols *isis*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure alternate IS-IS topologies.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies*
- *Example: Configuring IS-IS Multicast Topology*

## traceoptions (Protocols IS-IS)

---

|                            |                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>name</i> &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;;<br/>}</pre>                                                                                                                                   |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                              |
| <b>Description</b>         | Configure IS-IS protocol-level tracing options. To specify more than one tracing operation, include multiple <b>flag</b> statements.                                                                                                                                                                                                                        |



**NOTE:** The **traceoptions** statement is not supported on QFabric systems.

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b> | The default IS-IS protocol-level tracing options are those inherited from the routing protocols <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b> | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>name</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks (" "). All files are placed in the directory <b>/var/log</b>. We recommend that you place IS-IS tracing output in the file <b>isis-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 10 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one flag, include multiple <b>flag</b> statements.</p> <p><b>IS-IS Protocol-Specific Tracing Flags</b></p> |



- **csn**—Complete sequence number PDU (CSNP) packets
- **error**—Errored IS-IS packets
- **graceful-restart**—Graceful restart operation
- **hello**—Hello packets
- **ldp-synchronization**—Synchronization between IS-IS and LDP
- **lsp**—Link-state PDUs
- **lsp-generation**—Link-state PDU generation packets
- **packets**—All IS-IS protocol packets
- **psn**—Partial sequence number PDU (PSNP) packets
- **spf**—Shortest-path-first calculations

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations, including adjacency changes

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.                                   |
|                                 | routing-control and trace-control—To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | • <i>Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces</i>           |
|                                 | • <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i> |
|                                 | • <i>Example: Enabling Packet Checksums on IS-IS Interfaces</i>                                  |

## wide-metrics-only

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | wide-metrics-only;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <b>level</b> <i>level-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <b>level</b> <i>level-number</i> ],<br>[edit protocols isis <b>level</b> <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>level</b> <i>level-number</i> ]                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure IS-IS to generate metric values greater than 63 on a per IS-IS level basis.<br><br>Normally, IS-IS metrics can have values up to 63, and IS-IS generates two type, length, and value (TLV) tuples, one for an IS-IS adjacency and the second for an IP prefix. To allow IS-IS to support traffic engineering, a second pair of TLVs has been added to IS-IS, one for IP prefixes and the second for IS-IS adjacency and traffic engineering information. With these TLVs, IS-IS metrics can have values up to 16,777,215 ( $2^{24} - 1$ ).<br><br>To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, include the <b>wide-metrics-only</b> statement. |
| <b>Default</b>                  | By default, Junos OS supports the sending and receiving of wide metrics. Junos OS allows a maximum metric value of 63 and generates both pairs of TLVs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i></li> <li>• <i>te-metric</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



## CHAPTER 53

# Administration

- [Operational Commands on page 3141](#)

### Operational Commands

---

- [clear isis adjacency](#)
- [clear isis database](#)
- [clear isis overload](#)
- [clear isis statistics](#)
- [show isis adjacency](#)
- [show isis authentication](#)
- [show isis backup coverage](#)
- [show isis backup label-switched-path](#)
- [show isis backup spf results](#)
- [show isis database](#)
- [show isis hostname](#)
- [show isis interface](#)
- [show isis overview](#)
- [show isis route](#)
- [show isis spf](#)
- [show isis statistics](#)

## clear isis adjacency

---

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3142</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3142</a>                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax</b>                                     | <pre>clear isis adjacency &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;neighbor&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>clear isis adjacency &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;neighbor&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>                                | Remove entries from the IS-IS adjacency database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                                    | <p><b>none</b>—Remove all entries from the adjacency database.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear all adjacencies for the specified routing instance only.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear all adjacencies for the specified interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear adjacencies for the specified neighbor only.</p> |
| <b>Required Privilege Level</b>                   | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">show isis adjacency on page 3150</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>                      | <a href="#">clear isis adjacency on page 3142</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>                              | See <a href="#">show isis adjacency</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Sample Output

### clear isis adjacency

The following sample output displays IS-IS adjacency database information before and after the **clear isis adjacency** command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State          Hold (secs) SNPA
```

```
so-1/0/0.0    karaku1      3 Up                26
so-1/1/3.0    1921.6800.5080 3 Up                23
so-5/0/0.0    1921.6800.5080 3 Up                19
```

```
user@host> clear isis adjacency karaku1
```

```
user@host> show isis adjacency
```

```
IS-IS adjacency database:
```

| Interface  | System         | L State        | Hold (secs) | SNPA |
|------------|----------------|----------------|-------------|------|
| so-1/0/0.0 | karaku1        | 3 Initializing | 26          |      |
| so-1/1/3.0 | 1921.6800.5080 | 3 Up           | 24          |      |
| so-5/0/0.0 | 1921.6800.5080 | 3 Up           | 21          |      |

## clear isis database

---

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3144</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3144</a>                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax</b>                                     | <pre>clear isis database &lt;entries&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>clear isis database &lt;entries&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                               |
| <b>Description</b>                                | Remove the entries from the IS-IS link-state database, which contains prefixes and topology information.                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                                    | <b>none</b> —Remove all entries from the IS-IS link-state database for all routing instances.<br><br><b>entries</b> —(Optional) Name of the database entry.<br><br><b>instance <i>instance-name</i></b> —(Optional) Clear all entries for the specified routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | clear                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">show isis database on page 3164</a></li></ul>                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>                      | <a href="#">clear isis database on page 3144</a>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>                              | See <a href="#">show isis database</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                            |

## Sample Output

### clear isis database

The following sample output displays IS-IS link-state database information before and after the **clear isis database** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
crater.00-00          0x12   0x84dd             1139
  1 LSPs
IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
```



|                      |      |        |      |
|----------------------|------|--------|------|
| crater.00-00         | 0x19 | 0xe92c | 1134 |
| badlands.00-00       | 0x16 | 0x1454 | 985  |
| carlsbad.00-00       | 0x33 | 0x220b | 1015 |
| ranier.00-00         | 0x2e | 0xfc31 | 1007 |
| 1921.6800.5066.00-00 | 0x11 | 0x7313 | 566  |
| 1921.6800.5067.00-00 | 0x14 | 0xd9d4 | 939  |

6 LSPs

user@host> clear isis database

user@host> show isis database

IS-IS level 1 link-state database:

| LSP ID | Sequence | Checksum | Lifetime (secs) |
|--------|----------|----------|-----------------|
|--------|----------|----------|-----------------|

IS-IS level 2 link-state database:

| LSP ID | Sequence | Checksum | Lifetime (secs) |
|--------|----------|----------|-----------------|
|--------|----------|----------|-----------------|

## clear isis overload

---

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                             | <a href="#">Syntax on page 3146</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3146</a>                                                                                                                                                                                                                                                                                                                                                                      |
| Syntax                                     | <code>clear isis overload</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                                                                                                    |
| Syntax (EX Series Switches and QFX Series) | <code>clear isis overload</code><br><code>&lt;instance <i>instance-name</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                              |
| Release Information                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                            |
| Description                                | <p>Reset the IS-IS dynamic overload bit. This command can appear to not work, continuing to display <b>overload</b> after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.</p> <p>When other routers detect that the overload bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes.</p> |
| Options                                    | <p><b>none</b>—Reset the IS-IS dynamic overload bit.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>                                                                                                                                |
| Required Privilege Level                   | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Related Documentation                      | <ul style="list-style-type: none"><li>• <a href="#">show isis database on page 3164</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                   |
| List of Sample Output                      | <a href="#">clear isis overload on page 3146</a>                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Output Fields                              | See <a href="#">show isis database</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                         |

## Sample Output

### clear isis overload

The following sample output displays IS-IS database information before and after the **clear isis overload** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                               Sequence Checksum Lifetime Attributes
```

```
pro3-c.00-00          0x4   0x10db    1185 L1 L2 Overload
```

```
1 LSPs
```

```
IS-IS level 2 link-state database:
```

| LSP ID       | Sequence | Checksum | Lifetime | Attributes            |
|--------------|----------|----------|----------|-----------------------|
| pro3-c.00-00 | 0x5      | 0x429f   | 1185     | L1 L2 <b>Overload</b> |

```
pro2-a.00-00          0x91e  0x2589    874 L1 L2
```

```
pro2-a.02-00          0x1    0xcbc    874 L1 L2
```

```
3 LSPs
```

```
user@host> clear isis overload
```

```
user@host> show isis database
```

```
IS-IS level 1 link-state database:
```

| LSP ID       | Sequence | Checksum | Lifetime | Attributes |
|--------------|----------|----------|----------|------------|
| pro3-c.00-00 | 0xa      | 0x429e   | 1183     | L1 L2      |

```
1 LSPs
```

```
IS-IS level 2 link-state database:
```

| LSP ID       | Sequence | Checksum | Lifetime | Attributes |
|--------------|----------|----------|----------|------------|
| pro3-c.00-00 | 0xc      | 0x9c39   | 1183     | L1 L2      |
| pro2-a.00-00 | 0x91e    | 0x2589   | 783      | L1 L2      |
| pro2-a.02-00 | 0x1      | 0xcbc    | 783      | L1 L2      |

```
3 LSPs
```

## clear isis statistics

---

|                                            |                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                             | <a href="#">Syntax on page 3148</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3148</a>                                                                                                                                                                                                                                                                  |
| Syntax                                     | <code>clear isis statistics</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                              |
| Syntax (EX Series Switches and QFX Series) | <code>clear isis statistics</code><br><code>&lt;instance <i>instance-name</i>&gt;</code>                                                                                                                                                                                                                                                                                        |
| Release Information                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                        |
| Description                                | Set statistics about IS-IS traffic to zero.                                                                                                                                                                                                                                                                                                                                     |
| Options                                    | <b>none</b> —Set IS-IS traffic statistics to zero for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level                   | view                                                                                                                                                                                                                                                                                                                                                                            |
| Related Documentation                      | <ul style="list-style-type: none"><li>• <a href="#">show isis statistics on page 3194</a></li></ul>                                                                                                                                                                                                                                                                             |
| List of Sample Output                      | <a href="#">clear isis statistics on page 3148</a>                                                                                                                                                                                                                                                                                                                              |
| Output Fields                              | See <a href="#">show isis statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                   |

## Sample Output

### clear isis statistics

The following sample output displays IS-IS statistics before and after the **clear isis statistics** command is entered:

```
user@host> show isis statistics
IS-IS statistics for merino:
```

| PDU type | Received | Processed | Drops | Sent   | Rexmit |
|----------|----------|-----------|-------|--------|--------|
| LSP      | 12793    | 12793     | 0     | 8666   | 719    |
| IIH      | 116751   | 116751    | 0     | 118834 | 0      |
| CSNP     | 203956   | 203956    | 0     | 204080 | 0      |
| PSNP     | 7356     | 7350      | 6     | 8635   | 0      |
| Unknown  | 0        | 0         | 0     | 0      | 0      |
| Totals   | 340856   | 340850    | 6     | 340215 | 719    |

Total packets received: 340856 Sent: 340934

SNP queue length: 0 Drops: 0  
LSP queue length: 0 Drops: 0

SPF runs: 1064  
Fragments rebuilt: 1087  
LSP regenerations: 436  
Purges initiated: 0

user@host> clear isis statistics

user@host> show isis statistics  
IS-IS statistics for merino:

| PDU type | Received | Processed | Drops | Sent | Rexmit |
|----------|----------|-----------|-------|------|--------|
| LSP      | 0        | 0         | 0     | 0    | 0      |
| IIH      | 3        | 3         | 0     | 3    | 0      |
| CSNP     | 2        | 2         | 0     | 4    | 0      |
| PSNP     | 0        | 0         | 0     | 0    | 0      |
| Unknown  | 0        | 0         | 0     | 0    | 0      |
| Totals   | 5        | 5         | 0     | 7    | 0      |

Total packets received: 5 Sent: 7

SNP queue length: 0 Drops: 0  
LSP queue length: 0 Drops: 0

SPF runs: 0  
Fragments rebuilt: 0  
LSP regenerations: 0  
Purges initiated: 0

## show isis adjacency

---

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3150</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3150</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                                     | <pre>show isis adjacency &lt;system-id&gt; &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis adjacency &lt;system-id&gt; &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>                                | Display information about IS-IS neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                                    | <p><b>none</b>—Display standard information about IS-IS neighbors for all routing instances.</p> <p><b><i>system id</i></b>—(Optional) Display information about IS-IS neighbors for the specified intermediate system.</p> <p><b><i>brief   detail   extensive</i></b>—(Optional) Display standard information about IS-IS neighbors with the specified level of output.</p> <p><b><i>instance instance-name</i></b>—(Optional) Display information about IS-IS neighbors for the specified routing instance.</p> <p><b><i>logical-system (all   logical-system-name)</i></b>—(Optional) Display information about IS-IS neighbors for all logical systems or for a particular logical system.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">clear isis adjacency on page 3142</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                      | <a href="#">show isis adjacency on page 3152</a><br><a href="#">show isis adjacency brief on page 3152</a><br><a href="#">show isis adjacency detail on page 3153</a><br><a href="#">show isis adjacency extensive on page 3153</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>                              | <a href="#">Table 327 on page 3151</a> describes the output fields for the <b>show isis adjacency</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 327: show isis adjacency Output Fields

| Field Name                                | Field Description                                                                                                                                                                                                                               | Level of Output         |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Interface</b>                          | Interface through which the neighbor is reachable.                                                                                                                                                                                              | All levels              |
| <b>System</b>                             | System identifier ( <b>sysid</b> ), displayed as a name, if possible.                                                                                                                                                                           | <b>brief</b>            |
| <b>L or Level</b>                         | Level: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul> An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address. | All levels              |
| <b>State</b>                              | State of the adjacency: <b>Up</b> , <b>Down</b> , <b>New</b> , <b>One-way</b> , <b>Initializing</b> , or <b>Rejected</b> .                                                                                                                      | All levels              |
| <b>Hold (secs)</b>                        | Remaining hold time of the adjacency.                                                                                                                                                                                                           | <b>brief</b>            |
| <b>SNPA</b>                               | Subnetwork point of attachment (MAC address of the next hop).                                                                                                                                                                                   | <b>brief</b>            |
| <b>Expires in</b>                         | How long until the adjacency expires, in seconds.                                                                                                                                                                                               | <b>detail</b>           |
| <b>Priority</b>                           | Priority to become the designated intermediate system.                                                                                                                                                                                          | <b>detail extensive</b> |
| <b>Up/Down transitions</b>                | Count of adjacency status changes from <b>Up</b> to <b>Down</b> or from <b>Down</b> to <b>Up</b> .                                                                                                                                              | <b>detail</b>           |
| <b>Last transition</b>                    | Time of the last <b>Up/Down</b> transition.                                                                                                                                                                                                     | <b>detail</b>           |
| <b>Circuit type</b>                       | Bit mask of levels on this interface: 1=Level 1 router; 2=Level 2 router; 3=both Level 1 and Level 2 router.                                                                                                                                    | <b>detail</b>           |
| <b>Speaks</b>                             | Protocols supported by this neighbor.                                                                                                                                                                                                           | <b>detail extensive</b> |
| <b>MAC address</b>                        | MAC address of the interface.                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Topologies</b>                         | Supported topologies.                                                                                                                                                                                                                           | <b>detail extensive</b> |
| <b>Restart capable</b>                    | Whether a neighbor is capable of graceful restart: <b>Yes</b> or <b>No</b> .                                                                                                                                                                    | <b>detail extensive</b> |
| <b>Adjacency advertisement: Advertise</b> | This routing device has signaled to advertise this interface to its neighbors in their link-state PDUs.                                                                                                                                         | <b>detail extensive</b> |
| <b>Adjacency advertisement: Suppress</b>  | This neighbor has signaled not to advertise the interface in the routing device's outbound link-state PDUs.                                                                                                                                     | <b>detail extensive</b> |
| <b>IP addresses</b>                       | IP address of this neighbor.                                                                                                                                                                                                                    | <b>detail extensive</b> |

Table 327: show isis adjacency Output Fields (*continued*)

| Field Name     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Transition log | <p>List of recent transitions, including:</p> <ul style="list-style-type: none"> <li>• <b>When</b>—Time at which an IS-IS adjacency transition occurred.</li> <li>• <b>State</b>—Current state of the IS-IS adjacency (<b>up</b>, <b>down</b>, or <b>rejected</b>). <ul style="list-style-type: none"> <li>• <b>Up</b>—Adjacency is up and operational.</li> <li>• <b>Down</b>—Adjacency is down and not available.</li> <li>• <b>Rejected</b>—Adjacency has been rejected.</li> </ul> </li> <li>• <b>Event</b>—Type of transition that occurred. <ul style="list-style-type: none"> <li>• <b>Seenself</b>—Possible routing loop has been detected.</li> <li>• <b>Interface down</b>—IS-IS interface has gone down and is no longer available.</li> <li>• <b>Error</b>—Adjacency error.</li> </ul> </li> <li>• <b>Down reason</b>—Reason that an IS-IS adjacency is down: <ul style="list-style-type: none"> <li>• <b>3-Way Handshake Failed</b>—Connection establishment failed.</li> <li>• <b>Address Mismatch</b>—Address mismatch caused link failure.</li> <li>• <b>Aged Out</b>—Link expired.</li> <li>• <b>ISO Area Mismatch</b>—IS-IS area mismatch caused link failure.</li> <li>• <b>Bad Hello</b>—Unacceptable hello message caused link failure.</li> <li>• <b>BFD Session Down</b>—Bidirectional failure detection caused link failure.</li> <li>• <b>Interface Disabled</b>—IS-IS interface is disabled.</li> <li>• <b>Interface Down</b>—IS-IS interface is unavailable.</li> <li>• <b>Interface Level Disabled</b>—IS-IS level is disabled.</li> <li>• <b>Level Changed</b>—IS-IS level has changed on the adjacency.</li> <li>• <b>Level Mismatch</b>—Levels on adjacency are not compatible.</li> <li>• <b>MPLS LSP Down</b>—Label-switched path (LSP) is unavailable.</li> <li>• <b>MT Topology Changed</b>—IS-IS topology has changed.</li> <li>• <b>MT Topology Mismatch</b>—IS-IS topology is mismatched.</li> <li>• <b>Remote System ID Changed</b>—Adjacency peer system ID changed.</li> <li>• <b>Protocol Shutdown</b>—IS-IS protocol is disabled.</li> <li>• <b>CLI Command</b>—Adjacency brought down by user.</li> <li>• <b>Unknown</b>—Unknown.</li> </ul> </li> </ul> | extensive       |

## Sample Output

### show isis adjacency

```

user@host> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
at-2/3/0.0         ranier      3 Up          23

```

### show isis adjacency brief

The output for the **show isis adjacency brief** command is identical to that for the **show isis adjacency** command. For sample output, see [show isis adjacency on page 3152](#).



### show isis adjacency detail

```
user@host> show isis adjacency detail
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 21 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:09 ago
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 11.1.1.2
```

### show isis adjacency extensive

```
user@host> show isis adjacency extensive
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 22 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:16 ago
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 11.1.1.2
  Transition log:
    When           State      Event           Down reason
    Wed Nov  8 21:24:25  Up        Seenself
```

## show isis authentication

|                                                   |                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3154</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3154</a>                                                                                                                                                                                                                                         |
| <b>Syntax</b>                                     | <pre>show isis authentication &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                    |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis authentication &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>                        | <p>Command introduced in Junos OS Release 7.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for hitless authentication key rollover introduced in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                        |
| <b>Description</b>                                | Display information about IS-IS authentication.                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                    | <p><b>none</b>—Display information about IS-IS authentication.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display IS-IS authentication for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>                      | <a href="#">show isis authentication on page 3155</a><br><a href="#">show isis authentication (With Hitless Authentication Key Rollover Configured) on page 3155</a>                                                                                                                                                                                   |
| <b>Output Fields</b>                              | <p><a href="#">Table 328 on page 3154</a> describes the output fields for the <b>show isis authentication</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                     |

**Table 328: show isis authentication Output Fields**

| Field Name       | Field Description                                                                                                                                          |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b> | Interface name.                                                                                                                                            |
| <b>Level</b>     | IS-IS level.                                                                                                                                               |
| <b>IIH Auth</b>  | <p>IS-IS Hello (IIH) packet authentication type.</p> <p>Displays the name of the active keychain if hitless authentication key rollover is configured.</p> |
| <b>CSN Auth</b>  | Complete sequence number authentication type.                                                                                                              |

Table 328: show isis authentication Output Fields (*continued*)

| Field Name            | Field Description                            |
|-----------------------|----------------------------------------------|
| PSN Auth              | Partial sequence number authentication type. |
| L1 LSP Authentication | Layer 1 link-state PDU authentication type.  |
| L2 LSP Authentication | Layer 2 link-state PDU authentication type.  |

## Sample Output

### show isis authentication

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
at-2/3/0.0         1      Simple    Simple    Simple
                   2      MD5       MD5       MD5

L1 LSP Authentication: Simple
L2 LSP Authentication: MD5

```

### show isis authentication (With Hitless Authentication Key Rollover Configured)

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
so-0/1/3.0         2      hakrhello MD5       MD5

L2 LSP Authentication: MD5

```

## show isis backup coverage

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                     | <pre>show isis backup coverage &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis backup coverage &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                        | <p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                 |
| <b>Description</b>                                | Display information about the level of backup coverage available.                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                                    | <p><b>none</b>—Display information about the level of backup coverage available for all the nodes and prefixes in the network.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the level of backup coverage for a specific IS-IS routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Link and Node Protection for IS-IS Routes</i></li> <li>• <a href="#">show isis backup label-switched-path on page 3158</a></li> </ul>                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>                      | <a href="#">show isis backup coverage on page 3157</a>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Output Fields</b>                              | <p><a href="#">Table 329 on page 3156</a> lists the output fields for the <b>show isis backup coverage</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                     |

**Table 329: show isis backup coverage Output Fields**

| Field Name      | Field Description                                                                                            |
|-----------------|--------------------------------------------------------------------------------------------------------------|
| <b>Topology</b> | Type of topology or address family: <b>IPV4 Unicast</b> or <b>IPV6 Unicast</b> .                             |
| <b>Level</b>    | IS-IS level: <ul style="list-style-type: none"> <li>• 1—Level 1</li> <li>• 2—Level 2</li> </ul>              |
| <b>Node</b>     | By topology, the percentage of all routes configured on the node that are protected through backup coverage. |

Table 329: show isis backup coverage Output Fields (*continued*)

| Field Name | Field Description                                                                                      |
|------------|--------------------------------------------------------------------------------------------------------|
| IPv4       | Percentage of IPv4 unicast routes that are protected through backup coverage.                          |
| IPv6       | Percentage of IPv6 unicast routes that are protected through backup coverage.                          |
| CLNS       | Percentage of Connectionless Network Service (CLNS) routes that are protected through backup coverage. |

## Sample Output

show isis backup coverage

```

user@host> show isis backup coverage
Backup Coverage:
  Topology   Level  Node   IPv4   IPv6   CLNS
  IPV4 Unicast    2  28.57%  22.22%  0.00%  0.00%
  IPV6 Unicast    2   0.00%  0.00%  0.00%  0.00%

```

## show isis backup label-switched-path

|                                                   |                                                                                                                                                                                                                                                             |
|---------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                     | <b>show isis backup label-switched-path</b><br><b>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</b>                                                                                                                                             |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <b>show isis backup label-switched-path</b>                                                                                                                                                                                                                 |
| <b>Release Information</b>                        | Command introduced in Junos OS Release 9.5.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                        |
| <b>Description</b>                                | Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes.                                                                                                                                                    |
| <b>Options</b>                                    | <b>none</b> —Display information about MPLS LSPs designated as backup routes for IS-IS routes.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Link and Node Protection for IS-IS Routes</i></li> <li>• <a href="#">show isis backup coverage on page 3156</a></li> </ul>                                                                 |
| <b>List of Sample Output</b>                      | <a href="#">show isis backup label-switched-path on page 3159</a>                                                                                                                                                                                           |
| <b>Output Fields</b>                              | <a href="#">Table 330 on page 3158</a> lists the output fields for the <b>show isis backup label-switched-path</b> command. Output fields are listed in the approximate order in which they appear.                                                         |

**Table 330: show isis backup label-switched-path Output Fields**

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Backup MPLS LSPs</b> | List of MPLS LSPs designated as backup paths for IS-IS routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Egress</b>           | IP address of the egress routing device for the LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Status</b>           | State of the LSP: <ul style="list-style-type: none"> <li>• <b>Up</b>—The routing device can detect RSVP hello messages from the neighbor.</li> <li>• <b>Down</b>—The routing device has received one of the following indications:               <ul style="list-style-type: none"> <li>• Communication failure from the neighbor.</li> <li>• Communication from IGP that the neighbor is unavailable.</li> <li>• Change in the sequence numbers in the RSVP hello messages sent by the neighbor.</li> </ul> </li> <li>• <b>Deleted</b>—LSP is no longer available as a backup path.</li> </ul> |

Table 330: show isis backup label-switched-path Output Fields (*continued*)

| Field Name  | Field Description                                                                                                        |
|-------------|--------------------------------------------------------------------------------------------------------------------------|
| Last change | Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <i>hh:mm:ss</i> . |
| TE-metric   | Configured traffic engineering metric.                                                                                   |
| Metric      | Configured metric.                                                                                                       |

## Sample Output

### show isis backup label-switched-path

```
user@host> show isis backup label-switched-path
Backup MPLS LSPs:
f-to-g, Egress: 192.168.1.4, Status: up, Last change: 06:12:03
TE-metric: 9, Metric: 0
```

## show isis backup spf results

---

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                      | <pre>show isis backup spf results &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;no-coverage&gt; &lt;topology (ipv4-unicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Syntax (EX Series Switches) | <pre>show isis backup spf results &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;no-coverage&gt; &lt;topology (ipv4-unicast   unicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Release Information         | Command introduced in Junos OS Release 9.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Description                 | Display information about IS-IS shortest-path-first (SPF) calculations for backup paths.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Options                     | <p><b>none</b>—Display information about IS-IS SPF calculations for all backup paths for all destination nodes.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display SPF calculations for backup paths for the specified routing instance.</p> <p><b>level (1   2)</b>—(Optional) Display SPF calculations for the backup paths for the specified IS-IS level.</p> <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Display SPF calculations for the backup paths for all logical systems or on a particular logical system.</p> <p><b>no-coverage</b>—(Optional) Display SPF calculations only for destinations that do not have backup coverage.</p> <p><b>topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)</b>—(Optional) Display SPF calculations for backup paths for the specified topology only.</p> |
| Required Privilege Level    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Related Documentation       | <ul style="list-style-type: none"><li>• <i>Example: Configuring Link and Node Protection for IS-IS Routes</i></li><li>• <a href="#">show isis backup coverage on page 3156</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| List of Sample Output       | <a href="#">show isis backup spf results on page 3161</a><br><a href="#">show isis backup spf results no-coverage on page 3162</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Output Fields               | <a href="#">Table 331 on page 3161</a> lists the output fields for the <b>show isis backup spf results</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |



Table 331: show isis backup spf results Output Fields

| Field Name              | Field Description                                                                                  |
|-------------------------|----------------------------------------------------------------------------------------------------|
| <i>node-name</i>        | Name of the destination node.                                                                      |
| <b>Address</b>          | Address of the destination node.                                                                   |
| <b>Primary next-hop</b> | Interface and name of the node of the primary next hop to reach the destination.                   |
| <b>Root</b>             | Name of the next-hop neighbor.                                                                     |
| <b>Metric</b>           | Metric to the node.                                                                                |
| <b>Eligible</b>         | Indicates that the next-hop neighbor has been designated as a backup path to the destination node. |
| <b>Backup next-hop</b>  | Name of the interface of the backup next hop.                                                      |
| <b>SNPA</b>             | Subnetwork point of attachment (MAC address of the next hop).                                      |
| <b>LSP</b>              | Name of the MPLS label-switched path (LSP) designated as a backup path.                            |
| <b>Not eligible</b>     | Indicates that the next-hop neighbor cannot function as a backup path to the destination.          |
| <b>Reason</b>           | Describes why the next-hop neighbor is designated as <b>Not eligible</b> as a backup path.         |

## Sample Output

### show isis backup spf results

```

user@host> show isis backup spf results
IS-IS level 1 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0

```

```

        track-item: pro-bng3-j.00-00
        track-item: pro-bng3-i.00-00
        Not eligible, IPV4, Reason: Path loops
        Not eligible, IPV6, Reason: Path loops
    2 nodes

IS-IS level 2 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
2 nodes

```

### show isis backup spf results no-coverage

```

user@host> show isis backup spf results no-coverage
IS-IS level 1 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
2 nodes

IS-IS level 2 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21

```

```
Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
2 nodes
```

## show isis database

---

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                             | <a href="#">Syntax on page 3164</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3164</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Syntax                                     | <pre>show isis database &lt;system-id&gt; &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Syntax (EX Series Switches and QFX Series) | <pre>show isis database &lt;system-id&gt; &lt;brief   detail   extensive&gt; &lt;level (1   2)&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Release Information                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Description                                | Display the entries in the IS-IS link-state database, which contains data about PDU packets.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Options                                    | <p><b>none</b>—Display standard information about IS-IS link-state database entries for all routing instances.</p> <p><b><i>system id</i></b>—(Optional) Display IS-IS link-state database entries for the specified intermediate system.</p> <p><b><i>brief   detail   extensive</i></b>—(Optional) Display the specified level of output.</p> <p><b><i>instance instance-name</i></b>—(Optional) Display IS-IS link-state database entries for the specified routing instance.</p> <p><b><i>level (1   2)</i></b>—(Optional) Display IS-IS link-state database entries for the specified IS-IS level.</p> <p><b><i>logical-system (all   logical-system-name)</i></b>—(Optional) Display standard information about IS-IS link-state database entries for all logical systems or for a particular logical system.</p> |
| Required Privilege Level                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Related Documentation                      | <ul style="list-style-type: none"><li>• <a href="#">clear isis database on page 3144</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| List of Sample Output                      | <a href="#">show isis database on page 3166</a><br><a href="#">show isis database brief on page 3167</a><br><a href="#">show isis database detail on page 3167</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

[show isis database extensive on page 3167](#)

**Output Fields** [Table 332 on page 3165](#) describes the output fields for the **show isis database** command. Output fields are listed in the approximate order in which they appear. Fields that contain internal IS-IS information useful only in troubleshooting obscure problems are not described in the table. For more details about these fields, contact your customer support representative.

**Table 332: show isis database Output Fields**

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                            | Level of Output         |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Interface name</b> | Name of the interface on which the link-state PDU has been received; always <b>IS-IS</b> for this command.                                                                                                                                                                                                                                                                                   | All levels              |
| level                 | Level of intermediate system: <ul style="list-style-type: none"> <li>• <b>1</b>—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system.</li> <li>• <b>2</b>—Intermediate system routes between areas and toward other ASs.</li> </ul>                                                                                         | All levels              |
| LSP ID                | Link-state PDU identifier.                                                                                                                                                                                                                                                                                                                                                                   | All levels              |
| Sequence              | Sequence number of the link-state PDU.                                                                                                                                                                                                                                                                                                                                                       | All levels              |
| Checksum              | Checksum value of the link-state PDU.                                                                                                                                                                                                                                                                                                                                                        | All levels              |
| Lifetime (secs)       | Remaining lifetime of the link-state PDU, in seconds.                                                                                                                                                                                                                                                                                                                                        | All levels              |
| Attributes            | Attributes of the specified database: <b>L1</b> , <b>L2</b> , <b>Overload</b> , or <b>Attached</b> (L1 only).                                                                                                                                                                                                                                                                                | none <b>brief</b>       |
| # LSPs                | Total number of link-state PDUs in the specified link-state database.                                                                                                                                                                                                                                                                                                                        | none <b>brief</b>       |
| IP prefix             | Prefix advertised by this link-state PDU.                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b> |
| IS neighbor           | IS-IS neighbor of the advertising system.                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b> |
| ES neighbor           | (J Series routers only) An ES-IS neighbor of the advertising system.                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| IP prefix             | IPv4 prefix advertised by this link-state PDU.                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b> |
| V6 prefix             | IPv6 prefix advertised by this link-state PDU.                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b> |
| Metric                | Metric of the prefix or neighbor.                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| Header                | <ul style="list-style-type: none"> <li>• <b>LSP ID</b>—Link state PDU identifier of the header.</li> <li>• <b>Length</b>—Header length.</li> <li>• <b>Allocated Length</b>—Amount of length available for the header.</li> <li>• <b>Router ID</b>—Address of the local routing device.</li> <li>• <b>Remaining Lifetime</b>—Remaining lifetime of the link-state PDU, in seconds.</li> </ul> | <b>extensive</b>        |

Table 332: show isis database Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Packet     | <ul style="list-style-type: none"> <li>• <b>LSP ID</b>—The identifier for the link-state PDU.</li> <li>• <b>Length</b>—Packet length.</li> <li>• <b>Lifetime</b>—Remaining lifetime, in seconds.</li> <li>• <b>Checksum</b>—The checksum of the link-state PDU.</li> <li>• <b>Sequence</b>—The sequence number of the link-state PDU. Every time the link-state PDU is updated, this number increments.</li> <li>• <b>Attributes</b>—Packet attributes.</li> <li>• <b>NLPID</b>—Network layer protocol identifier.</li> <li>• <b>Fixed length</b>—Specifies the set length for the packet.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                            | extensive       |
| TLVs       | <ul style="list-style-type: none"> <li>• <b>Area Address</b>—Area addresses that the routing device can reach.</li> <li>• <b>Speaks</b>—Supported routing protocols.</li> <li>• <b>IP router id</b>—ID of the routing device (usually the IP address).</li> <li>• <b>IP address</b>—IPv4 address.</li> <li>• <b>Hostname</b>—Assigned name of the routing device.</li> <li>• <b>IP prefix</b>—IP prefix of the routing device.</li> <li>• <b>Metric</b>—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device.</li> <li>• <b>IP extended prefix</b>—Extended IP prefix of the routing device.</li> <li>• <b>IS neighbor</b>—Directly attached neighbor's name and metric.</li> <li>• <b>IS extended neighbor</b>—Directly attached neighbor's name, metric, IP address, local interface index, and remote interface index.</li> </ul> <p>The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.</p> | extensive       |

## Sample Output

### show isis database

```

user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x3     0x3167    1057 L1 L2
camaro.00-00          0x5     0x770e    1091 L1 L2
ranier.00-00          0x4     0xaa95    1091 L1 L2
glacier.00-00         0x4     0x206f    1089 L1 L2
glacier.02-00         0x1     0xd141    1089 L1 L2
badlands.00-00        0x3     0x87a2    1093 L1 L2
  6 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x6     0x8d6b    1096 L1 L2
camaro.00-00          0x9     0x877b    1101 L1 L2
ranier.00-00          0x8     0x855d    1103 L1 L2
glacier.00-00         0x7     0xf892    1098 L1 L2
glacier.02-00         0x1     0xd141    1089 L1 L2
badlands.00-00        0x6     0x562     1105 L1 L2
  6 LSPs

```

### show isis database brief

The output for the **show isis database brief** command is identical to that for the **show isis database** command. For sample output, see [show isis database on page 3166](#).

### show isis database detail

```
user@host> show isis database logical-system CE3 sisira.00-00 detail
```

IS-IS level 1 link-state database:

```
sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 975 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  ES neighbor: 0015.0015.0015           Metric:      10 Down
  ES neighbor: 0025.0025.0025           Metric:      10 Down
  ES neighbor: 0030.0030.0030           Metric:      10 Down
  ES neighbor: 0040.0040.0040           Metric:      10 Down
  ES neighbor: sisira                     Metric:       0
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
```

IS-IS level 2 link-state database:

```
sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 993 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 50.50.50.50/32              Metric:      10 Internal Up
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
  Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
  Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
  Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
  Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
  Metric:       0 Internal Up
```

### show isis database extensive

```
user@host> show isis database extensive
```

IS-IS level 1 link-state database:

```
Router-A.00-00 Sequence: 0x1, Checksum: 0xf75c, Lifetime: 1116 secs
```

IP prefix: 192.168.0.1/32                      Metric:              0 Internal Up

Header: LSP ID: Router-A.00-00, Length: 85 bytes  
Allocated length: 1492 bytes, Router ID: 192.168.0.1  
Remaining lifetime: 1116 secs, Level: 1, Interface: 0  
Estimated free bytes: 1353, Actual free bytes: 1407  
Aging timer expires in: 1116 secs  
Protocols: IP, IPv6

Packet: LSP ID: Router-A.00-00, Length: 85 bytes, Lifetime : 1200 secs  
Checksum: 0xf75c, Sequence: 0x1, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 18, Packet version: 1, Max area: 0

TLVs:

Area address: 49.0002 (3)  
LSP Buffer Size: 1492  
Speaks: IP  
Speaks: IPV6  
IP router id: 192.168.0.1  
IP address: 192.168.0.1  
Hostname: Router-A  
IP prefix: 192.168.0.1/32, Internal, Metric: default 0, Up  
IP extended prefix: 192.168.0.1/32 metric 0 up  
No queued transmissions

IS-IS level 2 link-state database:

Router-A.00-00 Sequence: 0x5, Checksum: 0x3196, Lifetime: 1158 secs  
IS neighbor: Router-B.02                      Metric:              10  
Two-way fragment: Router-B.02-00, Two-way first fragment: Router-B.02-00  
IS neighbor: Router-E.02                      Metric:              10  
Two-way fragment: Router-E.02-00, Two-way first fragment: Router-E.02-00  
IP prefix: 10.0.0.0/30                      Metric:              10 Internal Up  
IP prefix: 10.0.0.4/30                      Metric:              10 Internal Up  
IP prefix: 192.168.0.1/32                      Metric:              0 Internal Up

Header: LSP ID: Router-A.00-00, Length: 208 bytes  
Allocated length: 1492 bytes, Router ID: 192.168.0.1  
Remaining lifetime: 1158 secs, Level: 2, Interface: 0  
Estimated free bytes: 1233, Actual free bytes: 1284  
Aging timer expires in: 1158 secs  
Protocols: IP, IPv6

Packet: LSP ID: Router-A.00-00, Length: 208 bytes, Lifetime : 1198 secs  
Checksum: 0x3196, Sequence: 0x5, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

TLVs:

Area address: 49.0002 (3)  
LSP Buffer Size: 1492  
Speaks: IP  
Speaks: IPV6  
IP router id: 192.168.0.1  
IP address: 192.168.0.1  
Hostname: Router-A  
IP prefix: 192.168.0.1/32, Internal, Metric: default 0, Up  
IP prefix: 10.0.0.4/30, Internal, Metric: default 10, Up  
IP prefix: 10.0.0.0/30, Internal, Metric: default 10, Up  
IP extended prefix: 192.168.0.1/32 metric 0 up



```

IP extended prefix: 10.0.0.4/30 metric 10 up
IP extended prefix: 10.0.0.0/30 metric 10 up
IS neighbor: Router-E.02, Internal, Metric: default 10
IS neighbor: Router-B.02, Internal, Metric: default 10
IS extended neighbor: Router-E.02, Metric: default 10
  IP address: 10.0.0.1
    Local interface index: 101, Remote interface index: 0
IS extended neighbor: Router-B.02, Metric: default 10
  IP address: 10.0.0.5
    Local interface index: 102, Remote interface index: 0
No queued transmissions

Router-B.00-00 Sequence: 0x5, Checksum: 0xf8f, Lifetime: 1183 secs
  IS neighbor: Router-B.02                      Metric: 10
    Two-way fragment: Router-B.02-00, Two-way first fragment: Router-B.02-00
  IS neighbor: Router-C.02                      Metric: 10
    Two-way fragment: Router-C.02-00, Two-way first fragment: Router-C.02-00
IP prefix: 10.0.0.4/30                          Metric: 10 Internal Up
IP prefix: 10.0.0.8/30                          Metric: 10 Internal Up
IP prefix: 192.168.0.2/32                      Metric: 0 Internal Up

Header: LSP ID: Router-B.00-00, Length: 208 bytes
  Allocated length: 284 bytes, Router ID: 192.168.0.2
  Remaining lifetime: 1183 secs, Level: 2, Interface: 102
  Estimated free bytes: 114, Actual free bytes: 76
  Aging timer expires in: 1183 secs
  Protocols: IP, IPv6

Packet: LSP ID: Router-B.00-00, Length: 208 bytes, Lifetime : 1196 secs
  Checksum: 0xf8f, Sequence: 0x5, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

TLVs:
  Area address: 49.0002 (3)
  LSP Buffer Size: 1492
  Speaks: IP
  Speaks: IPV6
  IP router id: 192.168.0.2
  IP address: 192.168.0.2
  Hostname: Router-B
  IP prefix: 192.168.0.2/32, Internal, Metric: default 0, Up
  IP prefix: 10.0.0.4/30, Internal, Metric: default 10, Up
  IP prefix: 10.0.0.8/30, Internal, Metric: default 10, Up
  IP extended prefix: 192.168.0.2/32 metric 0 up
  IP extended prefix: 10.0.0.4/30 metric 10 up
  IP extended prefix: 10.0.0.8/30 metric 10 up
  IS neighbor: Router-B.02, Internal, Metric: default 10
  IS neighbor: Router-C.02, Internal, Metric: default 10
  IS extended neighbor: Router-B.02, Metric: default 10
    IP address: 10.0.0.6
      Local interface index: 108, Remote interface index: 0
  IS extended neighbor: Router-C.02, Metric: default 10
    IP address: 10.0.0.9
      Local interface index: 109, Remote interface index: 0
No queued transmissions

Router-B.02-00 Sequence: 0x1, Checksum: 0x3c7c, Lifetime: 1156 secs
  IS neighbor: Router-A.00                      Metric: 0
    Two-way fragment: Router-A.00-00, Two-way first fragment: Router-A.00-00
  IS neighbor: Router-B.00                      Metric: 0

```

Two-way fragment: Router-B.00-00, Two-way first fragment: Router-B.00-00

Header: LSP ID: Router-B.02-00, Length: 76 bytes  
Allocated length: 284 bytes, Router ID: 0.0.0.0  
Remaining lifetime: 1156 secs, Level: 2, Interface: 102  
Estimated free bytes: 208, Actual free bytes: 208  
Aging timer expires in: 1156 secs

Packet: LSP ID: Router-B.02-00, Length: 76 bytes, Lifetime : 1196 secs  
Checksum: 0x3c7c, Sequence: 0x1, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

TLVs:  
IS neighbor: Router-B.00, Internal, Metric: default 0  
IS neighbor: Router-A.00, Internal, Metric: default 0  
IS extended neighbor: Router-B.00, Metric: default 0  
IS extended neighbor: Router-A.00, Metric: default 0  
No queued transmissions

Router-C.00-00 Sequence: 0x5, Checksum: 0x255b, Lifetime: 1182 secs  
IS neighbor: Router-C.02 Metric: 10  
Two-way fragment: Router-C.02-00, Two-way first fragment: Router-C.02-00  
IS neighbor: Router-D.03 Metric: 10  
Two-way fragment: Router-D.03-00, Two-way first fragment: Router-D.03-00  
IP prefix: 10.0.0.8/30 Metric: 10 Internal Up  
IP prefix: 10.0.0.12/30 Metric: 10 Internal Up  
IP prefix: 192.168.0.3/32 Metric: 0 Internal Up

Header: LSP ID: Router-C.00-00, Length: 208 bytes  
Allocated length: 284 bytes, Router ID: 192.168.0.3  
Remaining lifetime: 1182 secs, Level: 2, Interface: 102  
Estimated free bytes: 114, Actual free bytes: 76  
Aging timer expires in: 1182 secs  
Protocols: IP, IPv6

Packet: LSP ID: Router-C.00-00, Length: 208 bytes, Lifetime : 1196 secs  
Checksum: 0x255b, Sequence: 0x5, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

TLVs:  
Area address: 49.0002 (3)  
LSP Buffer Size: 1492  
Speaks: IP  
Speaks: IPV6  
IP router id: 192.168.0.3  
IP address: 192.168.0.3  
Hostname: Router-C  
IP prefix: 192.168.0.3/32, Internal, Metric: default 0, Up  
IP prefix: 10.0.0.8/30, Internal, Metric: default 10, Up  
IP prefix: 10.0.0.12/30, Internal, Metric: default 10, Up  
IP extended prefix: 192.168.0.3/32 metric 0 up  
IP extended prefix: 10.0.0.8/30 metric 10 up  
IP extended prefix: 10.0.0.12/30 metric 10 up  
IS neighbor: Router-C.02, Internal, Metric: default 10  
IS neighbor: Router-D.03, Internal, Metric: default 10  
IS extended neighbor: Router-C.02, Metric: default 10  
IP address: 10.0.0.10  
Local interface index: 105, Remote interface index: 0  
IS extended neighbor: Router-D.03, Metric: default 10

IP address: 10.0.0.13  
 Local interface index: 106, Remote interface index: 0  
 No queued transmissions

Router-C.02-00 Sequence: 0x1, Checksum: 0xaa09, Lifetime: 1181 secs  
 IS neighbor: Router-B.00 Metric: 0  
 Two-way fragment: Router-B.00-00, Two-way first fragment: Router-B.00-00  
 IS neighbor: Router-C.00 Metric: 0  
 Two-way fragment: Router-C.00-00, Two-way first fragment: Router-C.00-00

Header: LSP ID: Router-C.02-00, Length: 76 bytes  
 Allocated length: 284 bytes, Router ID: 0.0.0.0  
 Remaining lifetime: 1181 secs, Level: 2, Interface: 102  
 Estimated free bytes: 208, Actual free bytes: 208  
 Aging timer expires in: 1181 secs

Packet: LSP ID: Router-C.02-00, Length: 76 bytes, Lifetime : 1194 secs  
 Checksum: 0xaa09, Sequence: 0x1, Attributes: 0x3 <L1 L2>  
 NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
 Packet type: 20, Packet version: 1, Max area: 0

TLVs:  
 IS neighbor: Router-C.00, Internal, Metric: default 0  
 IS neighbor: Router-B.00, Internal, Metric: default 0  
 IS extended neighbor: Router-C.00, Metric: default 0  
 IS extended neighbor: Router-B.00, Metric: default 0  
 No queued transmissions

Router-D.00-00 Sequence: 0x4, Checksum: 0x8ab7, Lifetime: 1180 secs  
 IS neighbor: Router-D.02 Metric: 10  
 Two-way fragment: Router-D.02-00, Two-way first fragment: Router-D.02-00  
 IS neighbor: Router-D.03 Metric: 10  
 Two-way fragment: Router-D.03-00, Two-way first fragment: Router-D.03-00  
 IP prefix: 10.0.0.12/30 Metric: 10 Internal Up  
 IP prefix: 10.0.0.20/30 Metric: 10 Internal Up  
 IP prefix: 192.168.0.4/32 Metric: 0 Internal Up

Header: LSP ID: Router-D.00-00, Length: 208 bytes  
 Allocated length: 284 bytes, Router ID: 192.168.0.4  
 Remaining lifetime: 1180 secs, Level: 2, Interface: 102  
 Estimated free bytes: 114, Actual free bytes: 76  
 Aging timer expires in: 1180 secs  
 Protocols: IP, IPv6

Packet: LSP ID: Router-D.00-00, Length: 208 bytes, Lifetime : 1192 secs  
 Checksum: 0x8ab7, Sequence: 0x4, Attributes: 0x3 <L1 L2>  
 NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
 Packet type: 20, Packet version: 1, Max area: 0

TLVs:  
 Area address: 49.0002 (3)  
 LSP Buffer Size: 1492  
 Speaks: IP  
 Speaks: IPV6  
 IP router id: 192.168.0.4  
 IP address: 192.168.0.4  
 Hostname: Router-D  
 IP prefix: 192.168.0.4/32, Internal, Metric: default 0, Up  
 IP prefix: 10.0.0.12/30, Internal, Metric: default 10, Up  
 IP prefix: 10.0.0.20/30, Internal, Metric: default 10, Up  
 IP extended prefix: 192.168.0.4/32 metric 0 up

```
IP extended prefix: 10.0.0.12/30 metric 10 up
IP extended prefix: 10.0.0.20/30 metric 10 up
IS neighbor: Router-D.02, Internal, Metric: default 10
IS neighbor: Router-D.03, Internal, Metric: default 10
IS extended neighbor: Router-D.02, Metric: default 10
  IP address: 10.0.0.22
  Local interface index: 115, Remote interface index: 0
IS extended neighbor: Router-D.03, Metric: default 10
  IP address: 10.0.0.14
  Local interface index: 114, Remote interface index: 0
No queued transmissions

Router-D.02-00 Sequence: 0x1, Checksum: 0xebbc, Lifetime: 1128 secs
IS neighbor: Router-D.00 Metric: 0
  Two-way fragment: Router-D.00-00, Two-way first fragment: Router-D.00-00
IS neighbor: Router-F.00 Metric: 0
  Two-way fragment: Router-F.00-00, Two-way first fragment: Router-F.00-00

Header: LSP ID: Router-D.02-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1128 secs, Level: 2, Interface: 101
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1128 secs

Packet: LSP ID: Router-D.02-00, Length: 76 bytes, Lifetime : 1160 secs
  Checksum: 0xebbc, Sequence: 0x1, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

TLVs:
  IS neighbor: Router-D.00, Internal, Metric: default 0
  IS neighbor: Router-F.00, Internal, Metric: default 0
  IS extended neighbor: Router-D.00, Metric: default 0
  IS extended neighbor: Router-F.00, Metric: default 0
No queued transmissions

Router-D.03-00 Sequence: 0x1, Checksum: 0x129b, Lifetime: 1180 secs
IS neighbor: Router-C.00 Metric: 0
  Two-way fragment: Router-C.00-00, Two-way first fragment: Router-C.00-00
IS neighbor: Router-D.00 Metric: 0
  Two-way fragment: Router-D.00-00, Two-way first fragment: Router-D.00-00

Header: LSP ID: Router-D.03-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1180 secs, Level: 2, Interface: 101
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1180 secs

Packet: LSP ID: Router-D.03-00, Length: 76 bytes, Lifetime : 1192 secs
  Checksum: 0x129b, Sequence: 0x1, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

TLVs:
  IS neighbor: Router-D.00, Internal, Metric: default 0
  IS neighbor: Router-C.00, Internal, Metric: default 0
  IS extended neighbor: Router-D.00, Metric: default 0
  IS extended neighbor: Router-C.00, Metric: default 0
No queued transmissions

Router-E.00-00 Sequence: 0x4, Checksum: 0x9da9, Lifetime: 1155 secs
```

```

IS neighbor: Router-E.02                      Metric:      10
  Two-way fragment: Router-E.02-00, Two-way first fragment: Router-E.02-00
IS neighbor: Router-F.02                      Metric:      20
  Two-way fragment: Router-F.02-00, Two-way first fragment: Router-F.02-00
IP prefix: 10.0.0.0/30                        Metric:      10 Internal Up
IP prefix: 10.0.0.16/30                       Metric:      20 Internal Up
IP prefix: 192.168.0.5/32                     Metric:       0 Internal Up

```

```

Header: LSP ID: Router-E.00-00, Length: 208 bytes
  Allocated length: 284 bytes, Router ID: 192.168.0.5
  Remaining lifetime: 1155 secs, Level: 2, Interface: 101
  Estimated free bytes: 114, Actual free bytes: 76
  Aging timer expires in: 1155 secs
  Protocols: IP, IPv6

```

```

Packet: LSP ID: Router-E.00-00, Length: 208 bytes, Lifetime : 1185 secs
  Checksum: 0x9da9, Sequence: 0x4, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

```

#### TLVs:

```

Area address: 49.0002 (3)
LSP Buffer Size: 1492
Speaks: IP
Speaks: IPV6
IP router id: 192.168.0.5
IP address: 192.168.0.5
Hostname: Router-E
IP prefix: 192.168.0.5/32, Internal, Metric: default 0, Up
IP prefix: 10.0.0.16/30, Internal, Metric: default 20, Up
IP prefix: 10.0.0.0/30, Internal, Metric: default 10, Up
IP extended prefix: 192.168.0.5/32 metric 0 up
IP extended prefix: 10.0.0.16/30 metric 20 up
IP extended prefix: 10.0.0.0/30 metric 10 up
IS neighbor: Router-E.02, Internal, Metric: default 10
IS neighbor: Router-F.02, Internal, Metric: default 20
IS extended neighbor: Router-E.02, Metric: default 10
  IP address: 10.0.0.2
  Local interface index: 112, Remote interface index: 0
IS extended neighbor: Router-F.02, Metric: default 20
  IP address: 10.0.0.17
  Local interface index: 111, Remote interface index: 0
No queued transmissions

```

```

Router-E.02-00 Sequence: 0x1, Checksum: 0xb4fa, Lifetime: 1130 secs
IS neighbor: Router-A.00                      Metric:       0
  Two-way fragment: Router-A.00-00, Two-way first fragment: Router-A.00-00
IS neighbor: Router-E.00                      Metric:       0
  Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00

```

```

Header: LSP ID: Router-E.02-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1130 secs, Level: 2, Interface: 101
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1130 secs

```

```

Packet: LSP ID: Router-E.02-00, Length: 76 bytes, Lifetime : 1161 secs
  Checksum: 0xb4fa, Sequence: 0x1, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

```

## TLVs:

IS neighbor: Router-E.00, Internal, Metric: default 0  
IS neighbor: Router-A.00, Internal, Metric: default 0  
IS extended neighbor: Router-E.00, Metric: default 0  
IS extended neighbor: Router-A.00, Metric: default 0

No queued transmissions

Router-F.00-00 Sequence: 0x5, Checksum: 0x94bd, Lifetime: 1153 secs  
IS neighbor: Router-D.02 Metric: 10  
Two-way fragment: Router-D.02-00, Two-way first fragment: Router-D.02-00  
IS neighbor: Router-F.02 Metric: 10  
Two-way fragment: Router-F.02-00, Two-way first fragment: Router-F.02-00  
IP prefix: 10.0.0.16/30 Metric: 10 Internal Up  
IP prefix: 10.0.0.20/30 Metric: 10 Internal Up  
IP prefix: 192.168.0.6/32 Metric: 0 Internal Up

Header: LSP ID: Router-F.00-00, Length: 208 bytes  
Allocated length: 284 bytes, Router ID: 192.168.0.6  
Remaining lifetime: 1153 secs, Level: 2, Interface: 101  
Estimated free bytes: 76, Actual free bytes: 76  
Aging timer expires in: 1153 secs  
Protocols: IP, IPv6

Packet: LSP ID: Router-F.00-00, Length: 208 bytes, Lifetime : 1183 secs  
Checksum: 0x94bd, Sequence: 0x5, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

## TLVs:

Area address: 49.0002 (3)  
LSP Buffer Size: 1492  
Speaks: IP  
Speaks: IPV6  
IP router id: 192.168.0.6  
IP address: 192.168.0.6  
Hostname: Router-F  
IP prefix: 192.168.0.6/32, Internal, Metric: default 0, Up  
IP prefix: 10.0.0.16/30, Internal, Metric: default 10, Up  
IP prefix: 10.0.0.20/30, Internal, Metric: default 10, Up  
IP extended prefix: 192.168.0.6/32 metric 0 up  
IP extended prefix: 10.0.0.16/30 metric 10 up  
IP extended prefix: 10.0.0.20/30 metric 10 up  
IS neighbor: Router-D.02, Internal, Metric: default 10  
IS neighbor: Router-F.02, Internal, Metric: default 10  
IS extended neighbor: Router-D.02, Metric: default 10  
IP address: 10.0.0.21  
Local interface index: 94, Remote interface index: 0  
IS extended neighbor: Router-F.02, Metric: default 10  
IP address: 10.0.0.18  
Local interface index: 93, Remote interface index: 0

No queued transmissions

Router-F.02-00 Sequence: 0x1, Checksum: 0xf5ae, Lifetime: 1153 secs  
IS neighbor: Router-E.00 Metric: 0  
Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00  
IS neighbor: Router-F.00 Metric: 0  
Two-way fragment: Router-F.00-00, Two-way first fragment: Router-F.00-00

Header: LSP ID: Router-F.02-00, Length: 76 bytes  
Allocated length: 284 bytes, Router ID: 0.0.0.0  
Remaining lifetime: 1153 secs, Level: 2, Interface: 101

Estimated free bytes: 208, Actual free bytes: 208  
Aging timer expires in: 1153 secs

Packet: LSP ID: Router-F.02-00, Length: 76 bytes, Lifetime : 1183 secs  
Checksum: 0xf5ae, Sequence: 0x1, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

TLVs:

IS neighbor: Router-F.00, Internal, Metric: default 0  
IS neighbor: Router-E.00, Internal, Metric: default 0  
IS extended neighbor: Router-F.00, Metric: default 0  
IS extended neighbor: Router-E.00, Metric: default 0  
No queued transmissions

## show isis hostname

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3176</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3176</a>                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax</b>                                     | <pre>show isis hostname &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax (EX Series Switches and QFX Series)</b> | show isis hostname                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>                                | <p>Display IS-IS hostname database information.</p> <p>This command displays the system ID-to-name cache. The output shows if the mapping has been learned by receipt of a Hostname TLV #137 (type dynamic) configured in Junos OS with the <b>set system host-name</b> command, or a static mapping defined in Junos OS with the <b>set system static-host-mapping hostname sysid</b> command (type static). The local router always has its type set to static even if <b>static-host-mapping</b> is not configured.</p> |
| <b>Options</b>                                    | <p><b>none</b>—Display IS-IS hostname database information.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>                      | <a href="#">show isis hostname on page 3177</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>                              | <p><a href="#">Table 333 on page 3176</a> describes the output fields for the <b>show isis hostname</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                               |

**Table 333: show isis hostname Output Fields**

| Field Name       | Field Description                                                                                                                                                                                                                                                                                       |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>System Id</b> | System identifier mapped to the hostname.                                                                                                                                                                                                                                                               |
| <b>Hostname</b>  | Hostname mapped to the system identifier.                                                                                                                                                                                                                                                               |
| <b>Type</b>      | <p>Type of mapping between system identifier and hostname.</p> <ul style="list-style-type: none"> <li><b>Dynamic</b>—Hostname mapping determined as described in RFC 2763, <i>Dynamic Hostname Exchange Mechanism for IS-IS</i>.</li> <li><b>Static</b>—Hostname mapping configured by user.</li> </ul> |




## Sample Output

show isis hostname

```
user@host> show isis hostname
IS-IS hostname database:
System Id      Hostname
1921.6800.4201 isis1
1921.6800.4202 isis2
1921.6800.4203 isis3
```

|                      | Type    |
|----------------------|---------|
| 1921.6800.4201 isis1 | Dynamic |
| 1921.6800.4202 isis2 | Static  |
| 1921.6800.4203 isis3 | Dynamic |

```
show isis interface
```

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                             | <a href="#">Syntax on page 3178</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3178</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Syntax                                     | <pre>show isis interface &lt;brief   detail   extensive&gt; &lt;interface-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Syntax (EX Series Switches and QFX Series) | <pre>show isis interface &lt;brief   detail   extensive&gt; &lt;interface-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Release Information                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Description                                | Display status information about IS-IS-enabled interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|                                            | <div><p><b>NOTE:</b> If the configured metric for an IS-IS level is above 63, and the <b>wide-metrics-only</b> statement is not configured, the <b>show isis interface detail</b> command and the <b>show isis interface extensive</b> command display 63 as the metric value for that level. Configure the <b>wide-metrics-only</b> statement to generate metric values greater than 63 on a per IS-IS level basis.</p><p>The <b>show isis interface</b> command displays the configured metric value for an IS-IS level irrespective of whether is configured or not.</p></div> |
| Options                                    | <p><b>none</b>—Display standard information about all IS-IS-enabled interfaces.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>interface-name</b>—(Optional) Display information about the specified interface only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>                                                                                                                                                                                                                                 |
| Required Privilege Level                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Related Documentation                      | <ul style="list-style-type: none"><li>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| List of Sample Output                      | <a href="#">show isis interface on page 3180</a><br><a href="#">show isis interface brief on page 3181</a><br><a href="#">show isis interface detail on page 3181</a><br><a href="#">show isis interface extensive on page 3181</a>                                                                                                                                                                                                                                                                                                                                                                                                                                 |

**Output Fields** Table 334 on page 3179 describes the output fields for the **show isis interface** command. Output fields are listed in the approximate order in which they appear.

**Table 334: show isis interface Output Fields**

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                            | Level of Output  |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <i>interface-name</i>    | Name of the interface.                                                                                                                                                                                                                                                                                                                                                       | detail           |
| <b>Designated router</b> | Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.                                                                                                                                                                                                               | detail           |
| <b>Index</b>             | Interface index assigned by the Junos OS kernel.                                                                                                                                                                                                                                                                                                                             | detail           |
| <b>State</b>             | Internal implementation information.                                                                                                                                                                                                                                                                                                                                         | detail           |
| <b>Circuit id</b>        | Circuit identifier.                                                                                                                                                                                                                                                                                                                                                          | detail           |
| <b>Circuit type</b>      | Circuit type: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul>                                                                                                                                                                                                                                  | detail           |
| <b>LSP interval</b>      | Interval between link-state PDUs sent from the interface.                                                                                                                                                                                                                                                                                                                    | detail           |
| <b>CSNP interval</b>     | Interval between complete sequence number PDUs sent from the interface.                                                                                                                                                                                                                                                                                                      | detail extensive |
| <b>Sysid</b>             | System identifier.                                                                                                                                                                                                                                                                                                                                                           | detail           |
| <b>Interface</b>         | Interface through which the adjacency is made.                                                                                                                                                                                                                                                                                                                               | none brief       |
| <b>L or Level</b>        | Level: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul> <p><b>NOTE:</b> The default IS-IS level on loopback interfaces are always same as the IS-IS level configured on other IS-IS interfaces in a router. You can also configure IS-IS level on loopback interfaces per your requirement.</p> | All levels       |
| <b>CirID</b>             | Circuit identifier.                                                                                                                                                                                                                                                                                                                                                          | none brief       |
| <b>Level 1 DR</b>        | Level 1 designated intermediate system.                                                                                                                                                                                                                                                                                                                                      | none brief       |
| <b>Level 2 DR</b>        | Level 2 designated intermediate system.                                                                                                                                                                                                                                                                                                                                      | none brief       |
| <b>L1/L2 Metric</b>      | Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0.                                                                                                                                                                                                                                                                                     | none brief       |

Table 334: show isis interface Output Fields (*continued*)

| Field Name                                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output         |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Adjacency advertisement: Advertise</b> | This routing device has signaled to advertise this interface to its neighbors in their label-switched paths (LSPs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b> |
| <b>Adjacency advertisement: Suppress</b>  | This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b> |
| <b>Adjacencies</b>                        | Number of adjacencies established on this interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail</b>           |
| <b>Priority</b>                           | Priority value for this interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail</b>           |
| <b>Metric</b>                             | Metric value for this interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b>           |
| <b>Hello(s) / Hello Interval</b>          | Interface's hello interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Hold(s) / Hold Time</b>                | Interface's hold time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Designated Router</b>                  | Router responsible for sending network link-state advertisements, which describe all the routing devices attached to the network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail</b>           |
| <b>Hello padding</b>                      | Type of hello padding: <ul style="list-style-type: none"> <li>• <b>Adaptive</b>—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state TLV. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface.</li> <li>• <b>Loose</b>—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state.</li> <li>• <b>Strict</b>—Padding is performed on all interface types and for all adjacency states, and is continuous.</li> </ul> | <b>extensive</b>        |
| <b>LDP sync state</b>                     | Current LDP synchronization state: <b>in sync</b> , <b>in holddown</b> , or <b>not supported</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>        |
| <b>reason</b>                             | Reason for being in the LDP sync state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>        |
| <b>config holdtime</b>                    | Configured value of the hold timer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>        |
| <b>remaining</b>                          | If the state is not in sync and the hold time is not infinity, then this field displays the remaining hold time in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>extensive</b>        |

## Sample Output

### show isis interface

```
user@host> show isis interface
```

IS-IS interface database:

| Interface  | L | CirID | Level 1 DR     | Level 2 DR     | L1/L2 Metric |
|------------|---|-------|----------------|----------------|--------------|
| at-2/3/0.0 | 3 | 0x1   | Point to Point | Point to Point | 10/10        |
| lo0.0      | 3 | 0x1   | Passive        | Passive        | 0/0          |

### show isis interface brief

The output for the **show isis interface brief** command is identical to that for the **show isis interface** command. For sample output, see [show isis interface on page 3180](#).

### show isis interface detail

```
user@host> show isis interface detail
```

IS-IS interface database:

at-2/3/0.0

Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3

LSP interval: 100 ms, CSNP interval: 5 s

| Level | Adjacencies | Priority | Metric | Hello (s) | Hold (s) | Designated Router |
|-------|-------------|----------|--------|-----------|----------|-------------------|
|-------|-------------|----------|--------|-----------|----------|-------------------|

|   |   |    |    |       |    |  |
|---|---|----|----|-------|----|--|
| 1 | 1 | 64 | 10 | 9.000 | 27 |  |
|---|---|----|----|-------|----|--|

|   |   |    |    |       |    |  |
|---|---|----|----|-------|----|--|
| 2 | 1 | 64 | 10 | 9.000 | 27 |  |
|---|---|----|----|-------|----|--|

lo0.0

Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0

LSP interval: 100 ms, CSNP interval: disabled

Adjacency advertisement: Advertise

Protection Type: Node Link, No eligible Backup

| Level | Adjacencies | Priority | Metric | Hello (s) | Hold (s) | Designated Router |
|-------|-------------|----------|--------|-----------|----------|-------------------|
|-------|-------------|----------|--------|-----------|----------|-------------------|

|   |   |    |   |         |  |  |
|---|---|----|---|---------|--|--|
| 1 | 0 | 64 | 0 | Passive |  |  |
|---|---|----|---|---------|--|--|

|   |   |    |   |         |  |  |
|---|---|----|---|---------|--|--|
| 2 | 0 | 64 | 0 | Passive |  |  |
|---|---|----|---|---------|--|--|

### show isis interface extensive

```
user@host> show isis interface extensive
```

IS-IS interface database:

xe-6/1/0.0

Index: 75, State: 0x6, Circuit id: 0x1, Circuit type: 2

LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding

Adjacency advertisement: Advertise

Level 1

Adjacencies: 0, Priority: 64, Metric: 10

Disabled

Level 2

Adjacencies: 1, Priority: 64, Metric: 10

Hello Interval: 20.000 s, Hold Time: 60 s

Designated Router: nemean.03

## show isis overview

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                     | <b>show isis overview</b><br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <b>show isis overview</b><br><instance <i>instance-name</i> >                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                        | Command introduced in Junos OS Release 8.5.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                            |
| <b>Description</b>                                | Display IS-IS overview information.                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                    | <b>none</b> —Display standard overview information about IS-IS for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display overview information for the specified routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>                      | <a href="#">show isis overview on page 3184</a>                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>                              | <a href="#">Table 335 on page 3182</a> lists the output fields for the <b>show isis overview</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                               |

**Table 335: show isis overview Output Fields**

| Field Name              | Field Description                                                                                                        |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Instance                | IS-IS routing instance.                                                                                                  |
| Router ID               | Router ID of the routing device.                                                                                         |
| Adjacency holddown      | Adjacency holddown capability: <b>enabled</b> or <b>disabled</b> .                                                       |
| Maximum Areas           | Maximum number of IS-IS areas advertised by the routing device.                                                          |
| LSP life time           | Lifetime of the link-state PDU, in seconds.                                                                              |
| Attached bit evaluation | Attached bit capability: <b>enabled</b> or <b>disabled</b> .                                                             |
| SPF delay               | Delay before performing consecutive shortest-path-first (SPF) calculations.                                              |
| SPF holddown            | Delay before performing additional SPF calculations after the maximum number of consecutive SPF calculations is reached. |

Table 335: show isis overview Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                       |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| SPF rapid runs                 | Maximum number of SPF calculations that can be performed in succession before the holddown timer begins.                |
| Overload bit at startup is set | Overload bit capability is enabled.                                                                                     |
| Overload high metrics          | Overload high metrics capability: <b>enabled</b> or <b>disabled</b> .                                                   |
| Overload timeout               | Time period after which overload is reset and the time that remains before the timer is set to expire.                  |
| Traffic engineering            | Traffic engineering capability: <b>enabled</b> or <b>disabled</b> .                                                     |
| Restart                        | Graceful restart capability: <b>enabled</b> or <b>disabled</b> .                                                        |
| Restart duration               | Time period for complete reacquisition of IS-IS neighbors.                                                              |
| Helper mode                    | Graceful restart helper capability: <b>enabled</b> or <b>disabled</b> .                                                 |
| Level                          | IS-IS level: <ul style="list-style-type: none"> <li>• 1—Level 1 information</li> <li>• 2—Level 2 information</li> </ul> |
| IPv4 is enabled                | IP Protocol version 4 capability is enabled.                                                                            |
| IPv6 is enabled                | IP Protocol version 6 capability is enabled.                                                                            |
| CLNS is enabled                | (J Series routers only) OSI CLNP capability is enabled.                                                                 |
| Internal route preference      | Preference value of internal routes.                                                                                    |
| External route preference      | Preference value of external routes.                                                                                    |
| Prefix export limit            | Number of prefixes allowed to be exported, as configured by the <a href="#">prefix-export-limit</a> statement.          |
| Prefix export count            | Number of prefixes exported.                                                                                            |
| Wide area metrics are enabled  | Wide area metrics capability is enabled.                                                                                |
| Narrow metrics are enabled     | Narrow metrics capability is enabled.                                                                                   |

## Sample Output

### show isis overview

```
user@host> show isis overview
Instance: master
  Router ID: 10.255.107.183
  Adjacency holddown: disabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
Level 1
  Internal route preference: 15
  External route preference: 160
  Wide metrics are enabled, Narrow metrics are enabled
Level 2
  Internal route preference: 18
  External route preference: 165
  Prefix export limit: 5, Prefix export count: 5
  Wide metrics are enabled
```



## show isis route

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3185</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3185</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax</b>                                     | <pre>show isis route &lt;destination&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis route &lt;destination&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                | Display the routes in the IS-IS routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                    | <p><b>none</b>—Display all routes in the IS-IS routing table for all supported address families for all routing instances.</p> <p><b><i>destination</i></b>—(Optional) Destination address for the route.</p> <p><b>inet   inet6</b>—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routes for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)</b>—(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>                      | <a href="#">show isis route logical-system on page 3186</a><br><a href="#">show isis route (CLNS) on page 3186</a><br><a href="#">show isis route on page 3187</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Output Fields</b>                              | <p><a href="#">Table 336 on page 3186</a> describes the output fields for the <b>show isis route</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 336: show isis route Output Fields

| Field Name             | Field Description                                                                                                                          |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Current version</b> | Number of the current version of the IS-IS routing table.                                                                                  |
| <b>L1</b>              | Version of Level 1 SPF that was run.                                                                                                       |
| <b>L2</b>              | Version of Level 2 SPF that was run.                                                                                                       |
| <b>Prefix</b>          | Destination of the route.                                                                                                                  |
| <b>L</b>               | IS-IS level: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul> |
| <b>Version</b>         | Version of SPF that generated the route.                                                                                                   |
| <b>Metric</b>          | Metric value associated with the route.                                                                                                    |
| <b>Type</b>            | Metric type: <b>int</b> (internal) or <b>ext</b> (external).                                                                               |
| <b>Interface</b>       | Interface to the next hop.                                                                                                                 |
| <b>Via</b>             | System identifier of the next hop, displayed as a name if possible.                                                                        |
| <b>ISO Routes</b>      | ISO routing table entries.                                                                                                                 |
| <b>snpa</b>            | MAC address.                                                                                                                               |

## Sample Output

### show isis route logical-system

```

user@host> show isis route logical-system ls1
IS-IS routing table           Current version: L1: 8 L2: 11
Prefix      L Version Metric Type Interface  Via
10.9.7.0/30  2    11    20 int  gr-0/2/0.0  h
10.9.201.1/32 2    11    60 int  gr-0/2/0.0  h
IPv6 Unicast IS-IS routing table   Current version: L1: 9 L2: 11
Prefix      L Version Metric Type Interface  Via
8009:3::a09:3200/126 2    11    20 int  gr-0/2/0.0  h

```

### show isis route (CLNS)

```

user@host> show isis route
IS-IS routing table           Current version: L1: 10 L2: 8
IPv4/IPv6 Routes
Prefix      L Version Metric Type Interface  Via
0.0.0.0/0   1    10    10 int  fe-0/0/1.0  ISIS.0
ISO Routes
Prefix L   Version Metric Type Interface  Via  snpa

```

```

0/0
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001/104
  1      10      0 int
47.0005.80ff.f800.0000.0108.0001.1921.6800.4001/152
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001.1921.6800.4002/152
  1      10      20 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0002/104
  1      10      0 int
47.0005.80ff.f800.0000.0108.0002.1921.6800.4001/152
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56

```

### show isis route

```
user@host> show isis route
```

```

IS-IS routing table          Current version: L1: 4 L2: 13
IPv4/IPv6 Routes
-----
Prefix                      L   Version  Metric Type Interface      NH   Via
10.255.71.52/32             2    13        10   int  ae0.0                 IPV4 camaro
10.255.71.238/32            2    13        20   int  so-6/0/0.0           IPV4 olympic
                             as0.0                 IPV4 glacier
10.255.71.239/32            2    13        20   int  so-6/0/0.0           IPV4 olympic
                             ae0.0                 IPV4 camaro
10.255.71.242/32            2    13        10   int  as0.0                 IPV4 glacier
10.255.71.243/32            2    13        10   int  so-6/0/0.0           IPV4 olympic
12.13.0.0/30                2    13        20   int  so-6/0/0.0           IPV4 olympic
12.15.0.0/30                2    13        20   int  so-6/0/0.0           IPV4 olympic
13.15.0.0/30                2    13        30   int  ae0.0                 IPV4 camaro
                             so-6/0/0.0           IPV4 olympic
                             as0.0                 IPV4 glacier
13.16.0.0/30                2    13        25   int  as0.0                 IPV4 glacier
14.15.0.0/30                2    13        20   int  ae0.0                 IPV4 camaro
192.2.1.0/30                2    13        30   int  so-6/0/0.0           IPV4 olympic
                             as0.0                 IPV4 glacier
1eee::/64                   2    13        30   int  so-6/0/0.0           IPV6 olympic
                             as0.0                 IPV6 glacier
abcd::10:255:71:52/128      2    13        10   int  ae0.0                 IPV6 camaro
abcd::10:255:71:238/128     2    13        20   int  so-6/0/0.0           IPV6 olympic

```

|                         |   |    |    |     |            |              |
|-------------------------|---|----|----|-----|------------|--------------|
|                         |   |    |    |     | as0.0      | IPv6 glacier |
| abcd::10:255:71:239/128 | 2 | 13 | 20 | int | so-6/0/0.0 | IPv6 olympic |
|                         |   |    |    |     | ae0.0      | IPv6 camaro  |
| abcd::10:255:71:242/128 | 2 | 13 | 10 | int | as0.0      | IPv6 glacier |
| abcd::10:255:71:243/128 | 2 | 13 | 10 | int | so-6/0/0.0 | IPv6 olympic |

## show isis spf

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3189</a><br><a href="#">Syntax (EX Series Switches) on page 3189</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                      | <pre>show isis spf (brief   log   results) &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (EX Series Switches)</b> | <pre>show isis spf (brief   log   results) &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                 | Display information about IS-IS shortest-path-first (SPF) calculations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                     | <p><b>brief</b>—Display an overview of SPF calculations.</p> <p><b>instance <i>instance instance-name</i></b>—(Optional) Display SPF calculations for the specified routing instance.</p> <p><b>level (1   2)</b>—(Optional) Display SPF calculations for the specified IS-IS level.</p> <p><b>log</b>—Display the log of SPF calculations.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>results</b>—Display the results of SPF calculations.</p> <p><b>topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)</b>—(Optional) Display SPF calculations for the specified topology only.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>       | <a href="#">show isis spf log on page 3190</a><br><a href="#">show isis spf results logical-system on page 3191</a><br><a href="#">show isis spf results (CLNS) on page 3192</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>               | <p><a href="#">Table 337 on page 3189</a> describes the output fields for the <b>show isis spf</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Table 337: show isis spf Output Fields**

| Field Name | Field Description    |
|------------|----------------------|
| Node       | System ID of a node. |

Table 337: show isis spf Output Fields (*continued*)

| Field Name     | Field Description                                                                       |
|----------------|-----------------------------------------------------------------------------------------|
| Metric         | Metric to the node.                                                                     |
| Interface      | Interface of the next hop.                                                              |
| Via            | System ID of the next hop.                                                              |
| SNPA           | Subnetwork point of attachment (MAC address of the next hop).                           |
| Start time     | (log option only) Time that the SPF computation started.                                |
| Elapsed (secs) | (log option only) Length of time, in seconds, required to complete the SPF computation. |
| Count          | (log option only) Number of times the SPF was triggered.                                |
| Reason         | (log option only) Reason that the SPF computation was completed.                        |

## Sample Output

### show isis spf log

```

user@host> show isis spf log logical-system lsl
IS-IS level 1 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000069    1 Reconfig
Fri Oct 31 12:41:18   0.000107    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000050    3 Address change on so-1/2/2.0
Fri Oct 31 12:41:23   0.000033    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28   0.000178    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59   0.000060    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30   0.000161    2 Multi area attachment change
Fri Oct 31 12:56:58   0.000198    1 Periodic SPF
Fri Oct 31 13:10:29   0.000209    1 Periodic SPF
IS-IS level 2 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000035    1 Reconfig
Fri Oct 31 12:41:18   0.000047    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000043    5 Address change on gr-0/2/0.0
Fri Oct 31 12:41:23   0.000022    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59   0.000144    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30   0.000257    3 New LSP skag.00-00
Fri Oct 31 12:54:37   0.000195    1 Periodic SPF
Fri Oct 31 12:55:50   0.000178    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55   0.000174    1 Updated LSP h.00-00
Fri Oct 31 12:55:58   0.000176    1 Updated LSP skag.00-00
Fri Oct 31 13:08:14   0.000198    1 Periodic SPF
IPv6 Unicast IS-IS level 1 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000028    1 Reconfig
Fri Oct 31 12:41:18   0.000043    3 Updated LSP fix.00-00

```

```

Fri Oct 31 12:41:18      0.000112    4 Updated LSP fix.00-00
Fri Oct 31 12:41:23      0.000059    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25      0.000041    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28      0.000103    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59      0.000040    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30      0.000118    2 Multi area attachment change
Fri Oct 31 12:56:08      0.000289    1 Periodic SPF
Fri Oct 31 13:11:07      0.000214    1 Periodic SPF
IPV6 Unicast IS-IS level 2 SPF log:

```

```

Start time      Elapsed (secs) Count Reason
Fri Oct 31 12:41:18      0.000027    1 Reconfig
Fri Oct 31 12:41:18      0.000039    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18      0.000049    6 Updated LSP fix.00-00
Fri Oct 31 12:41:23      0.000025    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25      0.000023    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59      0.000087    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30      0.000123    3 New LSP skag.00-00
Fri Oct 31 12:55:50      0.000121    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55      0.000121    1 Updated LSP h.00-00
Fri Oct 31 12:55:58      0.000121    1 Updated LSP skag.00-00
Fri Oct 31 13:09:46      0.000201    1 Periodic SPF
...

```

#### show isis spf results logical-system

```
user@host> show isis spf results logical-system ls1
```

```
IS-IS level 1 SPF results:
```

| Node    | Metric | Interface     | Via  | SNPA             |
|---------|--------|---------------|------|------------------|
| scat.00 | 10     | ge-1/1/0.0    | scat | 0:90:69:a6:48:9d |
|         | 20     | 10.9.1.0/30   |      |                  |
| fix.02  | 10     |               |      |                  |
| fix.00  | 0      |               |      |                  |
|         | 10     | 10.9.1.0/30   |      |                  |
|         | 10     | 10.9.5.0/30   |      |                  |
|         | 10     | 10.9.6.0/30   |      |                  |
|         | 20     | 10.9.7.0/30   |      |                  |
|         | 60     | 10.9.201.1/32 |      |                  |
| 3 nodes |        |               |      |                  |

```
IS-IS level 2 SPF results:
```

| Node    | Metric | Interface     | Via | SNPA |
|---------|--------|---------------|-----|------|
| skag.00 | 20     | gr-0/2/0.0    | h   |      |
|         | 30     | 10.9.7.0/30   |     |      |
| skag.02 | 20     | gr-0/2/0.0    | h   |      |
| h.00    | 10     | gr-0/2/0.0    | h   |      |
|         | 20     | 10.9.6.0/30   |     |      |
|         | 20     | 10.9.7.0/30   |     |      |
|         | 60     | 10.9.201.1/32 |     |      |
| fix.00  | 0      |               |     |      |
|         | 10     | 10.9.1.0/30   |     |      |
|         | 10     | 10.9.5.0/30   |     |      |
|         | 10     | 10.9.6.0/30   |     |      |
| 4 nodes |        |               |     |      |

```
IPV6 Unicast IS-IS level 1 SPF results:
```

| Node    | Metric | Interface            | Via  | SNPA             |
|---------|--------|----------------------|------|------------------|
| scat.00 | 10     | ge-1/1/0.0           | scat | 0:90:69:a6:48:9d |
|         |        | ge-1/1/0.0           | scat | 0:90:69:a6:48:9d |
|         | 20     | 8009:1::a09:1400/126 |      |                  |
| fix.02  | 10     |                      |      |                  |

```

fix.00      0
            10      8009:1::a09:1400/126
            10      8009:2::a09:1e00/126
            20      8009:3::a09:3200/126
            10      8009:4::a09:2800/126

3 nodes

IPv6 Unicast IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00    20      gr-0/2/0.0    h
           30      8009:3::a09:3200/126
           20      gr-0/2/0.0    h
skag.02    20      gr-0/2/0.0    h
           10      gr-0/2/0.0    h
h.00       10      gr-0/2/0.0    h
           20      8009:3::a09:3200/126
           20      8009:4::a09:2800/126
fix.00     0
           10      8009:1::a09:1400/126
           10      8009:2::a09:1e00/126
           10      8009:4::a09:2800/126

4 nodes

Multicast IS-IS level 1 SPF results:
Node      Metric      Interface      Via      SNPA
scat.00    10      ge-1/1/0.0    scat    0:90:69:a6:48:9d
fix.02     10
fix.00     0
3 nodes

Multicast IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00    20      gr-0/2/0.0    h
skag.02    20      gr-0/2/0.0    h
h.00       10      gr-0/2/0.0    h
fix.00     0
4 nodes
...

```

### show isis spf results (CLNS)

```

user@host> show isis spf results
IS-IS level 1 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00 10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           20      fe-0/0/1.0    toothache 0:12:0:34:0:56
           10      192.168.37.64/29
           20      192.168.37.64/29
           10      192.168.37.64/29
pro1-a.02 10
pro1-a.00 0
           0      10.255.245.1/32
           10      192.168.37.64/29
           0      192.168.37.64/29

3 nodes

IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00 10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           10      fe-0/0/1.0    toothache 0:12:0:34:0:56

```



|           |    |                                      |
|-----------|----|--------------------------------------|
|           | 20 | 10.255.245.1/32                      |
|           | 20 | 192.168.37.64/29                     |
|           | 20 | 47.0005.80ff.f800.0000.0109.0010/104 |
| pro1-a.02 | 10 |                                      |
| pro1-a.00 | 0  |                                      |
|           | 0  | 10.255.245.1/32                      |
|           | 10 | 192.168.37.64/29                     |
| 3 nodes   |    |                                      |

## show isis statistics

---

|                                                   |                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3194</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3194</a>                                                                                                                                                                                                                                       |
| <b>Syntax</b>                                     | show isis statistics<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches and QFX Series)</b> | show isis statistics<br><instance <i>instance-name</i> >                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                             |
| <b>Description</b>                                | Display statistics about IS-IS traffic.                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                    | <b>none</b> —Display IS-IS traffic statistics for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display statistics for the specified routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">clear isis statistics on page 3148</a></li></ul>                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                      | <a href="#">show isis statistics on page 3196</a>                                                                                                                                                                                                                                                                                                    |
| <b>Output Fields</b>                              | <a href="#">Table 338 on page 3195</a> describes the output fields for the <b>show isis statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                              |

Table 338: show isis statistics Output Fields

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PDU type                    | <p>PDU type:</p> <ul style="list-style-type: none"> <li>• <b>CSNP</b>—Complete sequence number PDUs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU.</li> <li>• <b>IIH</b>—IS-IS hello packets are broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems.</li> <li>• <b>LSP</b>—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area.</li> <li>• <b>PSNP</b>—Partial sequence number PDUs are sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device.</li> <li>• <b>Unknown</b>—The PDU type is unknown.</li> </ul> |
| Received                    | Number of PDUs received since IS-IS started or since the statistics were set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Processed                   | Number of PDUs received less the number dropped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Drops                       | Number of PDUs dropped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Sent                        | Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Rexmit                      | Number of PDUs retransmitted since IS-IS started or since the statistics were set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Total packets received/sent | Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| SNP queue length            | Number of CSPN and PSNP packets currently waiting in the queue for processing. This value is almost always 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| LSP queue length            | Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| SPF runs                    | Number of shortest-path-first (SPF) calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Fragments rebuilt           | Number of link-state PDU fragments that the local system has computed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| LSP regenerations           | Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Purges initiated            | Number of purges that the system initiated. A purge is initiated if the software decides that a link-state PDU must be removed from the network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Sample Output

### show isis statistics

```
user@host> show isis statistics
```

```
IS-IS statistics for merino:
```

| PDU type | Received | Processed | Drops | Sent   | Rexmit |
|----------|----------|-----------|-------|--------|--------|
| LSP      | 12227    | 12227     | 0     | 8184   | 683    |
| IIH      | 113808   | 113808    | 0     | 115817 | 0      |
| CSNP     | 198868   | 198868    | 0     | 198934 | 0      |
| PSNP     | 6985     | 6979      | 6     | 8274   | 0      |
| Unknown  | 0        | 0         | 0     | 0      | 0      |
| Totals   | 331888   | 331882    | 6     | 331209 | 683    |

```
Total packets received: 331888 Sent: 331892
```

```
SNP queue length:      0 Drops:      0
LSP queue length:      0 Drops:      0
```

```
SPF runs:              1014
Fragments rebuilt:     1038
LSP regenerations:     425
Purges initiated:      0
```

## PART 17

# OSPF

- [Overview on page 3199](#)
- [Configuration on page 3207](#)
- [Administration on page 3277](#)



## CHAPTER 54

# Overview

- [Layer 3 Protocols on page 3199](#)
- [OSPF Overview on page 3202](#)

## Layer 3 Protocols

- [Layer 3 Protocols Supported on EX Series Switches on page 3199](#)
- [Layer 3 Protocols Not Supported on EX Series Switches on page 3200](#)

## Layer 3 Protocols Supported on EX Series Switches

EX Series switches support the Junos OS Layer 3 features and configuration statements listed in [Table 314 on page 2939](#):

**Table 339: Supported Junos OS Layer 3 Protocol Statements and Features**

| Protocol           | Notes                                                                                                         | For More Information                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| BGP                | Fully supported.                                                                                              | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| BFD                | Fully supported.                                                                                              | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| ICMP               | Fully supported.                                                                                              | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| IGMPv1, v2, and v3 | Fully supported.                                                                                              | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |
| IS-IS              | Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on EX Series Switches" on page 2940. | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| MLD                | Fully supported (MLD versions 1 and 2).                                                                       | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |
| MPLS               | Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on EX Series Switches" on page 2940. | <a href="#">Junos OS MPLS Applications Configuration Guide</a>   |
| OSPFv1, v2 and v3  | Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on EX Series Switches" on page 2940. | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |

Table 339: Supported Junos OS Layer 3 Protocol Statements and Features (*continued*)

| Protocol | Notes                                                                                                | For More Information                                                                                            |
|----------|------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| PIM      | Fully supported.                                                                                     | <i>Junos OS Multicast Protocols Configuration Guide</i>                                                         |
| PPM      | Supported. See <i>EX Series Switch Software Features Overview</i> for specific platform information. | <i>Junos OS Routing Protocols Configuration Guide</i>                                                           |
| RIP      | Fully supported.                                                                                     | <i>Junos OS Routing Protocols Configuration Guide</i>                                                           |
| RIPng    | Fully supported.                                                                                     | <i>Junos OS Routing Protocols Configuration Guide</i>                                                           |
| SNMP     | Fully supported.                                                                                     | <i>Junos OS Network Management Configuration Guide</i>                                                          |
| VRRP     | Fully supported.                                                                                     | See "Understanding VRRP on EX Series Switches" on page 2501. See also <i>Junos OS High Availability Guide</i> . |

- Related Documentation**
- [Layer 3 Protocols Not Supported on EX Series Switches](#) on page 2940
  - *EX Series Switch Software Features Overview*

## Layer 3 Protocols Not Supported on EX Series Switches

EX Series switches do not support the Junos OS Layer 3 protocols and features listed in [Table 315 on page 2940](#):

Table 340: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported

| Feature                                                                                                         | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DVMRP                                                                                                           | <ul style="list-style-type: none"> <li>• <b>dvmrp</b> and subordinate statements</li> </ul>                                                                                                                                                                                                          |
| Flow aggregation (cflowd)                                                                                       | <ul style="list-style-type: none"> <li>• <b>cflow</b> and subordinate statements</li> </ul>                                                                                                                                                                                                          |
| IPsec                                                                                                           | <ul style="list-style-type: none"> <li>• <b>[edit services]</b> statements related to IPsec</li> </ul>                                                                                                                                                                                               |
| IS-IS: <ul style="list-style-type: none"> <li>• ES-IS</li> <li>• IPv6 in multicast routing protocols</li> </ul> | <ul style="list-style-type: none"> <li>• <b>clns-routing</b> statement</li> <li>• <b>ipv6-multicast</b> statement</li> <li>• <b>lsp-interval</b> statement</li> <li>• <b>label-switched-path</b> statement</li> <li>• <b>lsp-lifetime</b> statement</li> <li>• <b>te-metric</b> statement</li> </ul> |
| Logical routers                                                                                                 | <ul style="list-style-type: none"> <li>• <b>logical-routers</b> and subordinate statements</li> </ul>                                                                                                                                                                                                |



Table 340: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)

| Feature                                                                                                                                                                                                                                                                                                                                                                                                                                      | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MPLS: <ul style="list-style-type: none"> <li>Fast Reroute (FRR)</li> <li>Label Distribution Protocol (LDP) (except on EX8200 switches)</li> <li>Layer 3 VPNs (except on EX8200 switches)</li> <li>Multiprotocol BGP (MP-BGP) for VPN-IPv4 family</li> <li>Pseudowire emulation (PWE3)</li> <li>Routing policy statements related to Layer 3 VPNs and MPLS (except on EX8200 switches)</li> <li>Virtual Private LAN Service (VPLS)</li> </ul> | <ul style="list-style-type: none"> <li><b>ldp</b> and all subordinate statements (except on EX8200 switches)</li> </ul>                                                                                                                                                                                                                                                                                           |
| Network Address Translation (NAT)                                                                                                                                                                                                                                                                                                                                                                                                            | <ul style="list-style-type: none"> <li><b>nat</b> and subordinate statements</li> <li>Policy statements related to NAT</li> </ul>                                                                                                                                                                                                                                                                                 |
| OSPF                                                                                                                                                                                                                                                                                                                                                                                                                                         | <ul style="list-style-type: none"> <li><b>demand-circuit</b> statement</li> <li><b>label-switched-path</b> and subordinate statements</li> <li><b>neighbor</b> statement within an OSPF area</li> <li><b>peer-interface</b> and subordinate statements within an OSPF area</li> <li><b>sham-link</b> statement</li> <li><b>te-metric</b> statement</li> </ul>                                                     |
| PPM                                                                                                                                                                                                                                                                                                                                                                                                                                          | <ul style="list-style-type: none"> <li>Not supported on EX2200 and EX3300 switches</li> </ul>                                                                                                                                                                                                                                                                                                                     |
| Routing instances: <ul style="list-style-type: none"> <li>Routing instance forwarding</li> </ul>                                                                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li><b>l2vpn</b> and subordinate statements (except on EX4500, EX4550, and EX8200 switches)</li> <li><b>ldp</b> and subordinate statements (except on EX8200 switches)</li> <li><b>vpls</b> and subordinate statements</li> </ul>                                                                                                                                              |
| Routed VLAN interfaces (RVIs)                                                                                                                                                                                                                                                                                                                                                                                                                | <ul style="list-style-type: none"> <li><b>family mpls</b> statement</li> </ul>                                                                                                                                                                                                                                                                                                                                    |
| SAP and SDP                                                                                                                                                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li><b>sap</b> and all subordinate statements</li> </ul>                                                                                                                                                                                                                                                                                                                       |
| General routing options in the <b>routing-options</b> hierarchy: <ul style="list-style-type: none"> <li>MPLS and label-switched-paths</li> </ul>                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li><b>auto-export</b> and subordinate statements</li> <li><b>dynamic-tunnels</b> and subordinate statements</li> <li><b>lsp-next-hop</b> and subordinate statements</li> <li><b>multicast</b> and subordinate statements</li> <li><b>p2mp-lsp-next-hop</b> and subordinate statements</li> <li><b>route-distinguisher-id</b> statement (except on EX8200 switches)</li> </ul> |

Table 340: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)

| Feature                                                                    | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Traffic sampling and forwarding in the <b>forwarding-options</b> hierarchy | <ul style="list-style-type: none"> <li>• <b>accounting</b> and subordinate statements</li> <li>• <b>family mpls</b> and <b>family multiservice</b> under <b>hash-key</b> hierarchy</li> <li>• Under <b>monitoring group-name</b> family <b>inet output</b> hierarchy: <ul style="list-style-type: none"> <li>• <b>cflowd</b> statement</li> <li>• <b>export-format-cflowd-version-5</b> statement</li> <li>• <b>flow-active-timeout</b> statement</li> <li>• <b>flow-export-destination</b> statement</li> <li>• <b>flow-inactive-timeout</b> statement</li> <li>• <b>interface</b> statement</li> </ul> </li> <li>• <b>port-mirroring</b> statement (On EX Series switches, port mirroring is implemented using the <b>analyzer</b> statement.)</li> <li>• <b>sampling</b> and subordinate statements</li> </ul> |

- Related Documentation**
- [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)
  - [EX Series Switch Software Features Overview](#)

## OSPF Overview

- [Understanding IPsec Authentication for OSPF Packets on EX Series Switches on page 3202](#)

### Understanding IPsec Authentication for OSPF Packets on EX Series Switches

IP Security (IPsec) provides a secure way to authenticate senders and encrypt IP version 4 (IPv4) traffic between network devices. IPsec offers network administrators for Juniper Networks EX Series Ethernet Switches and their users the benefits of data confidentiality, data integrity, sender authentication, and anti-replay services.

IPsec is a framework for ensuring secure private communication over IP networks and is based on standards developed by the International Engineering Task Force (IETF). IPsec provides security services at the network layer of the Open Systems Interconnection (OSI) model by enabling a system to select required security protocols, determine the algorithms to use for the security services, and implement any cryptographic keys required to provide the requested services. You can use IPsec to protect one or more paths between a pair of hosts, between a pair of security gateways (such as switches), or between a security gateway and a host.

OSPF version 3 (OSPFv3), unlike OSPF version 2 (OSPFv2), does not have a built-in authentication method and relies on IPsec to provide this functionality. You can secure specific OSPFv3 interfaces and protect OSPFv3 virtual links.

- [Authentication Algorithms on page 3203](#)
- [Encryption Algorithms on page 3203](#)
- [IPsec Protocols on page 3204](#)

- [Security Associations on page 3204](#)
- [IPsec Modes on page 3204](#)

### Authentication Algorithms

---

Authentication is the process of verifying the identity of the sender. Authentication algorithms use a shared key to verify the authenticity of the IPsec devices. The Juniper Networks Junos operating system (Junos OS) uses the following authentication algorithms:

- Message Digest 5 (MD5) uses a one-way hash function to convert a message of arbitrary length to a fixed-length message digest of 128 bits. Because of the conversion process, it is mathematically infeasible to calculate the original message by computing it backwards from the resulting message digest. Likewise, a change to a single character in the message will cause it to generate a very different message digest number.

To verify that the message has not been tampered with, Junos OS compares the calculated message digest against a message digest that is decrypted with a shared key. Junos OS uses the MD5 hashed message authentication code (HMAC) variant that provides an additional level of hashing. MD5 can be used with an authentication header (AH) and Encapsulating Security Payload (ESP).

- Secure Hash Algorithm 1 (SHA-1) uses a stronger algorithm than MD5. SHA-1 takes a message of less than 264 bits in length and produces a 160-bit message digest. The large message digest ensures that the data has not been changed and that it originates from the correct source. Junos OS uses the SHA-1 HMAC variant that provides an additional level of hashing. SHA-1 can be used with AH, ESP, and Internet Key Exchange (IKE).

### Encryption Algorithms

---

Encryption encodes data into a secure format so that it cannot be deciphered by unauthorized users. As with authentication algorithms, a shared key is used with encryption algorithms to verify the authenticity of IPsec devices. Junos OS uses the following encryption algorithms:

- Data Encryption Standard cipher-block chaining (DES-CBC) is a symmetric secret-key block algorithm. DES uses a key size of 64 bits, where 8 bits are used for error detection and the remaining 56 bits provide encryption. DES performs a series of simple logical operations on the shared key, including permutations and substitutions. CBC takes the first block of 64 bits of output from DES, combines this block with the second block, feeds this back into the DES algorithm, and repeats this process for all subsequent blocks.
- Triple DES-CBC (3DES-CBC) is an encryption algorithm that is similar to DES-CBC but provides a much stronger encryption result because it uses three keys for 168-bit (3 x 56-bit) encryption. 3DES works by using the first key to encrypt the blocks, the second key to decrypt the blocks, and the third key to reencrypt the blocks.

## IPsec Protocols

---

IPsec protocols determine the type of authentication and encryption applied to packets that are secured by the switch. Junos OS supports the following IPsec protocols:

- **AH**—Defined in *RFC 2402*, AH provides connectionless integrity and data origin authentication for IPv4. It also provides protection against replays. AH authenticates as much of the IP header as possible, as well as the upper-level protocol data. However, some IP header fields might change in transit. Because the value of these fields might not be predictable by the sender, they cannot be protected by AH. In an IP header, AH can be identified with a value of 51 in the Protocol field of an IPv4 packet.
- **ESP**—Defined in *RFC 2406*, ESP can provide encryption and limited traffic flow confidentiality or connectionless integrity, data origin authentication, and an anti-replay service. In an IP header, ESP can be identified with a value of 50 in the Protocol field of an IPv4 packet.

## Security Associations

---

An IPsec consideration is the type of security association (SA) that you wish to implement. An SA is a set of IPsec specifications that are negotiated between devices that are establishing an IPsec relationship. These specifications include preferences for the type of authentication, encryption, and IPsec protocol to be used when establishing the IPsec connection. An SA can be either unidirectional or bidirectional, depending on the choices made by the network administrator. An SA is uniquely identified by a Security Parameter Index (SPI), an IPv4 or IPv6 destination address, and a security protocol (AH or ESP) identifier.

## IPsec Modes

---

Junos OS supports the following IPsec modes:

- **Tunnel mode** is supported for both AH and ESP in Junos OS. In tunnel mode, the SA and associated protocols are applied to tunneled IPv4 or IPv6 packets. For a tunnel mode SA, an outer IP header specifies the IPsec processing destination and an inner IP header specifies the ultimate destination for the packet. The security protocol header appears after the outer IP header and before the inner IP header. In addition, there are slight differences for tunnel mode when you implement it with AH and ESP:
  - For AH, portions of the outer IP header are protected, as well as the entire tunneled IP packet.
  - For ESP, only the tunneled packet is protected, not the outer header.

When one side of an SA is a security gateway (such as a switch), the SA must use tunnel mode. However, when traffic (for example, SNMP commands or BGP sessions) is destined for a switch, the system acts as a host. Transport mode is allowed in this case because the system does not act as a security gateway and does not send or receive transit traffic.



**NOTE:** Tunnel mode is not supported for OSPF v3 control packet authentication.

- Transport mode provides an SA between two hosts. In transport mode, the protocols provide protection primarily for upper-layer protocols. A transport mode security protocol header appears immediately after the IP header and any options and before any higher-layer protocols (for example, TCP or UDP). There are slight differences for transport mode when you implement it with AH and ESP:
  - For AH, selected portions of the IP header are protected, as well as selected portions of the extension headers and selected options within the IPv4 header.
  - For ESP, only the higher-layer protocols are protected, not the IP header or any extension headers preceding the ESP header.

**Related  
Documentation**

- [Using IPsec to Secure OSPFv3 Networks \(CLI Procedure\) on page 3211](#)
- [Configuring an OSPF Network \(J-Web Procedure\) on page 3207](#)



# Configuration

- [Configuration Tasks on page 3207](#)
- [Configuration Statements on page 3212](#)

## Configuration Tasks

---

- [Configuring an OSPF Network \(J-Web Procedure\) on page 3207](#)
- [Using IPsec to Secure OSPFv3 Networks \(CLI Procedure\) on page 3211](#)

### Configuring an OSPF Network (J-Web Procedure)

---



**NOTE:** This topic applies only to the J-Web Application package.

---

You can use the J-Web interface to create multiarea OSPF networks on an EX Series switch.

To configure a multiarea OSPF network:

1. Select **Configure > Routing > OSPF**.
- 



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

---

2. Click one of the following options:
  - **Add**—Adds an OSPF area. Enter information into the configuration page as described in [Table 341 on page 3208](#).
  - **Edit**—Modifies an existing OSPF area. Enter information into the configuration page as described in [Table 341 on page 3208](#).
  - **Delete**—Deletes an existing OSPF area.

3. To modify OSPF global settings, click **Edit**. Enter information as described in [Table 342 on page 3209](#).
4. To disable OSPF, click **Disable**.

**Table 341: OSPF Routing Configuration Summary**

| Field                 | Function                                                                                                                                                                                                                                    | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>General tab</b>    |                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Area Id               | Uniquely identifies the area within its AS.                                                                                                                                                                                                 | <p>Type a 32-bit numeric identifier for the area. Type an integer or select and edit the value.</p> <p>If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the area is <b>0.0.0.3</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Area Ranges           | Specifies a range of IP addresses for an area when sending summary link advertisements (within an area).                                                                                                                                    | <p>To add a range:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b>.</li> <li>2. Type the area range.</li> <li>3. Specify the subnet mask.</li> <li>4. To override the metric for the IP address range, type a specific metric value.</li> <li>5. If you do not want to display the routes that are contained within a summary, select <b>Restrict advertisements of this area range</b>.</li> <li>6. If you want a summary of a route to be advertised only when an exact match is made with the configured summary range, select <b>Enforce exact match for advertisement of this area range</b>.</li> <li>7. Click <b>OK</b>.</li> </ol> <p>To modify an existing area range, select the area range, click <b>Edit</b>, and edit the value. Click <b>OK</b>.</p> <p>To delete an area range, select the area range and click <b>Delete</b>.</p> |
| Area Type             | Designates the type of OSPF area. <ul style="list-style-type: none"> <li>• <b>regular</b>—A regular OSPF area, including the backbone area</li> <li>• <b>stub</b>—A stub area</li> <li>• <b>nssa</b>—A not-so-stubby area (NSSA)</li> </ul> | <p>Select the type of OSPF area you are creating from the list.</p> <p>If you select <b>stub</b>:</p> <ol style="list-style-type: none"> <li>1. Enter the default metric.</li> <li>2. To flood summary LSAs into the stub area, select the check box.</li> </ol> <p>If you select <b>nssa</b>:</p> <ol style="list-style-type: none"> <li>1. Specify the metric type.</li> <li>2. Enter the default metric.</li> <li>3. To flood summary LSAs into the nssa area, select the check box.</li> <li>4. To flood Type-7 LSAs into the nssa area, select the check box.</li> </ol>                                                                                                                                                                                                                                                                                    |
| <b>Interfaces tab</b> |                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |



Table 341: OSPF Routing Configuration Summary (*continued*)

| Field               | Function                                                                                                                                                    | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interfaces          | Specifies the interfaces to be associated with the OSPF configuration                                                                                       | <p>To associate an interface with the configuration, select the interface from the list, select <b>Associate</b> and click <b>OK</b>.</p> <p>To edit an interface's configuration:</p> <ol style="list-style-type: none"> <li>1. Select the interface from the list and click <b>Edit</b>.</li> <li>2. Specify the cost of an OSPF interface.</li> <li>3. Specify the traffic engineering metric.</li> <li>4. Specify how often the routing device sends hello packets from the interface.</li> <li>5. Specify how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements to an interface's neighbors.</li> <li>6. To enable OSPF on the interface, select the check box.</li> <li>7. To inform other protocols about neighbor down events, select the check box.</li> <li>8. To treat the interface as a secondary interface, select the check box.</li> <li>9. To only advertise OSPF, select the check box.</li> <li>10. Click <b>OK</b>.</li> </ol> |
| <b>Policies tab</b> |                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Import Policy       | Specifies one or more policies to control which routes learned from an area are used to generate summary link-state advertisements (LSAs) into other areas. | <p>Click <b>Add</b> to add an import policy.</p> <p>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.</p> <p>Click <b>Remove</b> to remove an import policy.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Export Policy       | Specifies one or more policies to control which summary LSAs are flooded into an area.                                                                      | <p>Click <b>Add</b> to add an export policy.</p> <p>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.</p> <p>Click <b>Remove</b> to remove an export policy.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

Table 342: Edit OSPF Global Settings

| Field                     | Function                                                                                                     | Your Action                        |
|---------------------------|--------------------------------------------------------------------------------------------------------------|------------------------------------|
| <b>General tab</b>        |                                                                                                              |                                    |
| Router Id                 | Specifies the ID for the routing device.                                                                     | Type or select and edit the value. |
| RIB Group                 | Installs the routes learned from OSPF routing instances into routing tables in the OSPF routing table group. | Select a value.                    |
| Internal Route Preference | Specifies the route preference for internal groups.                                                          | Type or select and edit the value. |

Table 342: Edit OSPF Global Settings (*continued*)

| Field                     | Function                                                                                                                                                                                                                                                                                                                                | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| External Route Preference | Specifies the route preference for external groups.                                                                                                                                                                                                                                                                                     | Type or select and edit the value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Graceful Restart          | Configures graceful restart for OSPF.                                                                                                                                                                                                                                                                                                   | <p>To configure graceful restart:</p> <ol style="list-style-type: none"> <li>1. Specify the estimated time to send out purged grace LSAs over all the interfaces.</li> <li>2. Specified the estimated time to reacquire a full OSPF neighbor from each area.</li> <li>3. To disable <b>No Strict LSA Checking</b>, select the check box.</li> <li>4. To disable graceful restart helper capability, select the check box. Helper mode is enabled by default.</li> <li>5. Click <b>OK</b>.</li> </ol>                               |
| SPF Options               | Configure options for running the shortest-path-first (SPF) algorithm. You can configure a delay for when to run the SPF algorithm after a network topology change is detected, the maximum number of times the SPF algorithm can run in succession, and a hold-down interval after the SPF algorithm runs the maximum number of times. | <p>To configure SPF:</p> <ol style="list-style-type: none"> <li>1. Specify the time interval between the detection of a topology change and when the SPF algorithm runs.</li> <li>2. Specify the time interval to hold down, or wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession.</li> <li>3. Specify the maximum number of times the SPF algorithm can run in succession. After the maximum is reached, the hold-down interval begins.</li> </ol> |
| <b>Policies tab</b>       |                                                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Import Policy             | Specifies one or more policies to control which routes learned from an area are used to generate summary link-state advertisements (LSAs) into other areas.                                                                                                                                                                             | <p>Click <b>Add</b> to add an import policy.</p> <p>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.</p> <p>Click <b>Remove</b> to remove an import policy.</p>                                                                                                                                                                                                                                                                                                               |
| Export Policy             | Specifies one or more policies to control which summary LSAs are flooded into an area.                                                                                                                                                                                                                                                  | <p>Click <b>Add</b> to add an export policy.</p> <p>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.</p> <p>Click <b>Remove</b> to remove an export policy.</p>                                                                                                                                                                                                                                                                                                               |
| <b>Trace Options tab</b>  |                                                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| File Name                 | Specifies the name of the file to receive the output of the tracing operation.                                                                                                                                                                                                                                                          | Type or select and edit the name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Number of Files           | Specifies the maximum number of trace files.                                                                                                                                                                                                                                                                                            | Type or select and edit the name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| File Size                 | Specifies the maximum size for each trace file.                                                                                                                                                                                                                                                                                         | Type or select and edit the name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 342: Edit OSPF Global Settings (*continued*)

| Field          | Function                                                         | Your Action                                                                                                                          |
|----------------|------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| World Readable | Specifies whether the trace file can be read by any user or not. | Select <b>True</b> to allow any user to read the file.<br><br>Select <b>False</b> to disallow all users being able to read the file. |
| Flags          | Specifies the tracing operation to perform.                      | Select a value from the list.                                                                                                        |

- Related Documentation**
- [Monitoring OSPF Routing Information on page 3277](#)
  - [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)

## Using IPsec to Secure OSPFv3 Networks (CLI Procedure)

OSPF version 3 (OSPFv3) does not have a built-in authentication method and relies on IP Security (IPsec) to provide this functionality. You can use IPsec to secure OSPFv3 interfaces on EX Series switches.

This topic includes:

- [Configuring Security Associations on page 3211](#)
- [Securing OPSFv3 Networks on page 3212](#)

### Configuring Security Associations

When you configure a security association (SA), include your choices for authentication, encryption, direction, mode, protocol, and security parameter index (SPI).

To configure a security association:

1. Specify a name for the security association:
 

```
[edit security ipsec]
user@switch# set security-association sa-name
```
2. Specify the mode of the security association:
 

```
[edit security ipsec security-association sa-name]
user@switch# set mode transport
```
3. Specify the type of security association:
 

```
[edit security ipsec security-association sa-name]
user@switch# set type manual
```
4. Specify the direction of the security association:
 

```
[edit security ipsec security-association sa-name]
user@switch# set direction bidirectional
```
5. Specify the value of the security parameter index:
 

```
[edit security ipsec security-association sa-name]
user@switch# set spi spi-value
```
6. Specify the type of authentication to be used:
 

```
[edit security ipsec security-association sa-name]
```

```
user@switch# set authentication algorithm type
```

7. Specify the encryption algorithm and key:

```
[edit security ipsec security-association sa-name]
```

```
user@switch# set encryption algorithm algorithm key type
```

---

## Securing OPSFv3 Networks

You can secure the OPSFv3 network by applying the SA to the OPSFv3 configuration.

To secure the OPSFv3 network:

```
[edit protocols ospf3 area area-number interface interface-name]
```

```
user@switch# set ipsec-sa sa-name
```

### Related Documentation

- [Understanding IPsec Authentication for OSPF Packets on EX Series Switches on page 3202](#)
- [Configuring an OSPF Network \(J-Web Procedure\) on page 3207](#)
- [Junos OS System Basics Configuration Guide](#)

---

## Configuration Statements

- [area on page 3214](#)
- [area-range on page 3216](#)
- [bandwidth-based-metrics on page 3218](#)
- [bfd-liveness-detection \(Protocols OSPF\) on page 3220](#)
- [dead-interval on page 3224](#)
- [default-lsa on page 3225](#)
- [default-metric on page 3226](#)
- [disable \(OSPF\) on page 3228](#)
- [domain-id on page 3229](#)
- [domain-vpn-tag on page 3230](#)
- [export \(Protocols OSPF\) on page 3231](#)
- [external-preference \(Protocols OSPF\) on page 3232](#)
- [graceful-restart \(Protocols OSPF\) on page 3233](#)
- [hello-interval \(Protocols OSPF\) on page 3235](#)
- [ignore-lsp-metrics on page 3236](#)
- [import \(Protocols OSPF\) on page 3237](#)
- [inter-area-prefix-export on page 3238](#)
- [inter-area-prefix-import on page 3239](#)
- [interface \(Protocols OSPF\) on page 3240](#)
- [interface-type \(Protocols OSPF\) on page 3242](#)
- [lsp-metric-into-summary on page 3243](#)
- [metric \(Protocols OSPF Interface\) on page 3244](#)

- [metric-type on page 3246](#)
- [no-nssa-abr on page 3247](#)
- [no-rfc-1583 on page 3248](#)
- [nssa on page 3249](#)
- [ospf on page 3250](#)
- [ospf3 on page 3251](#)
- [overload \(Protocols OSPF\) on page 3252](#)
- [passive \(Protocols OSPF\) on page 3254](#)
- [preference \(Protocols OSPF\) on page 3255](#)
- [prefix-export-limit \(Protocols OSPF\) on page 3256](#)
- [priority \(Protocols OSPF\) on page 3257](#)
- [realm on page 3258](#)
- [reference-bandwidth \(Protocols OSPF\) on page 3259](#)
- [retransmit-interval \(OSPF\) on page 3260](#)
- [rib-group \(Protocols OSPF\) on page 3261](#)
- [route-type-community on page 3262](#)
- [shortcuts \(Protocols OSPF\) on page 3263](#)
- [spf-options \(Protocols OSPF\) on page 3264](#)
- [stub on page 3266](#)
- [summaries on page 3267](#)
- [traceoptions \(Protocols OSPF\) on page 3268](#)
- [traffic-engineering \(OSPF\) on page 3271](#)
- [transit-delay \(OSPF\) on page 3273](#)
- [type-7 on page 3274](#)
- [virtual-link on page 3275](#)

## area

**Syntax**

```

area area-id {
    interface interface-name {
        passive;
        topology (ipv4-multicast | name) {
            disable;
        }
    }
    virtual-link neighbor-id router-id transit-area area-id {
        topology (ipv4-multicast | name) {
            disable;
        }
    }
}

```

**Hierarchy Level**

```

[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast |
  ipv4-multicast | ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
  (ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
  ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast |
  ipv4-multicast | ipv6-multicast)]

```

**Release Information**

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for the **realm** statement introduced in Junos OS Release 9.2.

Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description**

Specify the area identifier for this routing device to use when participating in OSPF routing. All routing devices in an area must use the same area identifier to establish adjacencies.

Specify multiple **area** statements to configure the routing device as an area border router. An area border router does not automatically summarize routes between areas. Use the **area-range** statement to configure route summarization. By definition, an area border router must be connected to the backbone area either through a physical link or through a virtual link. To create a virtual link, include the **virtual-link** statement.

To specify that the routing device is directly connected to the OSPF backbone, include the **area 0.0.0.0** statement.

All routing devices on the backbone must be contiguous. If they are not, use the **virtual-link** statement to create the appearance of connectivity to the backbone.

You can also configure any interface that belongs to one or more topologies to advertise the direct interface addresses without actually running OSPF on that interface. By default, OSPF must be configured on an interface in order for direct interface addresses to be advertised as interior routes.



**NOTE:** If you configure an interface with the `passive` statement, it applies to all the topologies to which the interface belongs. You cannot configure an interface as passive for only one specific topology and have it remain active for any other topologies to which it belongs.

|                                 |                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <b><i>area-id</i></b> —Area identifier. The identifier can be up to 32 bits. It is common to specify the area number as a simple integer or an IP address. Area number <b>0.0.0.0</b> is reserved for the OSPF backbone area.     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>OSPF Areas and Router Functionality Overview</i></li> <li>• <i>Understanding Multiple Address Families for OSPFv3</i></li> <li>• <a href="#">virtual-link on page 3275</a></li> </ul> |

## area-range

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <b>area-range</b> <i>network/mask-length</i> <exact> <override-metric <i>metric</i> > <restrict>;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>nssa</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>nssa</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>],</p> <p>[edit protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <b>nssa</b>],</p> <p>[edit protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>nssa</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>         | <p>(Area border routers only) For an area, summarize a range of IP addresses when sending summary link advertisements (within an area). To summarize multiple ranges, include multiple <b>area-range</b> statements.</p> <p>For a not-so-stubby area (NSSA), summarize a range of IP addresses when sending NSSA link-state advertisements. The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas. To specify multiple prefixes, include multiple <b>area-range</b> statements. All external routes learned within the area that do not fall into one of the prefixes are advertised individually to other areas.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>             | By default, area border routing devices do not summarize routes being sent from one area to other areas, but rather send all routes explicitly.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>             | <p><b>exact</b>—(Optional) Summarization of a route is advertised only when an exact match is made with the configured summary range.</p> <p><b>mask-length</b>—Number of significant bits in the network mask.</p> <p><b>network</b>—IP address. You can specify one or more IP addresses.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



**override-metric *metric***—(Optional) Override the metric for the IP address range and configure a specific metric value.

**restrict**—(Optional) Do not advertise the configured summary. This hides all routes that are contained within the summary, effectively creating a route filter.

**Range:** 1 through 16,777,215

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |                                                                                                                                  |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements</i></li></ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------|

## bandwidth-based-metrics

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>bandwidth-based-metrics {     bandwidth <i>value</i>;     metric <i>number</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf   ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</pre> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>         | <p>Specify a set of bandwidth threshold values and associated metric values for an OSPF interface or for a topology on an OSPF interface. When the bandwidth of an interface changes, Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>             | <p><b>bandwidth <i>value</i></b>—Specify the bandwidth threshold in bits per second.</p> <p><b>Range:</b> 9600 through 1,000,000,000,000,000</p> <p><b>metric <i>number</i></b>—Specify a metric value to associate with a specific bandwidth value.</p> <p><b>Range:</b> 1 through 65,535</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



**NOTE:** You must also configure a static metric value for the OSPF interface or topology with the metric statement. Junos OS uses this value to calculate the cost of a route from the OSPF interface or topology if the bandwidth for the interface is higher than of any bandwidth threshold values configured for bandwidth-based metrics.

|                                 |                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth</i></li><li>• <a href="#">metric on page 3244</a></li><li>• <i>Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth</i></li></ul> |

## bfd-liveness-detection (Protocols OSPF)

```
Syntax    bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            full-neighbors-only
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
```

| Hierarchy Level | Configuration Command                                                                                                                                                                                                                |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                 | [edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i> ],                                                                                              |
|                 | [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i> ],                                                |
|                 | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i> ],                                               |
|                 | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i> ], |
|                 | [edit protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i> ],                                                                                                                                         |
|                 | [edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i> ],                                                                                           |
|                 | [edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i> ],                                                                                          |
|                 | [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i> ]                                             |

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>detection-time threshold</b> and <b>transmit-interval threshold</b> options added in Junos OS Release 8.2.</p> <p>Support for logical systems introduced in Junos OS Release 8.3.</p> <p><b>no-adaptation</b> option introduced in Junos OS Release 9.0.</p> <p><b>no-adaptation</b> option introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 introduced in Junos OS Release 9.3.</p> <p>Support for OSPFv3 introduced in Junos OS Release 9.3 for EX Series switches.</p> <p><b>full-neighbors-only</b> option introduced in Junos OS Release 9.5.</p> <p><b>full-neighbors-only</b> option introduced in Junos OS Release 9.5 for EX Series switches.</p> |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

**authentication algorithm**, **authentication key-chain**, and **authentication loose-check** options introduced in Junos OS Release 9.6.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure bidirectional failure detection timers and authentication for OSPF.

The remaining statements are explained separately.

**Options**    **authentication algorithm** *algorithm-name* —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

**authentication key-chain** *key-chain-name* —Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

**detection-time threshold** *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**full-neighbors-only**—Establish BFD sessions only for OSPF neighbors in the full state. The default behavior is to establish BFD sessions for all OSPF neighbors.

**minimum-interval** *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

**Range:** 1 through 255,000 milliseconds

**minimum-receive-interval** *milliseconds*—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

**Range:** 1 through 255,000 milliseconds

**multiplier** *number*—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation**—Specify that BFD sessions should not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold** *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval** *milliseconds*—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established

a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

**Default:** **automatic**

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                                                                              |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Example: Configuring BFD for OSPF</i></li><li>• <i>Example: Configuring BFD Authentication for OSPF</i></li></ul> |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|

## dead-interval

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dead-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>virtual-link</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>virtual-link</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <b>virtual-link</b>],</p> <p>[edit protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>virtual-link</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Specify how long OSPF waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>seconds</b>—Interval to wait.</p> <p><b>Range:</b> 1 through 65,535 seconds</p> <p><b>Default:</b> Four times the hello interval—40 seconds (broadcast and point-to-point networks); 120 seconds (nonbroadcast multiple access (NBMA) networks)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring OSPF Timers</i></li> <li>• <i>Configuring RSVP and OSPF for LMP Peer Interfaces</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



- [hello-interval on page 3235](#)

## default-lsa

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>default-lsa {   default-metric metric;   metric-type type;   type-7; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <a href="#">nssa</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">nssa</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <a href="#">nssa</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">nssa</a>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <a href="#">nssa</a>],</p> <p>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">nssa</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <a href="#">nssa</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">nssa</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>On area border routers only, for a not-so-stubby area (NSSA), inject a default link-state advertisement (LSA) with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OSPF Areas and Router Functionality Overview</a></li> <li>• <a href="#">Example: Configuring OSPF Not-So-Stubby Areas</a></li> <li>• <a href="#">nssa on page 3249</a></li> <li>• <a href="#">stub on page 3266</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## default-metric

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>default-metric <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <a href="#">area area-id</a> nssa <a href="#">default-lsa</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <a href="#">area area-id</a> <a href="#">stub</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> nssa <a href="#">default-lsa</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">stub</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <a href="#">area area-id</a> nssa <a href="#">default-lsa</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <a href="#">area area-id</a> <a href="#">stub</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> nssa <a href="#">default-lsa</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">stub</a>],</p> <p>[edit protocols (ospf   ospf3) <a href="#">area area-id</a> nssa <a href="#">default-lsa</a>],</p> <p>[edit protocols (ospf   ospf3) <a href="#">area area-id</a> <a href="#">stub</a>],</p> <p>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> nssa <a href="#">default-lsa</a>],</p> <p>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">stub</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <a href="#">area area-id</a> nssa <a href="#">default-lsa</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <a href="#">area area-id</a> <a href="#">stub</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> nssa <a href="#">default-lsa</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">stub</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | On area border routing devices only, for a stub area, inject a default route with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b><i>metric</i></b>—Metric value.</p> <p><b>Range:</b> 1 through 16,777,215</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OSPF Areas and Router Functionality Overview</a></li> <li>• <a href="#">nssa on page 3249</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

- [stub on page 3266](#)

## disable (OSPF)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   <b>ospf3</b>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf <b>area</b> <i>area-id</i> <b>peer-interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <b>virtual-link</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   <b>ospf3</b>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>virtual-link</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols (<b>ospf</b>   <b>ospf3</b>)],</p> <p>[edit protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) <b>virtual-link</b>],</p> <p>[edit protocols ospf <b>area</b> <i>area-id</i> <b>peer-interface</b> <i>interface-name</i>],</p> <p>[edit protocols ospf <i>area</i> <i>area-id</i> <b>virtual-link</b> neighbor-id <i>router-id</i> transit-area <i>area-id</i>],</p> <p>[edit protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   <b>ospf3</b>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>virtual-link</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | <p>Disable OSPF, an OSPF interface, or an OSPF virtual link.</p> <p>By default, control packets sent to the remote end of a virtual link must be forwarded using the default topology. In addition, the transit area path consists only of links that</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

are in the default topology. You can disable a virtual link for a configured topology, but not for a default topology. Include the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id* topology *name*]** hierarchy level.



**NOTE:** If you disable the virtual link by including the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id*]** hierarchy level, you disable the virtual link for all topologies, including the default topology. You cannot disable the virtual link only in the default topology.

|                                 |                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | The configured object is enabled (operational) unless explicitly disabled.                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>OSPF Configuration Overview</i></li> <li>• <i>Configuring RSVP and OSPF for LMP Peer Interfaces</i></li> </ul> |

## domain-id

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>domain-id <i>domain-id</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i>   <i>ospf3</i>)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i>   <i>ospf3</i>)]</code>                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Specify a domain ID for a route. The domain ID identifies the OSPF domain from which the route originated.                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b><i>domain-id</i></b>—You can specify either an IP address or an IP address and a local identifier using the following format: <b><i>ip-address:local-identifier</i></b>. If you do not specify a local identifier with the IP address, the identifier is assumed to have a value of 0.</p> <p><b>Default:</b> If the router ID is not configured in the routing instance, the router ID is derived from an interface address belonging to the routing instance.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Routing Between PE and CE Routers in Layer 3 VPNs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                  |

## domain-vpn-tag

---

|                                 |                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>domain-vpn-tag <i>number</i>;</code>                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ( <a href="#">ospf</a>   <a href="#">ospf3</a> )],<br>[edit routing-instances <i>routing-instance-name</i> protocols ( <a href="#">ospf</a>   <a href="#">ospf3</a> )] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                        |
| <b>Description</b>              | Set a virtual private network (VPN) tag for OSPFv2 external routes generated by the provider edge (PE) routing device.                                                                                                                                                           |
| <b>Options</b>                  | <i>number</i> —VPN tag.                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Routing Between PE and CE Routers in Layer 3 VPNs</i></li></ul>                                                                                                                                                           |

## export (Protocols OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Apply one or more policies to routes being exported from the routing table into OSPF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding OSPF Routing Policy</i></li> <li>• <i>Import and Export Policies for Network Summaries Overview</i></li> <li>• <a href="#">import on page 3237</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## external-preference (Protocols OSPF)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>external-preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast</code><br><code>    ipv6-multicast)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</code><br><code>[edit protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</code><br><code>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast }</code><br><code>  ipv4-multicast   ipv6-multicast)]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Set the route preference for OSPF external routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b><i>preference</i></b> —Preference value.<br><b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ )<br><b>Default:</b> 150                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Controlling OSPF Route Preferences</i></li><li>• <a href="#">preference on page 3255</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



## graceful-restart (Protocols OSPF)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> graceful-restart {   disable;   helper-disable (standard   restart-signaling   both);   no-strict-lsa-checking;   notify-duration <i>seconds</i>;   restart-duration <i>seconds</i>; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the <b>no-strict-lsa-checking</b> statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode <b>standard</b>, <b>restart-signaling</b>, and <b>both</b> options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>         | <p>Configure graceful restart for OSPF.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>             | <p><b>disable</b>—Disable graceful restart for OSPF.</p> <p><b>helper-disable (standard   restart-signaling   both)</b>—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The <b>standard</b>, <b>restart-signaling</b>, and <b>both</b> options are only supported for OSPFv2. Specify <b>standard</b> to disable helper mode for standard graceful restart (based on RFC 3623). Specify <b>restart-signaling</b> to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify <b>both</b> to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p><b>Default:</b> Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p><b>no-strict-lsa-checking</b>—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p> |



**NOTE:** The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both

statements at the same time, the routing device displays a warning message when you enter the `show protocols (ospf | ospf3)` command.

.....  
**notify-duration seconds**—Estimated time needed to send out purged grace LSAs over all the interfaces.

**Range:** 1 through 3600 seconds

**Default:** 30 seconds

**restart-duration seconds**—Estimated time needed to reacquire a full OSPF neighbor from each area.

**Range:** 1 through 3600 seconds

**Default:** 180 seconds

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Example: Configuring Graceful Restart for OSPF</i></li><li>• <i>Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart</i></li><li>• <i>Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart</i></li><li>• <i>Example: Disabling Strict LSA Checking for OSPF Graceful Restart</i></li><li>• <i>Junos OS High Availability Library for Routing Devices</i></li></ul> |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## hello-interval (Protocols OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hello-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <a href="#">virtual-link</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <a href="#">virtual-link</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <a href="#">virtual-link</a>],</p> <p>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <a href="#">virtual-link</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <a href="#">interface</a> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Specify how often the routing device sends hello packets out the interface. The hello interval must be the same for all routing devices on a shared logical IP network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>seconds</b>—Time between hello packets, in seconds.</p> <p><b>Range:</b> 1 through 255 seconds</p> <p><b>Default:</b> 10 seconds (broadcast and point-to-point networks); 30 seconds (nonbroadcast multiple access [NBMA] networks)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring OSPF Timers</i></li> <li>• <i>Configuring RSVP and OSPF for LMP Peer Interfaces</i></li> <li>• <a href="#">dead-interval on page 3224</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## ignore-lsp-metrics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ignore-lsp-metrics;                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ospf <a href="#">traffic-engineering shortcuts</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">traffic-engineering shortcuts</a> ],<br>[edit protocols ospf <a href="#">traffic-engineering</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">traffic-engineering shortcuts</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for (OSPFv3) introduced in Junos OS Release 9.4.<br>Support for (OSPFv3) introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                      |
| <b>Description</b>              | Ignore RSVP LSP metrics in OSPF traffic engineering shortcut calculations.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Enabling OSPF Traffic Engineering Support</i></li></ul>                                                                                                                                                                                                                                                                                                                                                               |

## import (Protocols OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Filter OSPF routes from being added to the routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding OSPF Routing Policy</i></li> <li>• <i>Import and Export Policies for Network Summaries Overview</i></li> <li>• <a href="#">export on page 3231</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## inter-area-prefix-export

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>inter-area-prefix-export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <i>area</i> <i>area-id</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <i>area</i> <i>area-id</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ip4-unicast   ipv4-multicast   ipv6-multicast) <i>area</i> <i>area-id</i>],</code><br><code>[edit protocols ospf3 <i>area</i> <i>area-id</i>],</code><br><code>[edit protocols ospf3 <i>realm</i> (ip4-unicast   ipv4-multicast   ipv6-multicast) <i>area</i> <i>area-id</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>area</i> <i>area-id</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ip4-unicast   ipv4-multicast   ipv6-multicast) <i>area</i> <i>area-id</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 9.1 for EX Series switches.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Apply an export policy for OSPFv3 to specify which interarea prefix link-state advertisements (LSAs) are flooded into an area.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b><i>policy-name</i></b> —Name of a policy configured at the <code>[edit policy-options policy-statement <i>policy-name</i> term <i>term-name</i>]</code> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Import and Export Policies for Network Summaries Overview</i></li><li>• <a href="#">inter-area-prefix-import on page 3239</a></li><li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## inter-area-prefix-import

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>inter-area-prefix-import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>area</b> <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <b>area</b> <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>],</p> <p>[edit protocols ospf3 <b>area</b> <i>area-id</i>],</p> <p>[edit protocols ospf3 <b>realm</b> (ip4-unicast   ipv4-multicast   ipv6-multicast)], <b>area</b> <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <b>area</b> <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Apply an import policy for OSPFv3 to specify which routes learned from an area are used to generate interarea prefixes into other areas.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b><i>policy-name</i></b> —Name of a policy configured at the [edit policy-options policy-statement <i>policy-name</i> term <i>term-name</i> ] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Import and Export Policies for Network Summaries Overview</i></li> <li>• <a href="#">inter-area-prefix-export on page 3238</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## interface (Protocols OSPF)

**Syntax** interface *interface-name* {  
 disable;  
 authentication *key* <*key-id identifier*>;  
 bfd-liveness-detection {  
 authentication {  
 algorithm *algorithm-name*;  
 key-chain *key-chain-name*;  
 loose-check;  
 }  
 detection-time {  
 threshold *milliseconds*;  
 }  
 minimum-interval *milliseconds*;  
 minimum-receive-interval *milliseconds*;  
 transmit-interval {  
 threshold *milliseconds*;  
 minimum-interval *milliseconds*;  
 }  
 multiplier *number*;  
 }  
 dead-interval *seconds*;  
 demand-circuit;  
 hello-interval *seconds*;  
 ipsec-sa *name*;  
 interface-type *type*;  
 ldp-synchronization {  
 disable;  
 hold-time *seconds*;  
 }  
 metric *metric*;  
 neighbor *address* <*eligible*>;  
 no-interface-state-traps;  
 passive;  
 poll-interval *seconds*;  
 priority *number*;  
 retransmit-interval *seconds*;  
 te-metric *metric*;  
 topology (ipv4-multicast | *name*) {  
 metric *metric*;  
 }  
 transit-delay *seconds*;  
}

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols (ospf | ospf3) *area area-id*],  
 [edit logical-systems *logical-system-name* protocols ospf3 *realm* (ipv4-unicast |  
 ipv4-multicast | ipv6-multicast) *area area-id*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 (ospf | ospf3) *area area-id*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 ospf3 *realm* (ipv4-unicast | ipv4-multicast | ipv6-multicast) *area area-id*],  
 [edit protocols (ospf | ospf3) *area area-id*],  
 [edit protocols ospf3 *realm* (ipv4-unicast | ipv4-multicast | ipv6-multicast) *area area-id*],



[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) *area area-id*],  
 [edit routing-instances *routing-instance-name* protocols ospf3 *realm* (ipv4-unicast |  
 ipv4-multicast | ipv6-multicast) *area area-id*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Support for the **topology** statement introduced in Junos OS Release 9.0.  
 Support for the **topology** statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Support for the **realm** statement introduced in Junos OS Release 9.2.  
 Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.  
 Support for the **no-interface-state-traps** statement introduced in Junos OS Release 10.3.  
 This statement is supported only for OSPFv2.  
 Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Enable OSPF routing on a routing device interface.

You must include at least one **interface** statement in the configuration to enable OSPF on the routing device.

**Options** *interface-name*—Name of the interface. Specify the interface by IP address or interface name for OSPFv2, or only the interface name for OSPFv3. Using both the interface name and IP address of the same interface produces an invalid configuration. To configure all interfaces, you can specify **all**. Specifying a particular interface and **all** produces an invalid configuration.



**NOTE:** For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as *interface-name*.

The remaining statements are explained separately.



**NOTE:** You cannot run both OSPF and ethernet-tcc encapsulation between two Juniper Networks routing devices.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

**Related Documentation**

- *OSPF Configuration Overview*
- *Example: Configuring Multitopology Routing Based on Applications*
- *Example: Configuring Multitopology Routing Based on a Multicast Source*
- *Example: Configuring Multiple Address Families for OSPFv3*
- *neighbor*

## interface-type (Protocols OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interface-type (nbma   p2mp   p2p);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-multicast   ipv4-unicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-multicast   ipv4-unicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm (ipv4-multicast   ipv4-unicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-multicast   ipv4-unicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 for interface type <b>p2p</b> only introduced in Junos OS Release 9.4. You cannot configure other interface types for OSPFv3.</p> <p>Support for OSPFv3 for interface type <b>p2p</b> only introduced in Junos OS Release 9.4 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Specify the type of interface.</p> <p>By default, the software chooses the correct interface type based on the type of physical interface. Therefore, you should never have to set the interface type. The exception to this is for NBMA interfaces, which default to an interface type of point-to-multipoint. To have these interfaces explicitly run in Nonbroadcast multiaccess (NBMA) mode, configure the <b>nbma</b> interface type, using the IP address of the local ATM interface.</p> <p>In Junos OS Release 9.3 and later, a point-to-point interface can be an Ethernet interface without a subnet.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | The software chooses the correct interface type based on the type of physical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>nbma</b> (OSPFv2 only)—Nonbroadcast multiaccess (NBMA) interface.</p> <p><b>p2mp</b> (OSPFv2 only)—Point-to-multipoint interface.</p> <p><b>p2p</b>—Point-to-point interface.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

- Related Documentation**
- *About OSPF Interfaces*
  - *Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network*

## lsp-metric-into-summary

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>lsp-metric-into-summary;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <a href="#">traffic-engineering shortcuts</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <a href="#">traffic-engineering shortcuts</a>],</p> <p>[edit protocols (ospf   ospf3) <a href="#">traffic-engineering shortcuts</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <a href="#">traffic-engineering shortcuts</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 (<b>ospf3</b>) introduced in Junos OS Release 9.4.</p> <p>Support for OSPFv3 (<b>ospf3</b>) introduced in Junos OS Release 9.4 for EX Series switches.</p>                                                                                                                                                                                                                      |
| <b>Description</b>              | Advertise the LSP metric in summary LSAs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>OSPF Support for Traffic Engineering</i></li> <li>• <i>Example: Enabling OSPF Traffic Engineering Support</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                       |

## metric (Protocols OSPF Interface)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>metric <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <b>area</b> <i>area-id</i> sham-link-remote],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast   <i>name</i>)],</p> <p>[edit protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf <b>area</b> <i>area-id</i> sham-link-remote],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast   <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | <p>Specify the cost of an OSPF interface. The cost is a routing metric that is used in the link-state calculation.</p> <p>To set the cost of routes exported into OSPF, configure the appropriate routing policy.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>             | <p><b>metric</b>—Cost of the route.</p> <p><b>Range:</b> 1 through 65,535</p> <p><b>Default:</b> By default, the cost of an OSPF route is calculated by dividing the reference-bandwidth value by the bandwidth of the physical interface. Any specific value you configure for the <b>metric</b> overrides the default behavior of using the reference-bandwidth value to calculate the cost of the route for that interface.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Controlling the Cost of Individual OSPF Network Segments</i></li><li>• <i>Example: Configuring OSPFv2 Sham Links</i></li><li>• <i>Example: Configuring Multitopology Routing Based on Applications</i></li><li>• <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i></li><li>• <a href="#">bandwidth-based-metrics on page 3218</a></li><li>• <a href="#">reference-bandwidth on page 3259</a></li></ul> |

## metric-type

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric-type type;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>nssa default-lsa</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)) area <i>area-id</i> nssa<b>default-lsa</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>nssa default-lsa</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)) area <i>area-id</i> nssa <b>default-lsa</b>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <b>nssa default-lsa</b>],</p> <p>[edit protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)) area <i>area-id</i> nssa <b>default-lsa</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>nssa default-lsa</b>],</p> <p>[edit routing-instances <i>routing-instances</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)) area <i>area-id</i> nssa <b>default-lsa</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Specify the external metric type for the default LSA.</p> <p>The configured metric determines the method used to compute the cost to a destination:</p> <ul style="list-style-type: none"> <li>• The Type 1 external metric is equivalent to the link-state metric. The path cost uses the advertised external path cost and the path cost to the AS boundary router (the route is equal to the sum of all internal costs and the external cost).</li> <li>• The Type 2 external metric uses the cost assigned by the AS boundary router (the route is equal to the external cost alone). By default, OSPF uses the Type 2 external metric.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <b>type</b> —Metric type: 1 or 2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>OSPF Areas and Router Functionality Overview</i></li> <li>• <i>Example: Configuring OSPF Not-So-Stubby Areas</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## no-nssa-abr

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-nssa-abr;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Disable exporting Type 7 link-state advertisements into not-so-stubby-areas (NSSAs) for an autonomous system boundary router (ASBR) or an area border router (ABR).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring OSPF Not-So-Stubby Areas</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## no-rfc-1583

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-rfc-1583;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ( <b>ospf</b>   <b>ospf3</b> )],<br>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast  <br>ipv4-multicast   ipv6-multicast)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>( <b>ospf</b>   <b>ospf3</b> )],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],<br>[edit protocols ( <b>ospf</b>   <b>ospf3</b> )],<br>[edit protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],<br>[edit routing-instances <i>routing-instance-name</i> protocols ( <b>ospf</b>   <b>ospf3</b> )],<br>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast  <br>ipv4-multicast   ipv6-multicast)] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Disable compatibility with RFC 1583, <i>OSPF Version 2</i> . If the same external destination is advertised by AS boundary routers that belong to different OSPF areas, disabling compatibility with RFC 1583 can prevent routing loops.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default</b>                  | Compatibility with RFC 1583 is enabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control-level—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Disabling OSPFv2 Compatibility with RFC 1583</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



## nssa

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>nssa {   area-range network/mask-length &lt;restrict&gt; &lt;exact&gt; &lt;override-metric metric&gt;;   default-lsa {     default-metric metric;     metric-type type;     type-7;   }   (no-summaries   summaries); }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems logical-system-name protocols (ospf   ospf3) area area-id], [edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast     ipv4-multicast   ipv6-multicast)], [edit logical-systems logical-system-name routing-instances routing-instance-name protocols   (ospf   ospf3) area area-id], [edit logical-systems logical-system-name routing-instances routing-instance-name protocols   ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit protocols (ospf   ospf3) area area-id], [edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit routing-instances routing-instance-name protocols (ospf   ospf3) area area-id], [edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast     ipv4-multicast   ipv6-multicast)]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Configure a not-so-stubby area (NSSA). An NSSA allows external routes to be flooded within the area. These routes are then leaked into other areas.</p> <p>You cannot configure an area as being both a stub area and an NSSA.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>OSPF Areas and Router Functionality Overview</i></li> <li>• <i>Example: Configuring OSPF Not-So-Stubby Areas</i></li> <li>• <a href="#">stub on page 3266</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## ospf

---


|                                 |                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ospf { ... }                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                      |
| <b>Description</b>              | Enable OSPF routing on the routing device.<br><br>You must include the <b>ospf</b> statement to enable OSPF on the routing device.                                                                                                                                  |
| <b>Default</b>                  | OSPF is disabled on the routing device.                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>OSPF Configuration Overview</i></li><li>• <i>[edit protocols ospf] Hierarchy Level</i></li></ul>                                                                                                                         |

## ospf3

---

|                                 |                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ospf3 { ... }                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                      |
| <b>Description</b>              | Enable OSPFv3 routing on the routing device.<br><br>You must include the <b>ospf3</b> statement to enable OSPFv3.                                                                                                                                                   |
| <b>Default</b>                  | OSPFv3 is disabled on the routing device.                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>OSPF Configuration Overview</i></li> <li>• <i>[edit protocols ospf3] Hierarchy Level</i></li> </ul>                                                                                                                     |

## overload (Protocols OSPF)

|                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                | <pre>overload {     timeout <i>seconds</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>                                                                                                                                                                                       | <p>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   <b>ospf3</b>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf topology (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   <b>ospf3</b>)],</p> <p>[edit logical systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf topology (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<b>ospf</b>   <b>ospf3</b>)],</p> <p>[edit protocols ospf topology (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   <b>ospf3</b>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf topology (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> |
| <b>Release Information</b>                                                                                                                                                                                   | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                                                                                                                                                                           | <p>Configure the local routing device so that it appears to be overloaded. You might do this when you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <div>  <p><b>NOTE:</b> Traffic destined to directly attached interfaces continues to reach the routing device.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                                                                                                                                                                               | <p><b>timeout <i>seconds</i></b>—(Optional) Number of seconds at which the overloading is reset. If no timeout interval is specified, the routing device remains in overload state until the <b>overload</b> statement is deleted or a timeout is set.</p> <p><b>Range:</b> 60 through 1800 seconds</p> <p><b>Default:</b> 0 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |



**NOTE:** Multitopology Routing does not support the timeout option.

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring OSPF to Make Routing Devices Appear Overloaded*
- *Example: Configuring Multitopology Routing Based on Applications*
- *Example: Configuring Multitopology Routing Based on a Multicast Source*

## passive (Protocols OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> passive {     traffic-engineering {         remote-node-id address;     } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>traffic-engineering</b> and <b>remote-node-id address</b> statements introduced in Junos OS Release 8.0.</p> <p><b>traffic-engineering</b> and <b>remote-node-id address</b> statements introduced in Junos OS Release 8.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Advertise the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run.</p> <p>To configure an interface in OSPF passive traffic engineering mode, include the <b>traffic-engineering</b> statement. Configuring OSPF passive traffic engineering mode enables the dynamic discovery of OSPF AS boundary routers.</p> <p>Enable OSPF on an interface by including the <b>interface</b> statement at the [edit protocols (ospf   ospf3) area <i>area-id</i>] or the [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i>] hierarchy levels. Disable it by including the <b>disable</b> statement. To prevent OSPF from running on an interface, include the <b>passive</b> statement. These three states are mutually exclusive.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- Related Documentation**
- *Example: Configuring a Passive OSPF Interface*
  - *Example: Configuring OSPF Passive Traffic Engineering Mode*
  - [disable on page 3228](#)

## preference (Protocols OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Set the route preference for OSPF internal routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b><i>preference</i></b>—Preference value.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 10</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Controlling OSPF Route Preferences</i></li> <li>• <a href="#">external-preference on page 3232</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## prefix-export-limit (Protocols OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>prefix-export-limit <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf topology (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf topology (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit protocols ospf topology (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf topology (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure a limit to the number of prefixes exported into OSPF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b><i>number</i></b>—Prefix limit.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> None</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Limiting the Number of Prefixes Exported to OSPF</i></li> <li>• <i>Example: Configuring Multitopology Routing Based on Applications</i></li> <li>• <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |



## priority (Protocols OSPF)


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)) area <i>area-id</i> <b>interface</b> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Specify the routing device's priority for becoming the designated routing device. The routing device that has the highest priority value on the logical IP network or subnet becomes the network's designated router. You must configure at least one routing device on each logical IP network or subnet to be the designated router. You also should specify a routing device's priority for becoming the designated router on point-to-point interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>number</b>—Routing device's priority for becoming the designated router. A priority value of 0 means that the routing device never becomes the designated router. A value of 1 means that the routing device has the least chance of becoming a designated router.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 128</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>OSPF Designated Router Overview</i></li> <li>• <i>Example: Controlling OSPF Designated Router Election</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## realm

---

|                                 |                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>realm (ipv4-unicast   ipv4-multicast   ipv6-unicast) {<br/>    area area-id {<br/>        interface interface-name;<br/>    }<br/>}</pre>                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>ospf3</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>ospf3</b> ],<br>[edit protocols <b>ospf3</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>ospf3</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2.<br>Statement introduced in Junos OS Release 9.2 for EX Series switches.                                                                                                                                                                                                       |
| <b>Description</b>              | Configure OSPFv3 to advertise address families other than unicast IPv6. Junos OS maps each address family you configure to a separate realm with its own set of neighbors and link-state database.                                                                                                                          |
| <b>Options</b>                  | <b>ipv4-unicast</b> —Configure a realm for IPv4 unicast routes.<br><br><b>ipv4-multicast</b> —Configure a realm for IPv4 multicast routes.<br><br><b>ipv6-multicast</b> —Configure a realm for IPv6 multicast routes.<br><br>The remaining statements are explained separately.                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multiple Address Families for OSPFv3</i></li></ul>                                                                                                                                                                                                          |

## reference-bandwidth (Protocols OSPF)

|                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                               | <code>reference-bandwidth <i>reference-bandwidth</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                      | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                  | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                          | <p>Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula:</p> $\text{cost} = \text{ref-bandwidth} / \text{bandwidth}$                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                              | <p><b><i>reference-bandwidth</i></b>—Reference bandwidth, in bits per second.</p> <p><b>Range:</b> 9600 through 1,000,000,000,000 bits</p> <p><b>Default:</b> 100 Mbps (100,000,000 bits)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <div>  <p><b>NOTE:</b> The default behavior is to use the reference-bandwidth value to calculate the cost of OSPF interfaces. You can override this behavior for any OSPF interface by configuring a specific cost with the metric statement.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                             | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                | <ul style="list-style-type: none"> <li>• <a href="#">Example: Controlling the Cost of Individual OSPF Network Segments</a></li> <li>• <a href="#">metric on page 3244</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## retransmit-interval (OSPF)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>retransmit-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> virtual-link], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> virtual-link], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>], [edit protocols (ospf   ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf   ospf3) area <i>area-id</i> virtual-link], [edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> virtual-link], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</pre> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>         | Specify how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements (LSAs) to an interface's neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>             | <p><b>seconds</b>—Interval to wait.</p> <p><b>Range:</b> 1 through 65,535 seconds</p> <p><b>Default:</b> 5 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |



**NOTE:** You must configure LSA retransmit intervals to be equal to or greater than 3 seconds to avoid triggering a retransmit trap, because Junos OS delays LSA acknowledgments by up to 2 seconds.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring OSPF Timers*
- *Configuring RSVP and OSPF for LMP Peer Interfaces*

## rib-group (Protocols OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rib-group group-name;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],<br/> [edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],<br/> [edit protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],<br/> [edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],<br/> [edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],<br/> [edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Install routes learned from OSPF routing instances into routing tables in the OSPF routing table group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b>group-name</b> —Name of the routing table group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li> <li>• <i>Example: Importing Direct and Static Routes Into a Routing Instance</i></li> <li>• <i>Understanding Multiprotocol BGP</i></li> <li>• <a href="#">interface-routes on page 3474</a></li> <li>• <a href="#">rib-group on page 3509</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## route-type-community

---

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>route-type-community (iana   vendor);</code>                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i>   <i>ospf3</i>)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i>   <i>ospf3</i>)]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                 |
| <b>Description</b>              | Specify an extended community value to encode the OSPF route type. Each extended community is coded as an eight-octet value. This statement sets the most significant bit to either an IANA or vendor-specific route type.                                         |
| <b>Options</b>                  | <b>iana</b> —Encode a route type with the value <b>0x0306</b> . This is the default value.<br><b>vendor</b> —Encode the route type with the value <b>0x8000</b> .                                                                                                  |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Routing Between PE and CE Routers in Layer 3 VPNs</i></li></ul>                                                                                                                                             |

## shortcuts (Protocols OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | shortcuts {<br>lsp-metric-into-summary;<br>}                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <a href="#">traffic-engineering</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <a href="#">traffic-engineering</a> ],<br>[edit protocols (ospf   ospf3) <a href="#">traffic-engineering</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <a href="#">traffic-engineering</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for OSPFv3 ( <b>ospf3</b> ) introduced in Junos OS Release 9.4.<br>Support for OSPFv3 ( <b>ospf3</b> ) introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                              |
| <b>Description</b>              | Configure OSPF to use MPLS label-switched paths (LSPs) as shortcut next hops. By default, shortcut routes calculated through OSPFv2 are installed in the <b>inet.3</b> routing table, and shortcut routes calculated through OSPFv3 are installed in the <b>inet6.3</b> routing table.                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Enabling OSPF Traffic Engineering Support</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                         |

## spf-options (Protocols OSPF)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>spf-options {<br/>    delay <i>milliseconds</i>;<br/>    holddown <i>milliseconds</i>;<br/>    rapid-runs <i>number</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   <b>ospf3</b>)],<br/>[edit logical-systems <i>logical-system-name</i> protocols ospf topology (default   ipv4-multicast<br/>  <i>name</i>)],<br/>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast  <br/>ipv4-multicast   ipv6-multicast)],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>(<b>ospf</b>   <b>ospf3</b>)],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>ospf topology (default   ipv4-multicast   <i>name</i>)],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],<br/>[edit protocols (<b>ospf</b>   <b>ospf3</b>)],<br/>[edit protocols ospf topology (default   ipv4-multicast   <i>name</i>)],<br/>[edit protocols ospf3 <b>realm</b> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],<br/>[edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   <b>ospf3</b>)],<br/>[edit routing-instances <i>routing-instance-name</i> protocols ospf topology (default  <br/>ipv4-multicast   <i>name</i>)],<br/>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <b>realm</b> (ipv4-unicast  <br/>ipv4-multicast   ipv6-multicast)]</pre> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>         | <p>Configure options for running the shortest-path-first (SPF) algorithm. You can configure the following:</p> <ul style="list-style-type: none"><li>• A delay for when to run the SPF algorithm after a network topology change is detected.</li><li>• The maximum number of times the SPF algorithm can run in succession.</li><li>• A hold-down interval after the SPF algorithm runs the maximum number of times.</li></ul> <p>Running the SPF algorithm is usually the beginning of a series of larger system-wide events. For example, the SPF algorithm can lead to interior gateway protocol (IGP) prefix changes, which then lead to BGP nexthop resolution changes. Consider what happens if there are rapid link changes in the network. The local routing device can become overwhelmed. This is why it sometimes makes sense to throttle the scheduling of the SPF algorithm.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



|                                 |                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <b>delay <i>milliseconds</i></b> —Time interval between the detection of a topology change and when the SPF algorithm runs.<br><b>Range:</b> 50 through 8000 milliseconds<br><b>Default:</b> 200 milliseconds                                                                                   |
|                                 | <b>holddown <i>milliseconds</i></b> —Time interval to hold down, or to wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession.<br><b>Range:</b> 2000 through 20,000 milliseconds<br><b>Default:</b> 5000 milliseconds |
|                                 | <b>rapid-runs <i>number</i></b> —Maximum number of times the SPF algorithm can run in succession. After the maximum is reached, the hold down interval begins.<br><b>Range:</b> 1 through 10<br><b>Default:</b> 3                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring SPF Algorithm Options for OSPF</i></li><li>• <i>Example: Configuring Multitopology Routing Based on Applications</i></li><li>• <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i></li></ul>          |

## stub

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | stub <default-metric <i>metric</i> > <(no-summaries   summaries)>;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <i>area area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <i>area area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (ospf   ospf3) <i>area area-id</i>],</p> <p>[edit protocols ospf3 <i>realm</i> (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <i>area area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Specify that this area not be flooded with AS external link-state advertisements (LSAs). You must include the <b>stub</b> statement when configuring all routing devices that are in the stub area.</p> <p>The backbone cannot be configured as a stub area.</p> <p>You cannot configure an area to be both a stub area and a not-so-stubby area (NSSA).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>no-summaries</b>—(Optional) Do not advertise routes into the stub area. If you include the <b>default-metric</b> option, only the default route is advertised.</p> <p><b>summaries</b>—(Optional) Flood summary LSAs into the stub area.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>OSPF Areas and Router Functionality Overview</i></li><li>• <i>Example: Configuring OSPF Stub and Totally Stubby Areas</i></li><li>• <a href="#">nssa on page 3249</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## summaries

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (summaries   no-summaries);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <a href="#">area area-id nssa</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id nssa</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <a href="#">area area-id nssa</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id nssa</i>],</p> <p>[edit protocols (ospf   ospf3) <a href="#">area area-id nssa</a>],</p> <p>[edit protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast)] area <i>area-id nssa</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <a href="#">area area-id nssa</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <a href="#">realm</a> (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id nssa</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure whether or not area border routers advertise summary routes into an not-so-stubby area (NSSA):</p> <ul style="list-style-type: none"> <li>• <b>summaries</b>—Flood summary link-state advertisements (LSAs) into the NSSA.</li> <li>• <b>no-summaries</b>—Prevent area border routers from advertising summaries into an NSSA. If <b>default-metric</b> is configured for an NSSA, a Type 3 LSA is injected into the area by default.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>OSPF Areas and Router Functionality Overview</i></li> <li>• <i>Example: Configuring OSPF Not-So-Stubby Areas</i></li> <li>• <a href="#">nssa on page 3249</a></li> <li>• <a href="#">stub on page 3266</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## traceoptions (Protocols OSPF)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols (<i>ospf</i>   <i>ospf3</i>)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast     ipv4-multicast   ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   (<i>ospf</i>   <i>ospf3</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   ospf3 <i>realm</i> (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit protocols (<i>ospf</i>   <i>ospf3</i>)], [edit protocols ospf3 <i>realm</i> (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i>   <i>ospf3</i>)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast     ipv4-multicast   ipv6-multicast)]</pre> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>         | <p>Configure OSPF protocol-level tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



**NOTE:** The **traceoptions** statement is not supported on QFabric systems.

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b> | The default OSPF protocol-level tracing options are those inherited from the routing protocols <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b> | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place OSPF tracing output in the file <b>ospf-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> |

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

#### OSPF Tracing Flags

- **database-description**—Database description packets, which are used in synchronizing the OSPF and OSPFv3 topological database.
- **error**—OSPF and OSPFv3 error packets.
- **event**—OSPF and OSPFv3 state transitions.
- **flooding**—Link-state flooding packets.
- **graceful-restart**—Graceful-restart events.
- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable.
- **ldp-synchronization**—Synchronization events between OSPF and LDP.
- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database.
- **lsa-analysis**—Link-state analysis. Specific to the Juniper Networks implementation of OSPF, Junos OS performs LSA analysis before running the shortest-path-first (SPF) algorithm. LSA analysis helps to speed the calculations performed by the SPF algorithm.
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database.
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **on-demand**—Trace demand circuit extensions.
- **packet-dump**—Content of selected packet types.
- **packets**—All OSPF packets.
- **restart-signaling**—(OSPFv2 only) Restart-signaling graceful restart events.
- **spf**—Shortest-path-first (SPF) calculations.

#### Global Tracing Flags

- **all**—All tracing operations.
- **general**—A combination of the **normal** and **route** trace operations.
- **normal**—All normal operations. If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions.
- **route**—Routing table changes.
- **state**—State transitions.
- **task**—Routing protocol task processing.
- **timer**—Routing protocol timer processing.

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.                                  |
|                                 | routing-control and trace-control—To add this statement to the configuration.                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Tracing OSPF Protocol Traffic</i></li></ul> |

## traffic-engineering (OSPF)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traffic-engineering {   &lt;advertise-unnumbered-interfaces&gt;;   &lt;credibility-protocol-preference&gt;;   ignore-lsp-metrics;   multicast-rpf-routes;   no-topology;   shortcuts {     lsp-metric-into-summary;   } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols ( <b>ospf</b>   ospf3)],<br>[edit protocols ( <b>ospf</b>   ospf3)]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>multicast-rpf-routes</b> option introduced in Junos OS Release 7.5.</p> <p><b>advertise-unnumbered-interfaces</b> option introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 (<b>ospf3</b>) introduced in Junos OS Release 9.4.</p> <p>Support for OSPFv3 (<b>ospf3</b>) introduced in Junos OS Release 9.4 for EX Series switches.</p> <p><b>credibility-protocol-preference</b> statement introduced in Junos OS Release 9.4.</p> <p><b>credibility-protocol-preference</b> statement introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>         | Enable the OSPF traffic engineering features.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>             | Traffic engineering support is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>             | <p><b>advertise-unnumbered-interfaces</b>—(Optional) (OSPFv2 only) Include the link-local identifier in the link-local traffic-engineering link-state advertisement. This statement must be included on both ends of an unnumbered link to allow an ingress LER to update the link in its traffic engineering database and use it for CSPF calculations. The link-local identifier is then used by RSVP to signal unnumbered interfaces as defined in RFC 3477.</p> <p><b>credibility-protocol-preference</b>—(Optional) (OSPFv2 only) Use the configured preference value for OSPF routes to calculate the traffic engineering database credibility value used to select IGP routes. Use this statement to override the default behavior, in which the traffic engineering database prefers IS-IS routes even if OSPF routes are configured with a lower, that is, preferred, preference value. For example, OSPF routes have a default preference value of 10, whereas IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502, whereas IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.</p> |

**multicast-rpf-routes**—(Optional) (OSPFv2 only) Install routes for multicast RPF checks into the **inet.2** routing table. The **inet.2** routing table consists of unicast routes used for multicast RPF lookup. RPF is an antispoofing mechanism used to check whether the packet is coming in on an interface that is also sending data back to the packet source.



**NOTE:** You must enable OSPF traffic engineering shortcuts to use the **multicast-rpf-routes** statement. You must not allow LSP advertisements into OSPF when configuring the **multicast-rpf-routes** statement.

**no-topology**—(Optional) (OSPFv2 only) Disable the dissemination of the link-state topology information.

The remaining statements are explained separately.



**CAUTION:** When the OSPF traffic engineering configuration is considerably modified, the routing table entries are deleted and the routing table is recreated. Changes to configuration that can cause this behavior include enabling or disabling:

- Traffic engineering shortcuts
- IGP shortcuts
- LDP tunneling
- Multiprotocol LSP
- Advertise summary metrics
- Multicast RPF routes

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Enabling OSPF Traffic Engineering Support*



## transit-delay (OSPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>transit-delay <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> <b>virtual-link</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> <b>virtual-link</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> <b>virtual-link</b>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)] area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> <b>virtual-link</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> <b>interface</b> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Set the estimated time required to transmit a link-state update on the interface. When calculating this time, make sure to account for transmission and propagation delays.</p> <p>You should never have to modify the transit delay time.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b><i>seconds</i></b>—Estimated time, in seconds.</p> <p><b>Range:</b> 1 through 65,535 seconds</p> <p><b>Default:</b> 1 second</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring OSPF Timers</i></li> <li>• <i>Configuring RSVP and OSPF for LMP Peer Interfaces</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## type-7

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>type-7;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> nssa <a href="#">default-lsa</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> nssa <a href="#">default-lsa</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> nssa <a href="#">default-lsa</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> nssa <a href="#">default-lsa</a>],</code><br><code>[edit protocols (ospf   ospf3) area <i>area-id</i> nssa <a href="#">default-lsa</a>],</code><br><code>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> nssa <a href="#">default-lsa</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> nssa <a href="#">default-lsa</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i> nssa <a href="#">default-lsa</a>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.<br>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Flood Type 7 default link-state advertisements (LSAs) if the <b>no-summaries</b> statement is configured.</p> <p>By default, when the <b>no-summaries</b> statement is configured, a Type 3 LSA is injected into not-so-stubby areas (NSSAs) for Junos OS Release 5.0 and later. To support backward compatibility with earlier Junos OS releases, include the <b>type-7</b> statement. This statement enables NSSA ABRs to advertise a Type 7 default LSA into the NSSA if you have also included the <b>no-summaries</b> statement in the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>OSPF Areas and Router Functionality Overview</i></li><li>• <i>Example: Configuring OSPF Not-So-Stubby Areas</i></li><li>• <a href="#">no-summaries on page 3267</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## virtual-link

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i> {   disable;   authentication key &lt;key-id identifier&gt;;   dead-interval <i>seconds</i>;   hello-interval <i>seconds</i>;   ipsec-sa <i>name</i>;   retransmit-interval <i>seconds</i>;   transit-delay <i>seconds</i>; }</pre>                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   ospf <b>area</b> <i>area-id</i>], [edit protocols (ospf   ospf3) <b>area</b> <i>area-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf <b>area</b> <i>area-id</i>]</pre>                                                                                                                                       |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>For backbone areas only, create a virtual link to use in place of an actual physical link. All area border routers and other routing devices on the backbone must be contiguous. If this is not possible and there is a break in OSPF connectivity, use virtual links to create connectivity to the OSPF backbone. When configuring virtual links, you must configure links on the two routing devices that form the end points of the link, and both of these routing devices must be area border routers. You cannot configure links through stub areas.</p> |
| <b>Options</b>                  | <p><b>neighbor-id <i>router-id</i></b>—IP address of the routing device at the remote end of the virtual link.</p> <p><b>transit-area <i>area-id</i></b>—Area identifier of the area through which the virtual link transits. Virtual links are not allowed to transit the backbone area.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>OSPF Areas and Router Functionality Overview</i></li> <li>• <i>Example: Configuring OSPF Virtual Links</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                 |



# Administration

- [Routine Monitoring on page 3277](#)
- [Operational Commands on page 3280](#)

## Routine Monitoring

- [Monitoring OSPF Routing Information on page 3277](#)

### Monitoring OSPF Routing Information

Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to monitor OSPF routing information on routing devices.

Action

To view OSPF routing information in the J-Web interface, select **Monitor > Routing > OSPF Information**.

To view OSPF routing information in the CLI, enter the following CLI commands:

- `show ospf neighbor`
- `show ospf interface`
- `show ospf statistics`

Meaning

[Table 343 on page 3277](#) summarizes key output fields in the OSPF routing display in the J-Web interface.

Table 343: Summary of Key OSPF Routing Output Fields

| Field           | Values                              | Additional Information |
|-----------------|-------------------------------------|------------------------|
| OSPF Interfaces |                                     |                        |
| Interface       | Name of the interface running OSPF. |                        |

Table 343: Summary of Key OSPF Routing Output Fields (*continued*)

| Field               | Values                                                                                                                                                                     | Additional Information                                                                                                                                                                          |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| State               | State of the interface: <b>BDR</b> , <b>Down</b> , <b>DR</b> , <b>DROther</b> , <b>Loop</b> , <b>PtToPt</b> , or <b>Waiting</b> .                                          | The <b>Down</b> state, indicating that the interface is not functioning, and <b>PtToPt</b> state, indicating that a point-to-point connection has been established, are the most common states. |
| Area                | Number of the area that the interface is in.                                                                                                                               |                                                                                                                                                                                                 |
| DR ID               | Address of the area's designated device.                                                                                                                                   |                                                                                                                                                                                                 |
| BDR ID              | Address of the area's backup designated device.                                                                                                                            |                                                                                                                                                                                                 |
| Neighbors           | Number of neighbors on this interface.                                                                                                                                     |                                                                                                                                                                                                 |
| Adjacency Count     | Number of devices in the area using the same area identifier.                                                                                                              |                                                                                                                                                                                                 |
| Stub Type           | The areas into which OSPF does not flood AS external advertisements                                                                                                        |                                                                                                                                                                                                 |
| Passive Mode        | In this mode the interface is present on the network but does not transmit or receive packets.                                                                             |                                                                                                                                                                                                 |
| Authentication Type | The authentication scheme for the backbone or area.                                                                                                                        |                                                                                                                                                                                                 |
| Interface Address   | The IP address of the interface.                                                                                                                                           |                                                                                                                                                                                                 |
| Address Mask        | The subnet mask or address prefix.                                                                                                                                         |                                                                                                                                                                                                 |
| MTU                 | The maximum transmission unit size.                                                                                                                                        |                                                                                                                                                                                                 |
| Interface Cost      | The path cost used to calculate the root path cost from any given LAN segment is determined by the total cost of each link in the path.                                    |                                                                                                                                                                                                 |
| Hello Interval      | How often the routing device sends hello packets out of the interface.                                                                                                     |                                                                                                                                                                                                 |
| Dead Interval       | The interval during which the routing device receives no hello packets from the neighbor.                                                                                  |                                                                                                                                                                                                 |
| Retransmit Interval | The interval for which the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements to an interface's neighbors. |                                                                                                                                                                                                 |
| OSPF Statistics     |                                                                                                                                                                            |                                                                                                                                                                                                 |
| Packets tab         |                                                                                                                                                                            |                                                                                                                                                                                                 |

Table 343: Summary of Key OSPF Routing Output Fields (*continued*)

| Field                    | Values                                                                                        | Additional Information                                                                                                                                                                                                                                                                                         |
|--------------------------|-----------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Sent                     | Displays the total number of packets sent.                                                    |                                                                                                                                                                                                                                                                                                                |
| Received                 | Displays the total number of packets received.                                                |                                                                                                                                                                                                                                                                                                                |
| <b>Details tab</b>       |                                                                                               |                                                                                                                                                                                                                                                                                                                |
| Flood Queue Depth        | Number of entries in the extended queue.                                                      |                                                                                                                                                                                                                                                                                                                |
| Total Retransmits        | Number of retransmission entries enqueued.                                                    |                                                                                                                                                                                                                                                                                                                |
| Total Database Summaries | Total number of database description packets.                                                 |                                                                                                                                                                                                                                                                                                                |
| <b>OSPF Neighbors</b>    |                                                                                               |                                                                                                                                                                                                                                                                                                                |
| Address                  | Address of the neighbor.                                                                      |                                                                                                                                                                                                                                                                                                                |
| Interface                | Interface through which the neighbor is reachable.                                            |                                                                                                                                                                                                                                                                                                                |
| State                    | State of the neighbor: <b>Attempt, Down, Exchange, ExStart, Full, Init, Loading, or 2way.</b> | Generally, only the <b>Down</b> state, indicating a failed OSPF adjacency, and the <b>Full</b> state, indicating a functional adjacency, are maintained for more than a few seconds. The other states are transitional states that a neighbor is in only briefly while an OSPF adjacency is being established. |
| ID                       | ID of the neighbor.                                                                           |                                                                                                                                                                                                                                                                                                                |
| Priority                 | Priority of the neighbor to become the designated router.                                     |                                                                                                                                                                                                                                                                                                                |
| Activity Time            | The activity time.                                                                            |                                                                                                                                                                                                                                                                                                                |
| Area                     | Area that the neighbor is in.                                                                 |                                                                                                                                                                                                                                                                                                                |
| Options                  | Option bits received in the hello packets from the neighbor.                                  |                                                                                                                                                                                                                                                                                                                |
| DR Address               | Address of the designated router.                                                             |                                                                                                                                                                                                                                                                                                                |
| BDR Address              | Address of the backup designated router.                                                      |                                                                                                                                                                                                                                                                                                                |
| Uptime                   | Length of time since the neighbor came up.                                                    |                                                                                                                                                                                                                                                                                                                |
| Adjacency                | Length of time since the adjacency with the neighbor was established.                         |                                                                                                                                                                                                                                                                                                                |

- Related Documentation**
- [Configuring an OSPF Network \(J-Web Procedure\) on page 3207](#)
  - [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)

## Operational Commands

---

- `clear (ospf | ospf3) database`
- `clear (ospf | ospf3) io-statistics`
- `clear (ospf | ospf3) neighbor`
- `clear (ospf | ospf3) statistics`
- `clear (ospf | ospf3) overload`
- `show (ospf | ospf3) interface`
- `show (ospf | ospf3) io-statistics`
- `show (ospf | ospf3) log`
- `show (ospf | ospf3) neighbor`
- `show (ospf | ospf3) overview`
- `show (ospf | ospf3) route`
- `show (ospf | ospf3) statistics`
- `show ospf database`
- `show ospf3 database`



## clear (ospf | ospf3) database

**List of Syntax**    [Syntax on page 3281](#)  
                           [Syntax \(EX Series Switch and QFX Series\) on page 3281](#)

**Syntax**    clear (ospf | ospf3) database  
                   <advertising-router (*router-id* | self) >  
                   <area *area-id* >  
                   <asbrsummary >  
                   <external >  
                   <instance *instance-name* >  
                   <inter-area-prefix >  
                   <inter-area-router >  
                   <intra-area-prefix >  
                   <link-local >  
                   <logical-system (all | *logical-system-name*) >  
                   <lsa-id *lsa-id* >  
                   <netsummary >  
                   <network >  
                   <nssa >  
                   <opaque-area >  
                   <purge >  
                   <realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) >  
                   <router >

**Syntax (EX Series Switch and QFX Series)**    clear (ospf | ospf3) database  
                                                           <advertising-router (*router-id* | self) >  
                                                           <area *area-id* >  
                                                           <asbrsummary >  
                                                           <external >  
                                                           <instance *instance-name* >  
                                                           <inter-area-prefix >  
                                                           <inter-area-router >  
                                                           <intra-area-prefix >  
                                                           <link-local >  
                                                           <lsa-id *lsa-id* >  
                                                           <netsummary >  
                                                           <network >  
                                                           <nssa >  
                                                           <opaque-area >  
                                                           <purge >  
                                                           <router >

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                   **advertising-router** *router-id*, **netsummary**, **network**, **nssa**, **opaque-area**, and **router** options added in Junos OS Release 8.3. You must use the **purge** command with these options.  
                                   **area** *area-id* option added in Junos OS Release 8.3.  
                                   Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                   **realm** option added in Junos OS Release 9.2.  
                                   **advertising-router** (*router-id* | **self**) option added in Junos OS Release 9.5.  
                                   **advertising-router** (*router-id* | **self**) option introduced in Junos OS Release 9.5 for EX Series switches.  
                                   Command introduced in Junos OS Release 11.3 for the QFX Series.

**purge** option (and all options that are dependent on the **purge** option) hidden in Junos OS Release 13.3.

**Description** With the master Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and sync the new database with the master Routing Engine.



**CAUTION:** You can also use the **purge** command with any of the options to **discard** rather than **delete** the specified LSA entries. This command is useful only for testing. Use it with care, because it causes significant network disruption.

**Options** **none**—Delete all LSAs other than the system's own LSAs, which are regenerated. To resynchronize the database, the system destroys all adjacent neighbors that are in the state **EXSTART** or higher. The neighbors are then reacquired and the databases are synchronized.

**advertising-router (router-id | self)**—(Hidden) Discard entries for the LSA entries advertised by the specified routing device or by this routing device.

**area area-id**—(Optional) Discard entries for the LSAs in the specified area.

**asbrsummary**—(Optional) Discard summary AS boundary router LSA entries.

**external**—(Optional) Discard external LSAs.

**instance instance-name**—(Optional) Delete or discard entries for the specified routing instance only.

**inter-area-prefix**—(OSPFv3 only) (Optional) Discard interarea prefix LSAs.

**inter-area-router**—(OSPFv3 only) (Optional) Discard interarea router LSAs.

**intra-area-prefix**—(OSPFv3 only) (Optional) Discard intra-area prefix LSAs.

**logical-system (all | logical-system-name)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**link-local**—(Optional) Delete link-local LSAs.

**lsa-id lsa-id**—(Optional) Discard the LSA entries with the specified LSA identifier.

**netsummary**—(Hidden) Discard summary network LSAs.

**network**—(Hidden) Discard network LSAs.

**nssa**—(Hidden) Discard not-so-stubby area (NSSA) LSAs.

**opaque-area**—(Hidden) Discard opaque area-scope LSAs.

**purge**—(Hidden) Discard all entries in the link-state advertisement database. All link-state advertisements are set to **MAXAGE** and are flooded. The database is repopulated when the originators of the link-state advertisements receive the **MAXAGE** link-state advertisements and reissue them.

**realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)**—(OSPFv3 only) (Optional) Delete the entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

**router**—(Hidden) Discard router LSAs.

**Required Privilege Level**

clear

**Related Documentation**

- [show ospf database on page 3322](#)
- [show ospf3 database on page 3330](#)

**List of Sample Output** [clear ospf database on page 3283](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear ospf database

```
user@host> clear ospf database
```

## clear (ospf | ospf3) io-statistics

---

|                                                 |                                                                                                                                                                                                                   |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3284</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3284</a>                                                                                                      |
| <b>Syntax</b>                                   | clear (ospf   ospf3) io-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                        |
| <b>Syntax (EX Series Switch and QFX Series)</b> | clear (ospf   ospf3) io-statistics                                                                                                                                                                                |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                          |
| <b>Description</b>                              | Clear Open Shortest Path First (OSPF) input and output statistics.                                                                                                                                                |
| <b>Options</b>                                  | <b>none</b> —Clear OSPF input and output statistics.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                 | clear                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                    | <a href="#">clear ospf io-statistics on page 3284</a>                                                                                                                                                             |
| <b>Output Fields</b>                            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                             |

### Sample Output

#### clear ospf io-statistics

```
user@host> clear ospf io-statistics
```

## clear (ospf | ospf3) neighbor

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3285</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3285</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                   | <pre>clear (ospf   ospf3) neighbor &lt;area <i>area-id</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;neighbor&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>clear (ospf   ospf3) neighbor &lt;area <i>area-id</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;neighbor&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                              | Tear down Open Shortest Path First (OSPF) neighbor connections.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                  | <p><b>none</b>—Tear down OSPF connections with all neighbors for all routing instances.</p> <p><b>area <i>area-id</i></b>—(Optional) Tear down neighbor connections for the specified area only.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Tear down neighbor connections for the specified routing instance only.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Tear down neighbor connections for the specified interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear the state of the specified neighbor only.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Clear the state of the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| <b>Required Privilege Level</b>                 | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"> <li>• <a href="#">show (ospf   ospf3) neighbor on page 3301</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>                    | <a href="#">clear ospf neighbor on page 3286</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

### Sample Output

`clear ospf neighbor`

```
user@host> clear ospf neighbor
```

## clear (ospf | ospf3) statistics

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3287</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3287</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax</b>                                   | clear (ospf   ospf3) statistics<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )><br><realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)>                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (EX Series Switch and QFX Series)</b> | clear (ospf   ospf3) statistics<br><instance <i>instance-name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>realm</b> option introduced in Junos OS Release 9.2.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                              | Clear Open Shortest Path First (OSPF) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                  | <b>none</b> —Clear OSPF statistics.<br><br><b>instance <i>instance-name</i></b> —(Optional) Clear statistics for the specified routing instance only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b> —(Optional) (OSPFv3 only) Clear statistics for the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default. |
| <b>Required Privilege Level</b>                 | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"> <li>• <a href="#">show (ospf   ospf3) statistics on page 3318</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>                    | <a href="#">clear ospf statistics on page 3287</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>                            | See <a href="#">show (ospf   ospf3) statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### clear ospf statistics

The following sample output displays OSPF statistics before and after the **clear ospf statistics** command is entered:

```
user@host> show ospf statistics
```

| Packet type | Total | Last 5 seconds |
|-------------|-------|----------------|
|-------------|-------|----------------|

|          | Sent | Received | Sent | Received |
|----------|------|----------|------|----------|
| Hello    | 3254 | 2268     | 3    | 1        |
| DbD      | 41   | 46       | 0    | 0        |
| LSReq    | 8    | 7        | 0    | 0        |
| LSUpdate | 212  | 154      | 0    | 0        |
| LSAck    | 65   | 98       | 0    | 0        |

|                          |   |                    |   |   |
|--------------------------|---|--------------------|---|---|
| DBDs retransmitted       | : | 3, last 5 seconds  | : | 0 |
| LSAs flooded             | : | 12, last 5 seconds | : | 0 |
| LSAs flooded high-prio   | : | 0, last 5 seconds  | : | 0 |
| LSAs retransmitted       | : | 0, last 5 seconds  | : | 0 |
| LSAs transmitted to nbr: | : | 3, last 5 seconds  | : | 0 |
| LSAs requested           | : | 5, last 5 seconds  | : | 0 |
| LSAs acknowledged        | : | 19, last 5 seconds | : | 0 |

|                      |   |   |
|----------------------|---|---|
| Flood queue depth    | : | 0 |
| Total rexmit entries | : | 0 |
| db summaries         | : | 0 |
| lsreq entries        | : | 0 |

Receive errors:

626 subnet mismatches

user@host> clear ospf statistics

user@host> show ospf statistics

| Packet type | Total |          | Last 5 seconds |          |
|-------------|-------|----------|----------------|----------|
|             | Sent  | Received | Sent           | Received |
| Hello       | 3     | 1        | 3              | 1        |
| DbD         | 0     | 0        | 0              | 0        |
| LSReq       | 0     | 0        | 0              | 0        |
| LSUpdate    | 0     | 0        | 0              | 0        |
| LSAck       | 0     | 0        | 0              | 0        |

|                          |   |                   |   |   |
|--------------------------|---|-------------------|---|---|
| DBDs retransmitted       | : | 0, last 5 seconds | : | 0 |
| LSAs flooded             | : | 0, last 5 seconds | : | 0 |
| LSAs flooded high-prio   | : | 0, last 5 seconds | : | 0 |
| LSAs retransmitted       | : | 0, last 5 seconds | : | 0 |
| LSAs transmitted to nbr: | : | 0, last 5 seconds | : | 0 |
| LSAs requested           | : | 0, last 5 seconds | : | 0 |
| LSAs acknowledged        | : | 0, last 5 seconds | : | 0 |

|                      |   |   |
|----------------------|---|---|
| Flood queue depth    | : | 0 |
| Total rexmit entries | : | 0 |
| db summaries         | : | 0 |
| lsreq entries        | : | 0 |

Receive errors:

None



## clear (ospf | ospf3) overload

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3289</a><br><a href="#">Syntax (EX Series Switches) on page 3289</a>                                                                                                                                                                                                                                                                                           |
| <b>Syntax</b>                      | clear (ospf   ospf3) overload<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switches)</b> | clear (ospf   ospf3) overload<br><instance <i>instance-name</i> >                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                  |
| <b>Description</b>                 | Clear the Open Shortest Path First (OSPF) overload bit and rebuild link-state advertisements (LSAs).                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                     | <p><b>none</b>—Clear the overload bit and rebuild LSAs for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear the overload bit and rebuild LSAs for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | clear                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>       | <a href="#">clear ospf overload on page 3289</a>                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>               | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                     |

### Sample Output

#### clear ospf overload

```
user@host> clear ospf overload
```

```
show (ospf | ospf3) interface
```

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                             | <a href="#">Syntax on page 3290</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3290</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Syntax                                     | <pre>show (ospf   ospf3) interface &lt;brief   detail   extensive&gt; &lt;area <i>area-id</i>&gt; &lt;<i>interface-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Syntax (EX Series Switches and QFX Series) | <pre>show (ospf   ospf3) interface &lt;brief   detail   extensive&gt; &lt;area <i>area-id</i>&gt; &lt;<i>interface-name</i>&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Release Information                        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>area</b> option introduced in Junos OS Release 9.2.</p> <p><b>area</b> option introduced in Junos OS Release 9.2 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Description                                | Display the status of OSPF interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Options                                    | <p><b>none</b>—Display standard information about the status of all OSPF interfaces for all routing instances</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>area <i>area-id</i></b>—(Optional) Display information about the interfaces that belong to the specified area.</p> <p><b><i>interface-name</i></b>—(Optional) Display information for the specified interface.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(OSPFv3 only) (Optional) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| Required Privilege Level                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**List of Sample Output** [show ospf interface brief on page 3293](#)  
[show ospf interface detail on page 3293](#)  
[show ospf3 interface detail on page 3293](#)  
[show ospf interface detail\(When Multiarea Adjacency Is Configured\) on page 3293](#)  
[show ospf interface area area-id on page 3295](#)  
[show ospf interface extensive \(When Flooding Reduction Is Enabled\) on page 3295](#)  
[show ospf interface extensive \(When LDP Synchronization Is Configured\) on page 3295](#)

**Output Fields** [Table 344 on page 3291](#) lists the output fields for the **show (ospf | ospf3) interface** command. Output fields are listed in the approximate order in which they appear.

**Table 344: show (ospf | ospf3) interface Output Fields**

| Field Name              | Field Description                                                                                                                                                                                    | Level of Output         |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Interface</b>        | Name of the interface running OSPF version 2 or OSPF version 3.                                                                                                                                      | All levels              |
| <b>State</b>            | State of the interface: <b>BDR</b> , <b>Down</b> , <b>DR</b> , <b>DRother</b> , <b>Loop</b> , <b>PtToPt</b> , or <b>Waiting</b> .                                                                    | All levels              |
| <b>Area</b>             | Number of the area that the interface is in.                                                                                                                                                         | All levels              |
| <b>DR ID</b>            | Address of the area's designated router.                                                                                                                                                             | All levels              |
| <b>BDR ID</b>           | Backup designated router for a particular subnet.                                                                                                                                                    | All levels              |
| <b>Nbrs</b>             | Number of neighbors on this interface.                                                                                                                                                               | All levels              |
| <b>Type</b>             | Type of interface: <b>LAN</b> , <b>NBMA</b> , <b>P2MP</b> , <b>P2P</b> , or <b>Virtual</b> .                                                                                                         | <b>detail extensive</b> |
| <b>Address</b>          | IP address of the neighbor.                                                                                                                                                                          | <b>detail extensive</b> |
| <b>Mask</b>             | Netmask of the neighbor.                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Prefix-length</b>    | (OSPFv3) IPv6 prefix length, in bits.                                                                                                                                                                | <b>detail extensive</b> |
| <b>OSPF3-Intf-Index</b> | (OSPFv3) OSPF version 3 interface index.                                                                                                                                                             | <b>detail extensive</b> |
| <b>MTU</b>              | Interface maximum transmission unit (MTU).                                                                                                                                                           | <b>detail extensive</b> |
| <b>Cost</b>             | Interface cost (metric).                                                                                                                                                                             | <b>detail extensive</b> |
| <b>DR addr</b>          | Address of the designated router.                                                                                                                                                                    | <b>detail extensive</b> |
| <b>BDR addr</b>         | Address of the backup designated router.                                                                                                                                                             | <b>detail extensive</b> |
| <b>Adj count</b>        | Number of adjacent neighbors.                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Secondary</b>        | Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface for only one area. | <b>detail extensive</b> |

Table 344: show (ospf | ospf3) interface Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output         |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Flood Reduction</b> | Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the <b>DoNotAge</b> bit set. As a result, LSAs are refreshed only when a change occurs.                                                                                                                                                                                                                                   | <b>extensive</b>        |
| <b>Priority</b>        | Router priority used in designated router (DR) election on this interface.                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b> |
| <b>Flood list</b>      | List of link-state advertisements (LSAs) that might be about to flood this interface.                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>        |
| <b>Ack list</b>        | Acknowledgment list. List of pending acknowledgments on this interface.                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>        |
| <b>Descriptor list</b> | List of packet descriptors.                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b>        |
| <b>Hello</b>           | Configured value for the hello timer.                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b> |
| <b>Dead</b>            | Configured value for the dead timer.                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> |
| <b>Auth type</b>       | (OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: <ul style="list-style-type: none"> <li>• <b>MD5</b>—The MD5 mechanism is configured in accordance with RFC 2328.</li> <li>• <b>None</b>—No authentication method is configured.</li> <li>• <b>Password</b>—A simple password (RFC 2328) is configured.</li> </ul>                                                                                                                | <b>detail extensive</b> |
| <b>Topology</b>        | (Multiarea adjacency) Name of topology: <b>default</b> or <b>name</b> .                                                                                                                                                                                                                                                                                                                                                                                             |                         |
| <b>LDP sync state</b>  | (OSPFv2 and LDP synchronization) Current state of LDP synchronization: <b>in sync</b> , <b>in holddown</b> , and <b>not supported</b> .                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>        |
| <b>reason</b>          | (OSPFv2 and LDP synchronization) Reason for the current state of LDP synchronization. The LDP session might be up or down, or adjacency might be up or down.                                                                                                                                                                                                                                                                                                        | <b>extensive</b>        |
| <b>config holdtime</b> | (OSPFv2 and LDP synchronization) Configured value of the hold timer.<br><br>If the state is not synchronized, and the hold time is not infinity, the <b>remaining</b> field displays the number of seconds that remain until the configured hold timer expires.                                                                                                                                                                                                     | <b>extensive</b>        |
| <b>IPSec SA name</b>   | (OSPFv2) Name of the IPSec security association name.                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b> |
| <b>Active key ID</b>   | (OSPFv2 and MD5) Number from <b>0</b> to <b>255</b> that uniquely identifies an MD5 key.                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Start time</b>      | (OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as <b>Start time 1970 Jan 01 00:00:00 PST</b> . | <b>detail extensive</b> |

Table 344: show (ospf | ospf3) interface Output Fields (*continued*)

| Field Name                   | Field Description                          | Level of Output  |
|------------------------------|--------------------------------------------|------------------|
| ReXmit                       | Configured value for the Retransmit timer. | detail extensive |
| Stub, Not Stub, or Stub NSSA | Type of area.                              | detail extensive |

## Sample Output

### show ospf interface brief

```

user@host> show ospf interface brief
Intf           State   Area      DR ID      BDR ID      Nbrs
at-5/1/0.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
ge-2/3/0.0     DR      0.0.0.0   192.168.4.16 192.168.4.15 1
lo0.0          DR      0.0.0.0   192.168.4.16 0.0.0.0     0
so-0/0/0.0     Down    0.0.0.0   0.0.0.0    0.0.0.0     0
so-6/0/1.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
so-6/0/2.0     Down    0.0.0.0   0.0.0.0    0.0.0.0     0
so-6/0/3.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1

```

### show ospf interface detail

```

user@host> show ospf interface detail
Interface      State   Area      DR ID      BDR ID      Nbrs
fe-0/0/1.0     BDR     0.0.0.0   192.168.37.12 10.255.245.215 1
Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40
DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
t1-0/2/1.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     0
Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604
Adj count 0
Hello 10, Dead 40, ReXmit 5, Not Stub
Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST
IPsec SA Name: sa

```

### show ospf3 interface detail

```

user@host> show ospf3 interface so-0/0/3.0 detail
Interface      State   Area      DR-ID      BDR-ID      Nbrs
so-0/0/3.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     1
Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64
OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1
Hello 10, Dead 40, ReXmit 5, Not Stub

```

### show ospf interface detail (When Multiarea Adjacency Is Configured)

```

user@host> show ospf interface detail
regress@router> show ospf interface detail
Interface      State   Area      DR ID      BDR ID      Nbrs
lo0.0          DR      0.0.0.0   10.255.245.2 0.0.0.0     0

Type: LAN, Address: 127.0.0.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
DR addr: 127.0.0.1, Adj count: 0, Priority: 128
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None

```

```

Topology default (ID 0) -> Cost: 0
lo0.0          DR          0.0.0.0          10.255.245.2    0.0.0.0          0

Type: LAN, Address: 10.255.245.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0
DR addr: 10.255.245.2, Adj count: 0, Priority: 128
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 0
so-0/0/0.0     PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0     PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          0

Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0     PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0     PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          0

Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0     PtToPt  1.1.1.1          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0     PtToPt  1.1.1.1          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0     PtToPt  2.2.2.2          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0     PtToPt  2.2.2.2          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary

```

```

Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

```

#### show ospf interface area area-id

```

user@host> show ospf interface area 1.1.1.1
Interface      State   Area      DR ID      BDR ID      Nbrs
so-0/0/0.0     PtToPt  1.1.1.1   0.0.0.0    0.0.0.0     1
so-1/0/0.0     PtToPt  1.1.1.1   0.0.0.0    0.0.0.0     1

```

#### show ospf interface extensive (When Flooding Reduction Is Enabled)

```

user@host> show ospf interface extensive
Interface      State   Area      DR ID      BDR ID      Nbrs
fe-0/0/0.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     0

Type: P2P, Address: 10.10.10.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
Adj count: 0
Secondary, Flood Reduction
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

```

#### show ospf interface extensive (When LDP Synchronization Is Configured)

```

user@host> show ospf interface extensive
Interface      State   Area      DR ID      BDR ID
Nbrs
so-1/0/3.0     Down    0.0.0.0   0.0.0.0    0.0.0.0
0
Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 65535
Adj count: 0
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
LDP sync state: in holddown, for: 00:00:08, reason: LDP down during config
config holddtime: 10 seconds, remaining: 1

```

## show (ospf | ospf3) io-statistics

|                                                 |                                                                                                                                                                                                      |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3296</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3296</a>                                                                                         |
| <b>Syntax</b>                                   | show (ospf   ospf3) io-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                            |
| <b>Syntax (EX Series Switch and QFX Series)</b> | show (ospf   ospf3) io-statistics                                                                                                                                                                    |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.             |
| <b>Description</b>                              | Display Open Shortest Path First (OSPF) input and output statistics.                                                                                                                                 |
| <b>Options</b>                                  | none—Display OSPF input and output statistics.<br><br>logical-system (all   <i>logical-system-name</i> )—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                 |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"> <li><a href="#">clear (ospf   ospf3) statistics on page 3287</a></li> </ul>                                                                                       |
| <b>List of Sample Output</b>                    | <a href="#">show ospf io-statistics on page 3297</a>                                                                                                                                                 |
| <b>Output Fields</b>                            | Table 345 on page 3296 lists the output fields for the <b>show ospf io-statistics</b> command. Output fields are listed in the approximate order in which they appear.                               |

**Table 345: show (ospf | ospf3) io-statistics Output Fields**

| Field Name      | Field Description                                                                                         |
|-----------------|-----------------------------------------------------------------------------------------------------------|
| Packets read    | Number of OSPF packets read since the last time the routing protocol was started.                         |
| average per run | Total number of packets divided by the total number of times the OSPF read operation is scheduled to run. |
| max run         | Maximum number of packets for a given run among all scheduled runs.                                       |
| Receive errors  | Number of faulty packets received with errors.                                                            |



## Sample Output

### show ospf io-statistics

```
user@host> show ospf io-statistics
```

```
Packets read: 7361, average per run: 1.00, max run: 1  
Receive errors:  
None
```

## show (ospf | ospf3) log

**List of Syntax** [Syntax on page 3298](#)  
[Syntax \(EX Series Switch and QFX Series\) on page 3298](#)

**Syntax** `show (ospf | ospf3) log`  
`<instance instance-name>`  
`<logical-system (all | logical-system-name)>`  
`<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>`  
`<topology topology-name>`

**Syntax (EX Series Switch and QFX Series)** `show (ospf | ospf3) log`  
`<instance instance-name>`  
`<topology topology-name>`

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
**topology** option introduced in Junos OS Release 9.0.  
**topology** option introduced in Junos OS Release 9.0 for EX Series switches.  
**realm** option introduced in Junos OS Release 9.2.  
 Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations.

**Options** **none**—Display entries in the OSPF log of SPF calculations for all routing instances.

**instance *instance-name***—(Optional) Display entries for the specified routing instance.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**topology *topology-name***—(Optional) (OSPFv2 only) Display entries for the specified topology.

**realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)**—(OSPFv3 only) (Optional) Display entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

**Required Privilege Level** view

**List of Sample Output** [show ospf log on page 3299](#)  
[show ospf log topology voice on page 3299](#)

**Output Fields** [Table 346 on page 3298](#) lists the output fields for the **show (ospf | ospf3) log** command. Output fields are listed in the approximate order in which they appear.

**Table 346: show (ospf | ospf3) log Output Fields**

| Field Name  | Field Description                                                                      |
|-------------|----------------------------------------------------------------------------------------|
| <b>When</b> | Time, in weeks ( <b>w</b> ) and days ( <b>d</b> ), since the SPF calculation was made. |

Table 346: show (ospf | ospf3) log Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                            |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type       | Type of calculation: Cleanup, External, Interarea, NSSA, Redist, SPF, Stub, Total, or Virtuallink.                                                                           |
| Elapsed    | Amount of time, in seconds, that elapsed during the operation, or the time required to complete the SPF calculation. The start time is the time displayed in the When field. |

## Sample Output

### show ospf log

```

user@host> show ospf log
When          Type          Elapsed
1w4d 17:25:58 Stub          0.000017
1w4d 17:25:58 SPF           0.000070
1w4d 17:25:58 Stub          0.000019
1w4d 17:25:58 Interarea     0.000054
1w4d 17:25:58 External      0.000005
1w4d 17:25:58 Cleanup       0.000203
1w4d 17:25:58 Total         0.000537
1w4d 17:24:48 SPF           0.000125
1w4d 17:24:48 Stub          0.000017
1w4d 17:24:48 SPF           0.000100
1w4d 17:24:48 Stub          0.000016
1w4d 17:24:48 Interarea     0.000056
1w4d 17:24:48 External      0.000005
1w4d 17:24:48 Cleanup       0.000238
1w4d 17:24:48 Total         0.000600
...

```

### show ospf log topology voice

```

user@host> show ospf log topology voice
Topology voice SPF log:

    Last instance of each event type
When          Type          Elapsed
00:06:11      SPF           0.000116
00:06:11      Stub           0.000114
00:06:11      Interarea      0.000126
00:06:11      External        0.000067
00:06:11      NSSA            0.000037
00:06:11      Cleanup         0.000186

    Maximum length of each event type
When          Type          Elapsed
00:13:43      SPF           0.000140
00:13:33      Stub           0.000116
00:13:43      Interarea      0.000128
00:13:33      External        0.000075
00:13:38      NSSA            0.000039
00:13:53      Cleanup         0.000657

    Last 100 events

```

| When     | Type      | Elapsed  |
|----------|-----------|----------|
| 00:13:53 | SPF       | 0.000090 |
| 00:13:53 | Stub      | 0.000041 |
| 00:13:53 | Interarea | 0.000123 |
| 00:13:53 | External  | 0.000040 |
| 00:13:53 | NSSA      | 0.000038 |
| 00:13:53 | Cleanup   | 0.000657 |
| 00:13:53 | Total     | 0.001252 |
| .        |           |          |
| .        |           |          |
| 00:06:11 | SPF       | 0.000116 |
| 00:06:11 | Stub      | 0.000114 |
| 00:06:11 | Interarea | 0.000126 |
| 00:06:11 | External  | 0.000067 |
| 00:06:11 | NSSA      | 0.000037 |
| 00:06:11 | Cleanup   | 0.000186 |
| 00:06:11 | Total     | 0.000818 |

## show (ospf | ospf3) neighbor

**List of Syntax**    [Syntax on page 3301](#)  
                          [Syntax \(EX Series Switches and QFX Series\) on page 3301](#)

**Syntax**    `show (ospf | ospf3) neighbor`  
                  `<brief | detail | extensive>`  
                  `<area area-id>`  
                  `<instance (all | instance-name)>`  
                  `<interface interface-name>`  
                  `<logical-system (all | logical-system-name)>`  
                  `<neighbor>`  
                  `<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>`

**Syntax (EX Series Switches and QFX Series)**    `show (ospf | ospf3) neighbor`  
                  `<brief | detail | extensive>`  
                  `<area area-id>`  
                  `<instance (all | instance-name)>`  
                  `<interface interface-name>`  
                  `<neighbor>`

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                  Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                  **instance all** option introduced in Junos OS Release 9.1.  
                                  **instance all** option introduced in Junos OS Release 9.1 for EX Series switches.  
                                  **area**, **interface**, and **realm** options introduced in Junos OS Release 9.2.  
                                  **area** and **interface** options introduced in Junos OS Release 9.2 for EX Series switches.  
                                  Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description**    Display information about OSPF neighbors.

CPU utilization might increase while the device learns its OSPF neighbors. We recommend that you use the **show (ospf | ospf3) neighbor** command after the device learns and establishes OSPF neighbor adjacencies. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the **show (ospf | ospf3) neighbor** command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the CLI.

**Options**    **none**—Display standard information about all OSPF neighbors for all routing instances.

**brief | detail | extensive**—(Optional) Display the specified level of output.

**area *area-id***—(Optional) Display information about the OSPF neighbors for the specified area.

**instance (all | *instance-name*)**—(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.

**interface *interface-name***—(Optional) Display information about OSPF neighbors for the specified logical interface.

**logical-system** (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

**neighbor**—(Optional) Display information about the specified OSPF neighbor.

**realm** (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Display information about the OSPF neighbors for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

**Required Privilege Level** view

**Related Documentation** • [clear \(ospf | ospf3\) neighbor on page 3285](#)

**List of Sample Output** [show ospf neighbor brief on page 3304](#)  
[show ospf neighbor detail on page 3304](#)  
[show ospf neighbor extensive on page 3305](#)  
[show ospf3 neighbor detail on page 3306](#)  
[show ospf neighbor area area-id on page 3306](#)  
[show ospf neighbor interface interface-name on page 3306](#)  
[show ospf3 neighbor instance all \(OSPFv3 Multiple Family Address Support Enabled\) on page 3306](#)

**Output Fields** [Table 347 on page 3302](#) lists the output fields for the **show (ospf | ospf3) neighbor** command. Output fields are listed in the approximate order in which they appear.

**Table 347: show (ospf | ospf3) neighbor Output Fields**

| Field Name       | Field Description                                  | Level of Output |
|------------------|----------------------------------------------------|-----------------|
| <b>Address</b>   | Address of the neighbor.                           | All levels      |
| <b>Interface</b> | Interface through which the neighbor is reachable. | All levels      |

Table 347: show (ospf | ospf3) neighbor Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output         |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>State</b>                          | <p>State of the neighbor:</p> <ul style="list-style-type: none"> <li>• <b>Attempt</b>—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor.</li> <li>• <b>Down</b>—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the <b>Down</b> state, although at a reduced frequency.</li> <li>• <b>Exchange</b>—Routing device is describing its entire link-state database by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged.</li> <li>• <b>ExStart</b>—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the master, and to determine the initial sequence number.</li> <li>• <b>Full</b>—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements.</li> <li>• <b>Init</b>—A hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state might occur, for example, because the routing device itself did not appear in the neighbor's hello packet.</li> <li>• <b>Loading</b>—Link-state request packets are sent to the neighbor to acquire more recent advertisements that have been discovered (but not yet received) in the <b>Exchange</b> state.</li> <li>• <b>2Way</b>—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state <b>2Way</b> or greater.</li> </ul> | All levels              |
| <b>ID</b>                             | Router ID of the neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels              |
| <b>Pri</b>                            | Priority of the neighbor to become the designated router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels              |
| <b>Dead</b>                           | Number of seconds until the neighbor becomes unreachable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels              |
| <b>Link state acknowledgment list</b> | Number of link-state acknowledgments received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>        |
| <b>Link state retransmission list</b> | <p>Total number of link-state advertisements retransmitted. For <b>extensive</b> output only, the following information is also displayed:</p> <ul style="list-style-type: none"> <li>• <b>Type</b>—Type of link advertisement: <b>ASBR</b>, <b>Sum</b>, <b>Extern</b>, <b>Network</b>, <b>NSSA</b>, <b>OpagArea</b>, <b>Router</b>, or <b>Summary</b>.</li> <li>• <b>LSA ID</b>—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device.</li> <li>• <b>Adv rtr</b>—Address of the routing device that sent the advertisement.</li> <li>• <b>Seq</b>—Link sequence number of the advertisement.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |

Table 347: show (ospf | ospf3) neighbor Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                | Level of Output  |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Neighbor-address</b> | (OSPFv3 only) If the neighbor uses virtual links, the <b>Neighbor-address</b> is the site-local, local, or global address. If the neighbor uses a physical interface, the <b>Neighbor-address</b> is an IPv6 link-local address. | detail extensive |
| <b>area</b>             | Area that the neighbor is in.                                                                                                                                                                                                    | detail extensive |
| <b>OSPF3-Intf-Index</b> | (OSPFv3 only) Displays the OSPFv3 interface index.                                                                                                                                                                               | detail extensive |
| <b>opt</b>              | Option bits received in the hello packets from the neighbor.                                                                                                                                                                     | detail extensive |
| <b>DR or DR-ID</b>      | Address of the designated router.                                                                                                                                                                                                | detail extensive |
| <b>BDR or BDR-ID</b>    | Address of the backup designated router.                                                                                                                                                                                         | detail extensive |
| <b>Up</b>               | Length of time since the neighbor came up.                                                                                                                                                                                       | detail extensive |
| <b>adjacent</b>         | Length of time since the adjacency with the neighbor was established.                                                                                                                                                            | detail extensive |

## Sample Output

### show ospf neighbor brief

```

user@host> show ospf neighbor brief
  Address      Intf      State      ID          Pri  Dead
192.168.254.225 fxp3.0    2Way       10.250.240.32 128  36
192.168.254.230 fxp3.0    Full       10.250.240.8  128  38
192.168.254.229 fxp3.0    Full       10.250.240.35 128  33
10.1.1.129      fxp2.0    Full       10.250.240.12 128  37
10.1.1.131      fxp2.0    Full       10.250.240.11 128  38
10.1.2.1        fxp1.0    Full       10.250.240.9  128  32
10.1.2.81       fxp0.0    Full       10.250.240.10 128  33

```

### show ospf neighbor detail

```

user@host> show ospf neighbor detail
  Address      Interface      State      ID          Pri  Dead
10.5.1.2      ge-1/2/0.1    Full       10.5.1.2    128  37
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:28, adjacent 05:17:36
Link state acknowledgment list: 3 entries

Link state retransmission list: 9 entries

10.5.10.2      ge-1/2/0.10    ExStart    10.5.1.38   128  34
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:28
master, seq 0xac1530f8, rexmit DBD in 3 sec
rexmit LSREQ in 0 sec
10.5.11.2      ge-1/2/0.11    Full       10.5.1.42   128  38
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:28, adjacent 05:26:46
Link state retransmission list: 1 entries

```



```

10.5.12.2      ge-1/2/0.12      ExStart  10.5.1.46      128   33
area 0.0.0.1, opt 0x42, DR 10.5.12.2, BDR 10.5.12.1
Up 06:09:28
master, seq 0xac188a68, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec

```

### show ospf neighbor extensive

```

user@host> show ospf neighbor extensive
Address      Interface      State      ID      Pri  Dead
10.5.1.2      ge-1/2/0.1     Full       10.5.1.2  128   33
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:42, adjacent 05:17:50
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
Summary 10.8.56.0    172.25.27.82 0x8000004d
Router  10.5.1.94    10.5.1.94    0x8000005c
Network 10.5.24.2    10.5.1.94    0x80000036
Summary 10.8.57.0    172.25.27.82 0x80000024
Extern  1.10.90.0    10.8.1.2     0x80000041
Extern  1.4.109.0     10.6.1.2     0x80000041
Router  10.5.1.190    10.5.1.190   0x8000005f
Network 10.5.48.2    10.5.1.190   0x8000003d
Summary 10.8.58.0    172.25.27.82 0x8000004d
Extern  1.10.91.0    10.8.1.2     0x80000041
Extern  1.4.110.0     10.6.1.2     0x80000041
Router  10.5.1.18     10.5.1.18    0x8000005f
Network 10.5.5.2     10.5.1.18    0x80000033
Summary 10.8.59.0    172.25.27.82 0x8000003a
Summary 10.8.62.0    172.25.27.82 0x80000025

10.5.10.2     ge-1/2/0.10     ExStart  10.5.1.38      128   38
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:42
master, seq 0xac1530f8, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec
10.5.11.2     ge-1/2/0.11     Full     10.5.1.42      128   33
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:42, adjacent 05:27:00
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
Summary 10.8.58.0    172.25.27.82 0x8000004d

```

|         |           |              |            |
|---------|-----------|--------------|------------|
| Extern  | 1.10.91.0 | 10.8.1.2     | 0x80000041 |
| Extern  | 1.1.247.0 | 10.5.1.2     | 0x8000003f |
| Extern  | 1.4.110.0 | 10.6.1.2     | 0x80000041 |
| Router  | 10.5.1.18 | 10.5.1.18    | 0x8000005f |
| Network | 10.5.5.2  | 10.5.1.18    | 0x80000033 |
| Summary | 10.8.59.0 | 172.25.27.82 | 0x8000003a |

**show ospf3 neighbor detail**

```

user@host> show ospf3 neighbor detail
ID          Interface          State    Pri    Dead
10.255.71.13 fe-0/0/2.0          Full     128    30
Neighbor-address fe80::290:69ff:fe9b:e002
area 0.0.0.0, opt 0x13, OSPF3-Intf-Index 2
DR-ID 10.255.71.13, BDR-ID 10.255.71.12
Up 02:51:43, adjacent 02:51:43

```

**show ospf neighbor area area-id**

```

user@host >show ospf neighbor area 1.1.1.1
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0        Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.55 so-1/0/0.0        Full     10.255.245.5 128    37
Area 1.1.1.1

```

**show ospf neighbor interface interface-name**

```

user@host >show ospf neighbor interface so-0/0/0.0
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0        Full     10.255.245.4 128    37
Area 0.0.0.0
192.168.37.47 so-0/0/0.0        Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.47 so-0/0/0.0        Full     10.255.245.4 128    32
Area 2.2.2.2

```

**show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled)**

```

user @host > show ospf3 neighbor instance all
Instance: ina
Realm: ipv6-unicast
ID          Interface          State    Pri    Dead
100.1.1.1    fe-0/0/2.0          Full     128    37
Neighbor-address fe80::217:cb00:c87c:8c03
Instance: inb
Realm: ipv4-unicast
ID          Interface          State    Pri    Dead
100.1.2.1    fe-0/0/2.1          Full     128    33
Neighbor-address fe80::217:cb00:c97c:8c03

```

## show (ospf | ospf3) overview

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3307</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3307</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                   | <pre>show (ospf   ospf3) overview &lt;brief   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show (ospf   ospf3) overview &lt;brief   extensive&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                      | <p>Command introduced in Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Database protection introduced in Junos 10.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                              | Display Open Shortest Path First (OSPF) overview information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                  | <p><b>none</b>—Display standard information about all OSPF neighbors for all routing instances.</p> <p><b>brief   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                    | <a href="#">show ospf overview on page 3309</a><br><a href="#">show ospf overview (With Database Protection) on page 3310</a><br><a href="#">show ospf3 overview (With Database Protection) on page 3310</a><br><a href="#">show ospf overview extensive on page 3310</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>                            | <p><a href="#">Table 348 on page 3308</a> lists the output fields for the <b>show ospf overview</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

Table 348: show ospf overview Output Fields

| Field name                       | Field Description                                                                                                                                                                        | Level of Output |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Instance</b>                  | OSPF routing instance.                                                                                                                                                                   | All levels      |
| <b>Router ID</b>                 | Router ID of the routing device.                                                                                                                                                         | All levels      |
| <b>Route table index</b>         | Route table index.                                                                                                                                                                       | All levels      |
| <b>Configured overload</b>       | Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires. | All levels      |
| <b>Topology</b>                  | Topology identifier.                                                                                                                                                                     | All levels      |
| <b>Prefix export count</b>       | Number of prefixes exported into OSPF.                                                                                                                                                   | All levels      |
| <b>Full SPF runs</b>             | Number of complete Shortest Path First calculations.                                                                                                                                     | All levels      |
| <b>SPF delay</b>                 | Delay before performing consecutive Shortest Path First calculations.                                                                                                                    | All levels      |
| <b>SPF holddown</b>              | Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.                                           | All levels      |
| <b>SPF rapid runs</b>            | Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.                                                                | All levels      |
| <b>LSA refresh time</b>          | Refresh period for link-state advertisement (in minutes).                                                                                                                                | All levels      |
| <b>Database protection state</b> | Current state of database protection.                                                                                                                                                    | All levels      |
| <b>Warning threshold</b>         | Threshold at which a warning message is logged (percentage of maximum LSA count).                                                                                                        | All levels      |
| <b>Non self-generated LSAs</b>   | Number of LSAs whose router ID is not equal to the local router ID: <b>Current</b> , <b>Warning</b> (threshold), and <b>Allowed</b> .                                                    | All levels      |
| <b>Ignore time</b>               | How long the database has been in the ignore state.                                                                                                                                      | All levels      |
| <b>Reset time</b>                | How long the database must stay out of the ignore or isolated state before it returns to normal operations.                                                                              | All levels      |
| <b>Ignore count</b>              | Number of times the database has been in the ignore state: <b>Current</b> and <b>Allowed</b> .                                                                                           | All levels      |
| <b>Restart</b>                   | Graceful restart capability: <b>enabled</b> or <b>disabled</b> .                                                                                                                         | All levels      |
| <b>Restart duration</b>          | Time period for complete reacquisition of OSPF neighbors.                                                                                                                                | All levels      |
| <b>Restart grace period</b>      | Time period for which the neighbors should consider the restarting routing device as part of the topology.                                                                               | All levels      |

Table 348: show ospf overview Output Fields (*continued*)

| Field name                    | Field Description                                                                                                                                                                                                                                                                                                                     | Level of Output  |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Graceful restart helper mode  | (OSPFv2) Standard graceful restart helper capability (based on RFC 3623): <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                         | All levels       |
| Restart-signaling helper mode | (OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                  | All levels       |
| Helper mode                   | (OSPFv3) Graceful restart helper capability: <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                                                      | All levels       |
| Trace options                 | OSPF-specific trace options.                                                                                                                                                                                                                                                                                                          | <b>extensive</b> |
| Trace file                    | Name of the file to receive the output of the tracing operation.                                                                                                                                                                                                                                                                      | <b>extensive</b> |
| Area                          | Area number. Area 0.0.0.0 is the backbone area.                                                                                                                                                                                                                                                                                       | All levels       |
| Stub type                     | Stub type of area: <b>Normal Stub</b> , <b>Not Stub</b> , or <b>Not so Stubby Stub</b> .                                                                                                                                                                                                                                              | All levels       |
| Authentication Type           | Type of authentication: <b>None</b> , <b>Password</b> , or <b>MD5</b> .<br><br><b>NOTE:</b> The <b>Authentication Type</b> field refers to the authentication configured at the <b>[edit protocols ospf area area-id]</b> level. Any authentication configured for an interface in this area will not affect the value of this field. | All levels       |
| Area border routers           | Number of area border routers.                                                                                                                                                                                                                                                                                                        | All levels       |
| Neighbors                     | Number of autonomous system boundary routers.                                                                                                                                                                                                                                                                                         | All levels       |

## Sample Output

### show ospf overview

```

user@host> show ospf overview
Instance: master
  Router ID: 10.255.245.6
  Route table index: 0
  Configured overload, expires in 118 seconds
  LSA refresh time: 50 minutes
  Restart: Enabled
    Restart duration: 20 sec
    Restart grace period: 40 sec
    Helper mode: enabled
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 0
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 1
  SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

```

### show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
  Router ID: 10.255.112.218
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Restart: Enabled
    Restart duration: 180 sec
    Restart grace period: 210 sec
    Graceful restart helper mode: Enabled
    Restart-signaling helper mode: Enabled
  Database protection state: Normal
    Warning threshold: 70 percent
    Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 1
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 160
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 70
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

### show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
  Router ID: 10.255.112.128
  Route table index: 0
  LSA refresh time: 50 minutes
  Database protection state: Normal
    Warning threshold: 80 percent
    Non self-generated LSAs: Current 3, Warning 8, Allowed 10
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 2
  Area: 0.0.0.0
    Stub type: Not Stub
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 7
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

### show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
  Router ID: 1.1.1.103
  Route table index: 0
  Full SPF runs: 13, SPF delay: 0.200000 sec
  LSA refresh time: 50 minutes
```

```
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
```

```
show (ospf | ospf3) route
```

- |                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3312</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3312</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax</b>                                   | <pre>show (ospf   ospf3) route &lt;brief   detail   extensive&gt; &lt;abr   asbr   extern   inter   intra&gt; &lt;destination&gt; &lt;instance (default   ipv4-multicast   <i>instance-name</i>)&gt; &lt;logical-system (default   ipv4-multicast   <i>logical-system-name</i>)&gt; &lt;network&gt; &lt;no-backup-coverage&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt; &lt;router&gt; &lt;topology (default   ipv4-multicast   <i>topology-name</i>)&gt; &lt;transit&gt;</pre>                                                                                                                               |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show (ospf   ospf3) route &lt;brief   detail   extensive&gt; &lt;abr   asbr   extern   inter   intra&gt; &lt;destination&gt; &lt;instance <i>instance-name</i> &lt;network&gt; &lt;no-backup-coverage&gt; &lt;router&gt; &lt;topology (default   ipv4-multicast   <i>topology-name</i>)&gt; &lt;transit&gt;</pre>                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>topology</b> option introduced in Junos OS Release 9.0.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                            |
| <b>Description</b>                              | Display the entries in the Open Shortest Path First (OSPF) routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                  | <p><b>none</b>—Display standard information about all entries in the OSPF routing table for all routing instances and all topologies.</p> <p><b><i>destination</i></b>—Display routes to the specified IP address (with optional destination prefix length).</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>abr</b>—(Optional) Display routes to area border routers.</p> <p><b>asbr</b>—(Optional) Display routes to autonomous system border routers.</p> <p><b>extern</b>—(Optional) Display external routes.</p> <p><b>inter</b>—(Optional) Display interarea routes.</p> |



**intra**—(Optional) Display intra-area routes.

**instance** (**default** | **ipv4-multicast** | *instance-name*)—(Optional) Display entries for the default routing instance, the IPv4 multicast routing instance, or for the specified routing instance.

**logical-system** (**default** | **ipv4-multicast** | *logical-system-name*)—(Optional) Perform this operation on the default logical system, the IPv4 multicast logical system, or on a particular logical system.

**network**—(Optional) Display routes to networks.

**no-backup-coverage**—(Optional) Display routes with no backup coverage.

**realm** (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Display entries in the routing table for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

**router**—(Optional) Display routes to all routers.

**topology** (**default** | **ipv4-multicast** | *topology-name*)—(OSPFv2 only) (Optional) Display routes for the default OSPF topology, IPv4 multicast topology, or for a particular topology.

**transit**—(Optional) (OSPFv3 only) Display OSPFv3 routes to pseudonodes.

**Required Privilege Level**

view

**List of Sample Output**

[show ospf route on page 3315](#)  
[show ospf route detail on page 3315](#)  
[show ospf3 route on page 3315](#)  
[show ospf3 route detail on page 3316](#)  
[show ospf route topology voice on page 3316](#)

**Output Fields**

[Table 349 on page 3313](#) list the output fields for the **show (ospf | ospf3) route** command. Output fields are listed in the approximate order in which they appear.

**Table 349: show (ospf | ospf3) route Output Fields**

| Field Name      | Field Description         | Output Level |
|-----------------|---------------------------|--------------|
| <b>Topology</b> | Name of the topology.     | All levels   |
| <b>Prefix</b>   | Destination of the route. | All levels   |

Table 349: show (ospf | ospf3) route Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Output Level |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| <b>Path type</b>          | How the route was learned: <ul style="list-style-type: none"> <li>• <b>Inter</b>—Interarea route</li> <li>• <b>Ext1</b>—External type 1 route</li> <li>• <b>Ext2</b>—External type 2 route</li> <li>• <b>Intra</b>—Intra-area route</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                              | All levels   |
| <b>Route type</b>         | The type of routing device from which the route was learned: <ul style="list-style-type: none"> <li>• <b>AS BR</b>—Route to AS border router.</li> <li>• <b>Area BR</b>—Route to area border router.</li> <li>• <b>Area/AS BR</b>—Route to router that is both an <b>Area BR</b> and <b>AS BR</b>.</li> <li>• <b>Network</b>—Network router.</li> <li>• <b>Router</b>—Route to a router that is neither an <b>Area BR</b> nor an <b>AS BR</b>.</li> <li>• <b>Transit</b>—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link.</li> <li>• <b>Discard</b>—Route to a summary discard.</li> </ul> | All levels   |
| <b>NH Type</b>            | Next-hop type: <b>LSP</b> or <b>IP</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels   |
| <b>Metric</b>             | Route's metric value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels   |
| <b>NH-interface</b>       | (OSPFv3 only) Interface through which the route's next hop is reachable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels   |
| <b>NH-addr</b>            | (OSPFv3 only) IPv6 address of the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels   |
| <b>NextHop Interface</b>  | (OSPFv2 only) Interface through which the route's next hop is reachable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels   |
| <b>Nexthop addr/label</b> | (OSPFv2 only) If the <b>NH Type</b> is <b>IP</b> , then it is the address of the next hop. If the <b>NH Type</b> is <b>LSP</b> , then it is the name of the label-switched path.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | All levels   |
| <b>Area</b>               | Area ID of the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | detail       |
| <b>Origin</b>             | Router from which the route was learned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | detail       |
| <b>Type 7</b>             | Route was learned through a not-so-stubby area (NSSA) link-state advertisement (LSA).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | detail       |
| <b>P-bit</b>              | Route was learned through NSSA LSA and the propagate bit was set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | detail       |
| <b>Fwd NZ</b>             | Forwarding address is nonzero. <b>Fwd NZ</b> is only displayed if the route is learned through an NSSA LSA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | detail       |

Table 349: show (ospf | ospf3) route Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Output Level  |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| <b>optional-capability</b> | Optional capabilities propagated in the router LSA. This field is in the output for intra-area router routes only (when <b>Route Type</b> is <b>Area BR</b> , <b>AS BR</b> , <b>Area/AS BR</b> , or <b>Router</b> ), not for interarea router routes or network routes. Three bits in this field are defined as follows: <ul style="list-style-type: none"> <li>• <b>0x4 (V)</b>—Routing device is at the end of a virtual active link.</li> <li>• <b>0x2 (E)</b>—Routing device is an autonomous system boundary router.</li> <li>• <b>0x1 (B)</b>—Routing device is an area border router.</li> </ul> | <b>detail</b> |
| <b>priority</b>            | The priority assigned to the prefix: <ul style="list-style-type: none"> <li>• <b>high</b></li> <li>• <b>medium</b></li> <li>• <b>low</b></li> </ul> <p><b>NOTE:</b> The <b>priority</b> field applies only to routes of type <b>Network</b>.</p>                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b> |

## Sample Output

### show ospf route

```

user@host> show ospf route
Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type   Type    Type                Interface    addr/label
10.255.71.12     Intra Router  IP    1       fe-0/0/2.0   192.16.22.86
10.255.71.13/32  Intra Network IP    0       lo0.0
192.168.222.84/30 Intra Network LSP   1       fe-0/0/2.0   lsp-ab

```

### show ospf route detail

```

user@host> show ospf route detail
Topology default Route Table:

Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type   Type    Type                Interface    addr/label
10.255.14.174    Inter AS BR  IP      210    t1-3/0/1.0
area 0.0.0.2, origin 10.255.14.185
10.255.14.178    Intra Router  IP      200    t3-3/1/3.0
area 0.0.0.2, origin 10.255.14.178, optional-capability 0x0
10.210.1.0/30    Intra Network IP      10     t3-3/1/2.0
area 0.0.0.2, origin 10.255.14.172, priority medium
100.1.1.1/32     Inter Network IP      210    t1-3/0/1.0
area 0.0.0.2, origin 10.255.14.185, priority low
112.3.1.0/24     Ext2  Network  IP      0     t1-3/0/1.0
area 0.0.0.0, origin 10.255.14.174, priority high
200.3.3.0/30     Inter Network IP      220    t1-3/0/1.0
area 0.0.0.2, origin 10.255.14.185, priority high

```

### show ospf3 route

```

user@host> show ospf3 route
Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type   Type    Type                Interface    addr/label

```

```

10.255.71.13      Intra Router      IP      1
  NH-interface fe-0/0/2.0, NH-addr fe80::290:69ff:fe9b:e002
10.255.71.13;0.0.0.2
10.255.245.1      Intra Router      IP      40 fxp1.1      192.168.36.17

  area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3      Intra AS BR      IP      1 fxp2.3      192.168.36.34

  area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32   Intra Network    IP      40 fxp1.1      192.168.36.17

  area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32   Intra Network    IP      0 lo0.0
  area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32   Intra Network    IP      1 fxp2.3      192.168.36.34

  area 0.0.0.0, origin 10.255.245.3, priority low
  Intra Transit    IP      1
  NH-interface fe-0/0/2.0
192::168:222:84/126 Intra Network    IP      1
  NH-interface fe-0/0/2.0
abcd::71:12/128   Intra Network    IP      0
  NH-interface lo0.0
abcd::71:13/128   Intra Network    LSP     1
  NH-interface fe-0/0/2.0, NH-addr lsp-cd

```

#### show ospf3 route detail

```

user@host> show ospf3 route detail
Prefix
Path      Route      NH      Metric
type      type
10.255.14.174 Intra Area/AS BR IP 110
  NH-interface so-1/2/2.0
  Area 0.0.0.0, Origin 10.255.14.174, Optional-capability 0x3
10.255.14.178 Intra Router IP 200
  NH-interface t3-3/1/3.0
  Area 0.0.0.0, Origin 10.255.14.178, Optional-capability 0x0
10.255.14.185;0.0.0.2 Intra Transit IP 200
  NH-interface t1-3/0/1.0
  NH-interface so-1/2/2.0
  Area 0.0.0.0, Origin 10.255.14.185
1000:1:1::1/128 Inter Network IP 110
  NH-interface so-1/2/2.0
  Area 0.0.0.0, Origin 10.255.14.174, Priority low
1001:2:1::/48 Ext1 Network IP 110
  NH-interface so-1/2/2.0
  Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority medium
1002:1:7::/48 Ext2 Network IP 0
  NH-interface so-1/2/2.0
  Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority low
1002:3:4::/48 Ext2 Network IP 0
  NH-interface so-1/2/2.0
  Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority high
abcd::10:255:14:172/128 Intra Network IP 0
  NH-interface lo0.0
  Area 0.0.0.0, Origin 10.255.14.172, Priority low

```

#### show ospf route topology voice

```
user@host show ospf route topology voice
```

Topology voice Route Table:

| Prefix          | Path<br>Type | Route<br>Type | NH<br>Type | Metric | NextHop<br>Interface | Nexthop<br>addr/label |
|-----------------|--------------|---------------|------------|--------|----------------------|-----------------------|
| 10.255.8.2      | Intra        | Router        | IP         | 1      | so-0/2/0.0           |                       |
| 10.255.8.3      | Intra        | Router        | IP         | 2      | so-0/2/0.0           |                       |
| 10.255.8.1/32   | Intra        | Network       | IP         | 0      | lo0.0                |                       |
| 10.255.8.2/32   | Intra        | Network       | IP         | 1      | so-0/2/0.0           |                       |
| 10.255.8.3/32   | Intra        | Network       | IP         | 2      | so-0/2/0.0           |                       |
| 192.168.8.0/29  | Intra        | Network       | IP         | 2      | so-0/2/0.0           |                       |
| 192.168.8.44/30 | Intra        | Network       | IP         | 2      | so-0/2/0.0           |                       |
| 192.168.8.46/32 | Intra        | Network       | IP         | 1      | so-0/2/0.0           |                       |
| 192.168.8.48/30 | Intra        | Network       | IP         | 1      | so-0/2/1.0           |                       |
| 192.168.8.52/30 | Intra        | Network       | IP         | 2      | so-0/2/0.0           |                       |
| 192.168.9.44/30 | Intra        | Network       | IP         | 1      | so-0/2/0.0           |                       |
| 192.168.9.45/32 | Intra        | Network       | IP         | 2      | so-0/2/0.0           |                       |

## show (ospf | ospf3) statistics

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3318</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3318</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                   | <pre>show (ospf   ospf3) statistics &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show (ospf   ospf3) statistics &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                              | Display OSPF statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                  | <p><b>none</b>—Display OSPF statistics for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display all statistics for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Display all statistics for the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p> |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"> <li><a href="#">clear (ospf   ospf3) statistics on page 3287</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                    | <a href="#">show ospf statistics on page 3320</a><br><a href="#">show ospf statistics logical-system all on page 3320</a><br><a href="#">show ospf3 statistics on page 3321</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>                            | <p><a href="#">Table 350 on page 3318</a> lists the output fields for the <b>show (ospf   ospf3) statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Table 350: show (ospf | ospf3) statistics Output Fields**

| Field Name                | Field Description                          |
|---------------------------|--------------------------------------------|
| Packet type               | Type of OSPF packet.                       |
| Total Sent/Total Received | Total number of packets sent and received. |

Table 350: show (ospf | ospf3) statistics Output Fields (*continued*)

| Field Name                                  | Field Description                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Last 5 seconds Sent/Last 5 seconds Received | Total number of packets sent and received in the last 5 seconds.                                                                                                                                                                                                                                                                                                                |
| DBDs retransmitted                          | Total number of database description packets retransmitted, and number retransmitted in the last 5 seconds.                                                                                                                                                                                                                                                                     |
| LSAs flooded                                | Total number of link-state advertisements flooded, and number flooded in the last 5 seconds.                                                                                                                                                                                                                                                                                    |
| LSAs flooded high-prio                      | <p>Total number of high priority link-state advertisements flooded, and number flooded in the last 5 seconds.</p> <p>A link-state advertisement is deemed a high priority if it has changed since it was last sent.</p>                                                                                                                                                         |
| LSAs retransmitted                          | Total number of link-state advertisements retransmitted, and number retransmitted in the last 5 seconds.                                                                                                                                                                                                                                                                        |
| LSAs transmitted to nbr                     | Total number of link-state advertisements transmitted to a neighbor, and number transmitted in the last 5 seconds.                                                                                                                                                                                                                                                              |
| LSAs requested                              | Total number of link-state advertisements requested by neighboring devices, and number requested in the last 5 seconds.                                                                                                                                                                                                                                                         |
| LSAs acknowledged                           | Total number of link-state advertisements acknowledged, and number acknowledged in the last 5 seconds.                                                                                                                                                                                                                                                                          |
| Flood queue depth                           | Total number of entries in the extended queue.                                                                                                                                                                                                                                                                                                                                  |
| Total rexmit entries                        | Total number of retransmission entries waiting to be sent from the OSPF routing instance.                                                                                                                                                                                                                                                                                       |
| db summaries                                | Total number of database description summaries waiting to be sent from the OSPF routing instance.                                                                                                                                                                                                                                                                               |
| lsreq entries                               | Total number of link-state request entries waiting to be sent from the OSPF routing instance.                                                                                                                                                                                                                                                                                   |
| Receive errors                              | <p>Number and type of receive errors. Some sample receive errors include:</p> <ul style="list-style-type: none"> <li>• mtu mismatches</li> <li>• no interface found</li> <li>• no virtual link found</li> <li>• nssa mismatches</li> <li>• stub area mismatches</li> <li>• subnet mismatches</li> </ul> <p>If there are no receive errors, the output displays <b>none</b>.</p> |

## Sample Output

### show ospf statistics

```

user@host> show ospf statistics
Packet type          Total
                   Sent      Received
Hello                31        14
  DbD                 9        10
  LSReq               2         2
LSUpdate             8        16
  LSAck              9         9

                   Last 5 seconds
                   Sent      Received
Hello                2         2
  DbD                 0         0
  LSReq               0         0
LSUpdate             0         0
  LSAck              0         0

DBDs retransmitted   :          3, last 5 seconds :          0
LSAs flooded         :         12, last 5 seconds :          0
LSAs flooded high-prio :          0, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          3, last 5 seconds :          0
LSAs requested       :          5, last 5 seconds :          0
LSAs acknowledged    :         19, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:
  862 no interface found
  115923 no virtual link found

```

### show ospf statistics logical-system all

```

user@host> show ospf statistics logical-system all
logical-system: C
OSPF instance is not running
-----

logical-system: B
Packet type          Total
                   Sent      Received
Hello              313740      313653
  DbD                3         2
  LSReq              1         1
LSUpdate            2752      1825
  LSAck             1821      2747

                   Last 5 seconds
                   Sent      Received
Hello                1         0
  DbD                 0         0
  LSReq               0         0
LSUpdate            0         0
  LSAck              0         0

DBDs retransmitted   :          0, last 5 seconds :          0
LSAs flooded         :        2741, last 5 seconds :          0
LSAs flooded high-prio :         10, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          2, last 5 seconds :          0
LSAs requested       :          1, last 5 seconds :          0
LSAs acknowledged    :       1831, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:

```



```

None
-----

logical-system: A

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello                313698      313695         0         0
  DbD                  2         3         0         0
  LSReq                1         1         0         0
LSUpdate             1825      2752         0         0
LSAck                2747      1821         0         0

DBDs retransmitted   :          0, last 5 seconds :          0
LSAs flooded         :        1825, last 5 seconds :          0
LSAs flooded high-prio :         10, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:         1, last 5 seconds :          0
LSAs requested        :          2, last 5 seconds :          0
LSAs acknowledged    :       2748, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries         :          0

Receive errors:
None
-----

```

### show ospf3 statistics

```

user@host> show ospf3 statistics

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello                0         0         0         0
  DbD                  0         0         0         0
  LSReq                0         0         0         0
LSUpdate             0         0         0         0
LSAck                0         0         0         0

DBDs retransmitted   :          0, last 5 seconds :          0
LSAs flooded         :          0, last 5 seconds :          0
LSAs flooded high-prio :          0, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          0, last 5 seconds :          0
LSAs requested        :          0, last 5 seconds :          0
LSAs acknowledged    :          0, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries         :          0

Receive errors:
None

```

## show ospf database

---

**List of Syntax**   [Syntax on page 3322](#)

[Syntax \(EX Series Switches and QFX Series\) on page 3322](#)

**Syntax**   `show ospf database`  
          `<brief | detail | extensive | summary>`  
          `<advertising-router (address | self)>`  
          `<area area-id>`  
          `<asbrsummary>`  
          `<external>`  
          `<instance instance-name>`  
          `<link-local>`  
          `<logical-system (all | logical-system-name)>`  
          `<lsa-id lsa-id>`  
          `<netsummary>`  
          `<network>`  
          `<nssa>`  
          `<opaque-area>`  
          `<router>`

**Syntax (EX Series Switches and QFX Series)**   `show ospf database`  
                                  `<brief | detail | extensive | summary>`  
                                  `<advertising-router (address | self)>`  
                                  `<area area-id>`  
                                  `<asbrsummary>`  
                                  `<external>`  
                                  `<instance instance-name>`  
                                  `<link-local>`  
                                  `<lsa-id lsa-id>`  
                                  `<netsummary>`  
                                  `<network>`  
                                  `<nssa>`  
                                  `<opaque-area>`  
                                  `<router>`

**Release Information**   Command introduced before Junos OS Release 7.4.  
                              Command introduced in Junos OS Release 9.0 for EX Series switches.  
                              **advertising-router self (address | self)** option introduced in Junos OS Release 9.5.  
                              **advertising-router self (address | self)** option introduced in Junos OS Release 9.5 for EX Series switches.  
                              Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description**   Display the entries in the OSPF version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets.

**Options**   **none**—Display standard information about entries in the OSPFv2 link-state database for all routing instances.

**brief | detail | extensive | summary**—(Optional) Display the specified level of output.

**advertising-router (address | self)**—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.

**area** *area-id*—(Optional) Display the LSAs in a particular area.

**asbrsummary**—(Optional) Display summary AS boundary router LSA entries.

**external**—(Optional) Display external LSAs.

**instance** *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

**link-local**—(Optional) Display information about link-local LSAs.

**logical-system** (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

**lsa-id** *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

**netsummary**—(Optional) Display summary network LSAs.

**network**—(Optional) Display information about network LSAs.

**nssa**—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

**opaque-area**—(Optional) Display opaque area-scope LSAs.

**router**—(Optional) Display information about router LSAs.

**Required Privilege Level**

view

**Related Documentation**

- [clear \(ospf | ospf3\) database on page 3281](#)

**List of Sample Output**

[show ospf database on page 3325](#)  
[show ospf database brief on page 3325](#)  
[show ospf database detail on page 3325](#)  
[show ospf database extensive on page 3327](#)  
[show ospf database summary on page 3329](#)

**Output Fields**

[Table 351 on page 3323](#) describes the output fields for the **show ospf database** command. Output fields are listed in the approximate order in which they appear.

**Table 351: show ospf database Output Fields**

| Field Name     | Field Description                                                                                                                                        | Level of Output |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>area</b>    | Area number. Area 0.0.0.0 is the backbone area.                                                                                                          | All levels      |
| <b>Type</b>    | Type of link advertisement: <b>ASBRSum</b> , <b>Extern</b> , <b>Network</b> , <b>NSSA</b> , <b>OpaqArea</b> , <b>Router</b> , or <b>Summary</b> .        | All levels      |
| <b>ID</b>      | LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device. | All levels      |
| <b>Adv Rtr</b> | Address of the routing device that sent the advertisement.                                                                                               | All levels      |

Table 351: show ospf database Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Seq</b>                              | Link sequence number of the advertisement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels              |
| <b>Age</b>                              | Time elapsed since the LSA was originated, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels              |
| <b>Opt</b>                              | Optional OSPF capabilities associated with the LSA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | All levels              |
| <b>Cksum</b>                            | Checksum value of the LSA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels              |
| <b>Len</b>                              | Length of the advertisement, in bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels              |
| <b>Router</b>                           | Router link-state advertisement information: <ul style="list-style-type: none"> <li><b>bits</b>—Flags describing the routing device that generated the LSP.</li> <li><b>link count</b>—Number of links in the advertisement.</li> <li><b>id</b>—ID of a routing device or subnet on the link.</li> <li><b>data</b>—For stub networks, the subnet mask. Otherwise, the IP address of the routing device that generated the LSP.</li> <li><b>type</b>—Type of link. It can be <b>PointToPoint</b>, <b>Transit</b>, <b>Stub</b>, or <b>Virtual</b>.</li> <li><b>TOS count</b>—Number of type-of-service (ToS) entries in the advertisement.</li> <li><b>TOS 0 metric</b>—Metric for ToS 0.</li> <li><b>TOS</b>—Type-of-service (ToS) value.</li> <li><b>metric</b>—Metric for the ToS.</li> </ul> | <b>detail extensive</b> |
| <b>Network</b>                          | Network link-state advertisement information: <ul style="list-style-type: none"> <li><b>mask</b>—Network mask.</li> <li><b>attached router</b>—ID of the attached neighbor.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Summary</b>                          | Summary link-state advertisement information: <ul style="list-style-type: none"> <li><b>mask</b>—Network mask.</li> <li><b>TOS</b>—Type-of-service (ToS) value.</li> <li><b>metric</b>—Metric for the ToS.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b> |
| <b>Gen timer</b>                        | How long until the LSA is regenerated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b>        |
| <b>Aging timer</b>                      | How long until the LSA expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>        |
| <b>Installed <i>hh:mm:ss</i> ago</b>    | How long ago the route was installed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b>        |
| <b>expires in <i>hh:mm:ss</i></b>       | How long until the route expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>extensive</b>        |
| <b>sent <i>hh:mm:ss</i> ago</b>         | How long ago the LSA was sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>extensive</b>        |
| <b>Last changed <i>hh:mm:ss</i> ago</b> | How long ago the route was changed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>        |

Table 351: show ospf database Output Fields (*continued*)

| Field Name   | Field Description                                                                  | Level of Output |
|--------------|------------------------------------------------------------------------------------|-----------------|
| Change count | Number of times the route has changed.                                             | extensive       |
| Ours         | Indicates that this is a local advertisement.                                      | extensive       |
| Router LSAs  | Number of router link-state advertisements in the link-state database.             | summary         |
| Network LSAs | Number of network link-state advertisements in the link-state database.            | summary         |
| Summary LSAs | Number of summary link-state advertisements in the link-state database.            | summary         |
| NSSA LSAs    | Number of not-so-stubby area link-state advertisements in the link-state database. | summary         |

## Sample Output

### show ospf database

```

user@host> show ospf database
OSPF link state database, Area 0.0.0.1
  Type      ID            Adv Rtr      Seq          Age  Opt  Cksum  Len
Router     10.255.70.103   10.255.70.103 0x80000002   215  0x20 0x4112  48
Router     *10.255.71.242  10.255.71.242 0x80000002   214  0x20 0x11b1  48
Summary    *23.1.1.0       10.255.71.242 0x80000002   172  0x20 0x6d72  28
Summary    *24.1.1.0       10.255.71.242 0x80000002   177  0x20 0x607e  28
NSSA       *33.1.1.1       10.255.71.242 0x80000002   217  0x28 0x73bd  36

      OSPF link state database, Area 0.0.0.2
  Type      ID            Adv Rtr      Seq          Age  Opt  Cksum  Len
Router     10.255.71.52   10.255.71.52  0x80000004   174  0x20 0xd021  36
Router     *10.255.71.242  10.255.71.242 0x80000003   173  0x20 0xe191  36
Network    *23.1.1.1       10.255.71.242 0x80000002   173  0x20 0x9c76  32
Summary    *12.1.1.0       10.255.71.242 0x80000001   217  0x20 0xfeec  28
Summary    *24.1.1.0       10.255.71.242 0x80000002   177  0x20 0x607e  28
NSSA       *33.1.1.1       10.255.71.242 0x80000001   222  0x28 0xe047  36

      OSPF link state database, Area 0.0.0.3
  Type      ID            Adv Rtr      Seq          Age  Opt  Cksum  Len
Router     10.255.71.238   10.255.71.238 0x80000003   179  0x20 0x3942  36
Router     *10.255.71.242  10.255.71.242 0x80000003   177  0x20 0xf37d  36
Network    *24.1.1.1       10.255.71.242 0x80000002   177  0x20 0xc591  32
Summary    *12.1.1.0       10.255.71.242 0x80000001   217  0x20 0xfeec  28
Summary    *23.1.1.0       10.255.71.242 0x80000002   172  0x20 0x6d72  28
NSSA       *33.1.1.1       10.255.71.242 0x80000001   222  0x28 0xeb3b  36

```

### show ospf database brief

The output for the **show ospf database brief** command is identical to that for the **show ospf database** command. For sample output, see [show ospf database on page 3325](#).

### show ospf database detail

```
user@host> show ospf database detail
```

```

    OSPF link state database, Area 0.0.0.1
    Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router 10.255.70.103  10.255.70.103  0x80000002  261  0x20 0x4112  48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Router *10.255.71.242  10.255.71.242  0x80000002  260  0x20 0x11b1  48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Summary *23.1.1.0      10.255.71.242  0x80000002  218  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0      10.255.71.242  0x80000002  223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA *33.1.1.1        10.255.71.242  0x80000002  263  0x28 0x73bd  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0

```

```

    OSPF link state database, Area 0.0.0.2
    Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router 10.255.71.52   10.255.71.52   0x80000004  220  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router *10.255.71.242  10.255.71.242  0x80000003  219  0x20 0xe191  36
  bits 0x3, link count 1
  id 23.1.1.1, data 23.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *23.1.1.1      10.255.71.242  0x80000002  219  0x20 0x9c76  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.52
Summary *12.1.1.0      10.255.71.242  0x80000001  263  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0      10.255.71.242  0x80000002  223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA *33.1.1.1        10.255.71.242  0x80000001  268  0x28 0xe047  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0

```

```

    OSPF link state database, Area 0.0.0.3
    Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router 10.255.71.238  10.255.71.238  0x80000003  225  0x20 0x3942  36
  bits 0x0, link count 1
  id 24.1.1.1, data 24.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router *10.255.71.242  10.255.71.242  0x80000003  223  0x20 0xf37d  36
  bits 0x3, link count 1
  id 24.1.1.1, data 24.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *24.1.1.1      10.255.71.242  0x80000002  223  0x20 0xc591  32
  mask 255.255.255.0
  attached router 10.255.71.242

```

```

    attached router 10.255.71.238
Summary *12.1.1.0          10.255.71.242    0x80000001    263    0x20 0xfeec    28
    mask 255.255.255.0
    TOS 0x0, metric 1
Summary *23.1.1.0          10.255.71.242    0x80000002    218    0x20 0x6d72    28
    mask 255.255.255.0
    TOS 0x0, metric 1
NSSA  *33.1.1.1            10.255.71.242    0x80000001    268    0x28 0xeb3b    36
    mask 255.255.255.255
    Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0

```

### show ospf database extensive

```

user@host> show ospf database extensive
    OSPF link state database, Area 0.0.0.1
Type      ID          Adv Rtr      Seq      Age      Opt  Cksum  Len
Router    10.255.70.103    10.255.70.103  0x80000002  286    0x20 0x4112  48
    bits 0x0, link count 2
    id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
    TOS count 0, TOS 0 metric 1
    id 12.1.1.0, data 255.255.255.0, Type Stub (3)
    TOS count 0, TOS 0 metric 1
    Aging timer 00:55:14
    Installed 00:04:43 ago, expires in 00:55:14
    Last changed 00:04:43 ago, Change count: 2
Router  *10.255.71.242    10.255.71.242    0x80000002    285    0x20 0x11b1  48
    bits 0x3, link count 2
    id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
    TOS count 0, TOS 0 metric 1
    id 12.1.1.0, data 255.255.255.0, Type Stub (3)
    TOS count 0, TOS 0 metric 1
    Gen timer 00:45:15
    Aging timer 00:55:15
    Installed 00:04:45 ago, expires in 00:55:15, sent 00:04:43 ago
    Last changed 00:04:45 ago, Change count: 2, Ours
Summary *23.1.1.0          10.255.71.242    0x80000002    243    0x20 0x6d72    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:57
    Aging timer 00:55:57
    Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0          10.255.71.242    0x80000002    248    0x20 0x607e    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:52
    Aging timer 00:55:52
    Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
NSSA  *33.1.1.1            10.255.71.242    0x80000002    288    0x28 0x73bd    36
    mask 255.255.255.255
    Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0
    Gen timer 00:45:12
    Aging timer 00:55:12
    Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:48 ago
    Last changed 00:04:48 ago, Change count: 2, Ours

    OSPF link state database, Area 0.0.0.2
Type      ID          Adv Rtr      Seq      Age      Opt  Cksum  Len
Router    10.255.71.52    10.255.71.52    0x80000004    245    0x20 0xd021  36
    bits 0x0, link count 1

```

```

id 23.1.1.1, data 23.1.1.2, Type Transit (2)
TOS count 0, TOS 0 metric 1
Aging timer 00:55:55
Installed 00:04:02 ago, expires in 00:55:55
Last changed 00:04:02 ago, Change count: 2
Router *10.255.71.242 10.255.71.242 0x80000003 244 0x20 0xe191 36
bits 0x3, link count 1
id 23.1.1.1, data 23.1.1.1, Type Transit (2)
TOS count 0, TOS 0 metric 1
Gen timer 00:45:56
Aging timer 00:55:56
Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
Last changed 00:04:04 ago, Change count: 2, Ours
Network *23.1.1.1 10.255.71.242 0x80000002 244 0x20 0x9c76 32
mask 255.255.255.0
attached router 10.255.71.242
attached router 10.255.71.52
Gen timer 00:45:56
Aging timer 00:55:56
Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
Last changed 00:04:04 ago, Change count: 1, Ours
Summary *12.1.1.0 10.255.71.242 0x80000001 288 0x20 0xfeec 28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:04 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0 10.255.71.242 0x80000002 248 0x20 0x607e 28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:04 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1 10.255.71.242 0x80000001 293 0x28 0xe047 36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:04 ago
Last changed 00:04:53 ago, Change count: 1, Ours

OSPF link state database, Area 0.0.0.3
Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.71.238 10.255.71.238 0x80000003 250 0x20 0x3942 36
bits 0x0, link count 1
id 24.1.1.1, data 24.1.1.2, Type Transit (2)
TOS count 0, TOS 0 metric 1
Aging timer 00:55:50
Installed 00:04:07 ago, expires in 00:55:50
Last changed 00:04:07 ago, Change count: 2
Router *10.255.71.242 10.255.71.242 0x80000003 248 0x20 0xf37d 36
bits 0x3, link count 1
id 24.1.1.1, data 24.1.1.1, Type Transit (2)
TOS count 0, TOS 0 metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 2, Ours
Network *24.1.1.1 10.255.71.242 0x80000002 248 0x20 0xc591 32

```



```

mask 255.255.255.0
attached router 10.255.71.242
attached router 10.255.71.238
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 1, Ours
Summary *12.1.1.0      10.255.71.242    0x80000001    288    0x20 0xfeec    28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:13 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243    0x20 0x6d72    28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA  *33.1.1.1        10.255.71.242    0x80000001    293    0x28 0xeb3b    36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:13 ago
Last changed 00:04:53 ago, Change count: 1, Ours

```

#### show ospf database summary

```

user@host> show ospf database summary
Area 0.0.0.1:
  2 Router LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.2:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.3:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Externals:
Interface fe-2/2/1.0:
Interface ge-0/3/2.0:
Interface so-0/1/2.0:
Interface so-0/1/2.0:

```

```
show ospf3 database
```

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                             | <a href="#">Syntax on page 3330</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3330</a>                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Syntax                                     | <pre>show ospf3 database &lt;brief   detail   extensive   summary&gt; &lt;advertising-router (<i>address</i>   self)&gt; &lt;area <i>area-id</i>&gt; &lt;external&gt; &lt;instance <i>instance-name</i>&gt; &lt;inter-area-prefix&gt; &lt;inter-area-router&gt; &lt;intra-area-prefix&gt; &lt;link&gt; &lt;link-local&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;lsa-id <i>lsa-id</i>&gt; &lt;network&gt; &lt;nssa&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt; &lt;router&gt;</pre> |
| Syntax (EX Series Switches and QFX Series) | <pre>show ospf3 database &lt;brief   detail   extensive   summary&gt; &lt;advertising-router (<i>address</i>   self)&gt; &lt;area <i>area-id</i>&gt; &lt;external&gt; &lt;instance <i>instance-name</i>&gt; &lt;inter-area-prefix&gt; &lt;inter-area-router&gt; &lt;intra-area-prefix&gt; &lt;link&gt; &lt;link-local&gt; &lt;lsa-id <i>lsa-id</i>&gt; &lt;network&gt; &lt;nssa&gt; &lt;router&gt;</pre>                                                                                                                          |
| Release Information                        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p><b>advertising-router</b> (<i>address</i>   <b>self</b>) option introduced in Junos Release 9.5.</p> <p><b>advertising-router</b> (<i>address</i>   <b>self</b>) option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                  |
| Description                                | <p>Display the entries in the OSPF version 3 (OSPFv3) link-state database, which contains data about link-state advertisement (LSA) packets.</p>                                                                                                                                                                                                                                                                                                                                                                                  |
| Options                                    | <p><b>none</b>—Display standard information about all entries in the OSPFv3 link-state database.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p>                                                                                                                                                                                                                                                                                                                         |

**advertising-router** (*address* | *self*)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.

**area** *area-id*—(Optional) Display the LSAs in a particular area.

**external**—(Optional) Display external LSAs.

**instance** *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

**inter-area-prefix**—(Optional) Display information about interarea-prefix LSAs.

**inter-area-router**—(Optional) Display information about interarea-router LSAs.

**intra-area-prefix**—(Optional) Display information about intra-area-prefix LSAs.

**link**—(Optional) Display information about link LSAs.

**link-local**—(Optional) Display information about link-local LSAs.

**logical-system** (*all* | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

**lsa-id** *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

**network**—(Optional) Display information about network LSAs.

**nssa**—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

**realm** (*ipv4-multicast* | *ipv4-unicast* | *ipv6-multicast*)—(Optional) Display information about the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family other than IPv6 unicast, which is the default.

**router**—(Optional) Display information about router LSAs.

**Required Privilege Level**

view

**Related Documentation**

- [clear \(ospf | ospf3\) database on page 3281](#)

**List of Sample Output**

[show ospf3 database brief on page 3336](#)  
[show ospf3 database extensive on page 3336](#)  
[show ospf3 database summary on page 3339](#)

**Output Fields**

[Table 352 on page 3331](#) lists the output fields for the **show ospf3 database** command. Output fields are listed in the approximate order in which they appear.

**Table 352: show ospf3 database Output Fields**

| Field Name                                        | Field Description                                 | Level of Output        |
|---------------------------------------------------|---------------------------------------------------|------------------------|
| OSPF link state database, area <i>area-number</i> | Entries in the link-state database for this area. | brief detail extensive |

Table 352: show ospf3 database Output Fields (*continued*)

| Field Name                                                                       | Field Description                                                                                                                                                                                                                                                                                                                                                          | Level of Output               |
|----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| <b>OSPF AS SCOPE<br/>link state database</b>                                     | Entries in the AS scope link-state database.                                                                                                                                                                                                                                                                                                                               | <b>brief detail extensive</b> |
| <b>OSPF Link-Local<br/>link state database,<br/>interface<br/>interface-name</b> | Entries in the link-local link-state database for this interface.                                                                                                                                                                                                                                                                                                          | <b>brief detail extensive</b> |
| <b>area</b>                                                                      | Area number. Area <b>0.0.0.0</b> is the backbone area.                                                                                                                                                                                                                                                                                                                     | All levels                    |
| <b>Type</b>                                                                      | Type of link advertisement: <b>Extern</b> , <b>InterArPfx</b> , <b>InterArRtr</b> , <b>IntraArPrx</b> , <b>Link</b> , <b>Network</b> , <b>NSSA</b> , or <b>Router</b> .                                                                                                                                                                                                    | <b>brief detail extensive</b> |
| <b>ID</b>                                                                        | Link identifier included in the advertisement. An asterisk (*) preceding the identifier marks database entries that originated from the local routing device.                                                                                                                                                                                                              | <b>brief detail extensive</b> |
| <b>Adv Rtr</b>                                                                   | Address of the routing device that sent the advertisement.                                                                                                                                                                                                                                                                                                                 | <b>brief detail extensive</b> |
| <b>Seq</b>                                                                       | Link sequence number of the advertisement.                                                                                                                                                                                                                                                                                                                                 | <b>brief detail extensive</b> |
| <b>Age</b>                                                                       | Time elapsed since the LSA was originated, in seconds.                                                                                                                                                                                                                                                                                                                     | <b>brief detail extensive</b> |
| <b>Cksum</b>                                                                     | Checksum value of the LSA.                                                                                                                                                                                                                                                                                                                                                 | <b>brief detail extensive</b> |
| <b>Len</b>                                                                       | Length of the advertisement, in bytes.                                                                                                                                                                                                                                                                                                                                     | <b>brief detail extensive</b> |
| <b>Router (Router Link-State Advertisements)</b>                                 |                                                                                                                                                                                                                                                                                                                                                                            |                               |
| <b>bits</b>                                                                      | Flags describing the routing device that generated the LSP.                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>       |
| <b>Options</b>                                                                   | Option bits carried in the router LSA.                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>       |
| <b>For Each Router Link</b>                                                      |                                                                                                                                                                                                                                                                                                                                                                            |                               |
| <b>Type</b>                                                                      | Type of interface. The value of all other output fields describing a routing device interface depends on the interface's type: <ul style="list-style-type: none"> <li>• <b>PointToPoint (1)</b>—Point-to-point connection to another routing device.</li> <li>• <b>Transit (2)</b>—Connection to a transit network.</li> <li>• <b>Virtual (4)</b>—Virtual link.</li> </ul> | <b>detail extensive</b>       |
| <b>Loc-if-id</b>                                                                 | Local interface ID assigned to the interface that uniquely identifies the interface with the routing device.                                                                                                                                                                                                                                                               | <b>detail extensive</b>       |
| <b>Nbr-if-id</b>                                                                 | Interface ID of the neighbor's interface for this routing device link.                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>       |
| <b>Nbr-rtr-id</b>                                                                | Router ID of the neighbor routing device (for type 2 interfaces, the attached link's designated router).                                                                                                                                                                                                                                                                   | <b>detail extensive</b>       |
| <b>Metric</b>                                                                    | Cost of the router link.                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>       |

Table 352: show ospf3 database Output Fields (*continued*)

| Field Name                                                     | Field Description                                                                                                                                                                                                                                            | Level of Output  |
|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Gen timer                                                      | How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                          | extensive        |
| Aging timer                                                    | How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                                 | extensive        |
| Installed <i>nn:nn:nn</i> ago                                  | How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                           | extensive        |
| expires in <i>nn:nn:nn</i>                                     | How long until the route expires, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                               | extensive        |
| sent <i>nn:nn:nn</i> ago                                       | Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .                                                                                                       | extensive        |
| Ours                                                           | Indicates that this is a local advertisement.                                                                                                                                                                                                                | extensive        |
| <b>Network (Network Link-State Advertisements)</b>             |                                                                                                                                                                                                                                                              |                  |
| Options                                                        | Option bits carried in the network LSA.                                                                                                                                                                                                                      | detail extensive |
| Attached Router                                                | Router IDs of each of the routing devices attached to the link. Only routing devices that are fully adjacent to the designated router are listed. The designated router includes itself in this list.                                                        | detail extensive |
| <b>InterArPfx (Interarea-Prefix Link-State Advertisements)</b> |                                                                                                                                                                                                                                                              |                  |
| Prefix                                                         | IPv6 address prefix.                                                                                                                                                                                                                                         | detail extensive |
| Prefix-options                                                 | Option bit associated with the prefix.                                                                                                                                                                                                                       | detail extensive |
| Metric                                                         | Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range. | detail extensive |
| Gen timer                                                      | How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                          | extensive        |
| Aging timer                                                    | How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                                 | extensive        |
| Installed <i>nn:nn:nn</i> ago                                  | How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                           | extensive        |
| expires in <i>nn:nn:nn</i>                                     | How long until the route expires, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                               | extensive        |
| sent <i>nn:nn:nn</i> ago                                       | Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .                                                                                                       | extensive        |
| Ours                                                           | Indicates that this is a local advertisement.                                                                                                                                                                                                                | extensive        |
| <b>InterArRtr (Interarea-Router Link-State Advertisements)</b> |                                                                                                                                                                                                                                                              |                  |
| Dest-router-id                                                 | Router ID of the routing device described by the LSA.                                                                                                                                                                                                        | detail extensive |

Table 352: show ospf3 database Output Fields (*continued*)

| Field Name                                         | Field Description                                                                                                                                                                                                                                            | Level of Output         |
|----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>options</b>                                     | Optional capabilities supported by the routing device.                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Metric</b>                                      | Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range. | <b>detail extensive</b> |
| <b>Prefix</b>                                      | IPv6 address prefix.                                                                                                                                                                                                                                         | <b>extensive</b>        |
| <b>Prefix-options</b>                              | Option bit associated with the prefix.                                                                                                                                                                                                                       | <b>extensive</b>        |
| <b>Extern (External Link-State Advertisements)</b> |                                                                                                                                                                                                                                                              |                         |
| <b>Prefix</b>                                      | IPv6 address prefix.                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Prefix-options</b>                              | Option bit associated with the prefix.                                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Metric</b>                                      | Cost of the route, which depends on the value of <b>Type</b> .                                                                                                                                                                                               | <b>detail extensive</b> |
| <b>Type <i>n</i></b>                               | Type of external metric: <b>Type 1</b> or <b>Type 2</b> .                                                                                                                                                                                                    | <b>detail extensive</b> |
| <b>Aging timer</b>                                 | How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                                 | <b>extensive</b>        |
| <b>Installed <i>nn:nn:nn</i> ago</b>               | How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                           | <b>extensive</b>        |
| <b>expires in <i>nn:nn:nn</i></b>                  | How long until the route expires, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                               | <b>extensive</b>        |
| <b>sent <i>nn:nn:nn</i> ago</b>                    | Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .                                                                                                       | <b>extensive</b>        |
| <b>Link (Link-State Advertisements)</b>            |                                                                                                                                                                                                                                                              |                         |
| <b>IPv6-Address</b>                                | IPv6 link-local address on the link for which this link LSA originated.                                                                                                                                                                                      | <b>detail extensive</b> |
| <b>Options</b>                                     | Option bits carried in the link LSA.                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>priority</b>                                    | Router priority of the interface attaching the originating routing device to the link.                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Prefix-count</b>                                | Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.                                                                                                              | <b>detail extensive</b> |
| <b>Prefix</b>                                      | IPv6 address prefix.                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Prefix-options</b>                              | Option bit associated with the prefix.                                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Gen timer</b>                                   | How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                                          | <b>extensive</b>        |

Table 352: show ospf3 database Output Fields (*continued*)

| Field Name                                               | Field Description                                                                                                                                                                                                              | Level of Output  |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Aging timer                                              | How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                   | extensive        |
| Installed <i>nn:nn:nn</i> ago                            | How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .                                                                                                                                             | extensive        |
| expires in <i>nn:nn:nn</i>                               | How long until the route expires, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                 | extensive        |
| sent <i>nn:nn:nn</i> ago                                 | Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .                                                                         | extensive        |
| Ours                                                     | Indicates that this is a local advertisement.                                                                                                                                                                                  | extensive        |
| IntraArPfx (Intra-Area-Prefix Link-State Advertisements) |                                                                                                                                                                                                                                |                  |
| Ref-lsa-type                                             | LSA type of the referenced LSA. <ul style="list-style-type: none"> <li><b>Router</b>—Address prefixes are associated with a router LSA.</li> <li><b>Network</b>—Address prefixes are associated with a network LSA.</li> </ul> | detail extensive |
| Ref-lsa-id                                               | Link-state ID of the referenced LSA.                                                                                                                                                                                           | detail extensive |
| Ref-router-id                                            | Advertising router ID of the referenced LSA.                                                                                                                                                                                   | detail extensive |
| Prefix-count                                             | Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.                                                                                | detail extensive |
| Prefix                                                   | IPv6 address prefix.                                                                                                                                                                                                           | detail extensive |
| Prefix-options                                           | Option bit associated with the prefix.                                                                                                                                                                                         | detail extensive |
| Metric                                                   | Cost of this prefix. Expressed in the same units as the interface costs in the router LSAs.                                                                                                                                    | detail extensive |
| Gen timer                                                | How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .                                                                                                                                            | extensive        |
| Aging timer                                              | How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                   | extensive        |
| Installed <i>hh:mm:ss</i> ago                            | How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .                                                                                                                                             | extensive        |
| expires in <i>hh:mm:ss</i>                               | How long until the route expires, in the format <i>hours:minutes:seconds</i> .                                                                                                                                                 | extensive        |
| sent <i>hh:mm:ss</i> ago                                 | Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .                                                                         | extensive        |
| <i>n</i> Router LSAs                                     | Number of router LSAs in the link-state database.                                                                                                                                                                              | summary          |
| <i>n</i> Network LSAs                                    | Number of network LSAs in the link-state database.                                                                                                                                                                             | summary          |

Table 352: show ospf3 database Output Fields (*continued*)

| Field Name                         | Field Description                                                        | Level of Output |
|------------------------------------|--------------------------------------------------------------------------|-----------------|
| <i>n</i> InterArPfx LSAs           | Number of interarea-prefix LSAs in the link-state database.              | summary         |
| <i>n</i> InterArRtr LSAs           | Number of interarea-router LSAs in the link-state database.              | summary         |
| <i>n</i> IntraArPfx LSAs           | Number of intra-area-prefix LSAs in the link-state database.             | summary         |
| Externals                          | Display of the external LSA database.                                    | summary         |
| <i>n</i> Extern LSAs               | Number of external LSAs in the link-state database.                      | summary         |
| Interface<br><i>interface-name</i> | Name of the interface for which link-local LSA information is displayed. | summary         |
| <i>n</i> Link LSAs                 | Number of link LSAs in the link-state database.                          | summary         |

## Sample Output

### show ospf3 database brief

```

user@host> show ospf3 database brief
      OSPF3 link state database, area 0.0.0.0
      Type      ID          Adv Rtr      Seq          Age  Cksum  Len
      Router    0.0.0.1      10.255.4.85  0x80000003   885  0xa697  40
      Router    *0.0.0.1     10.255.4.93  0x80000002   953  0xc677  40
      InterArPfx *0.0.0.2     10.255.4.93  0x80000001   910  0xb96f  44
      InterArRtr *0.0.0.1     10.255.4.93  0x80000001   910  0xe159  32
      IntraArPfx *0.0.0.1     10.255.4.93  0x80000002   432  0x788f  72

      OSPF3 link state database, area 0.0.0.1
      Type      ID          Adv Rtr      Seq          Age  Cksum  Len
      Router    *0.0.0.1     10.255.4.93  0x80000003   916  0xea40  40
      Router    0.0.0.1     10.255.4.97  0x80000006   851  0xc95b  40
      Network    0.0.0.2     10.255.4.97  0x80000002   916  0x4598  32
      InterArPfx *0.0.0.1     10.255.4.93  0x80000002   117  0xa980  44
      InterArPfx *0.0.0.2     10.255.4.93  0x80000002    62  0xd47e  44
      NSSA       0.0.0.1     10.255.4.97  0x80000002   362  0x45ee  44
      IntraArPfx 0.0.0.1     10.255.4.97  0x80000006   851  0x2f77  52

      OSPF3 AS SCOPE link state database
      Type      ID          Adv Rtr      Seq          Age  Cksum  Len
      Extern     0.0.0.1     10.255.4.85  0x80000002    63  0x9b86  44
      Extern     *0.0.0.1     10.255.4.93  0x80000001   910  0x59c9  44

      OSPF3 Link-Local link state database, interface ge-1/3/0.0
      Type      ID          Adv Rtr      Seq          Age  Cksum  Len
      Link       *0.0.0.2     10.255.4.93  0x80000003   916  0x4dab  64

```

### show ospf3 database extensive

```

user@host> show ospf3 database extensive
      OSPF3 link state database, area 0.0.0.0
      Type      ID          Adv Rtr      Seq          Age  Cksum  Len
      Router    0.0.0.1     10.255.4.85  0x80000003  1028  0xa697  40

```



```

bits 0x2, Options 0x13
Type PointToPoint (1), Metric 10
  Loc-If-Id 2, Nbr-If-Id 3, Nbr-Rtr-Id 10.255.4.93
Aging timer 00:42:51
Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
Router *0.0.0.1 10.255.4.93 0x80000002 1096 0xc677 40
bits 0x3, Options 0x13
Type PointToPoint (1), Metric 10
  Loc-If-Id 3, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.85
Gen timer 00:00:40
Aging timer 00:41:44
Installed 00:18:16 ago, expires in 00:41:44, sent 00:18:14 ago
Ours
InterArPfx *0.0.0.2 10.255.4.93 0x80000001 1053 0xb96f 44
Prefix feee::10:10:2:0/126
Prefix-options 0x0, Metric 10
Gen timer 00:17:02
Aging timer 00:42:26
Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
Ours
InterArPfx *0.0.0.3 10.255.4.93 0x80000001 1053 0x71d3 44
Prefix feee::10:255:4:97/128
Prefix-options 0x0, Metric 10
Gen timer 00:21:07
Aging timer 00:42:26
Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
Ours
InterArRtr *0.0.0.1 10.255.4.93 0x80000001 1053 0xe159 32
Dest-router-id 10.255.4.97, Options 0x19, Metric 10
Gen timer 00:29:18
Aging timer 00:42:26
Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
Ours
IntraArPfx 0.0.0.1 10.255.4.85 0x80000002 1028 0x2403 72
Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.85
Prefix-count 2
Prefix feee::10:255:4:85/128
  Prefix-options 0x2, Metric 0
Prefix feee::10:10:1:0/126
  Prefix-options 0x0, Metric 10
Aging timer 00:42:51
Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
IntraArPfx *0.0.0.1 10.255.4.93 0x80000002 575 0x788f 72
Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.93
Prefix-count 2
Prefix feee::10:255:4:93/128
  Prefix-options 0x2, Metric 0
Prefix feee::10:10:1:0/126
  Prefix-options 0x0, Metric 10
Gen timer 00:33:23
Aging timer 00:50:24
Installed 00:09:35 ago, expires in 00:50:25, sent 00:09:33 ago
OSPF3 link state database, area 0.0.0.1
Type ID Adv Rtr Seq Age Cksum Len
Router *0.0.0.1 10.255.4.93 0x80000003 1059 0xea40 40
bits 0x3, Options 0x19
Type Transit (2), Metric 10
  Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
Gen timer 00:08:51
Aging timer 00:42:20
Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago

```

```

Router      0.0.0.1          10.255.4.97      0x80000006   994  0xc95b  40
  bits 0x2, Options 0x19
  Type Transit (2), Metric 10
    Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
  Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
Network     0.0.0.2          10.255.4.97      0x80000002   1059 0x4598  32
  Options 0x11
  Attached router 10.255.4.97
  Attached router 10.255.4.93
  Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
InterArPfx *0.0.0.1          10.255.4.93      0x80000002   260  0xa980  44
  Prefix feee::10:10:1:0/126
  Prefix-options 0x0, Metric 10
  Gen timer 00:45:39
  Aging timer 00:55:39
  Installed 00:04:20 ago, expires in 00:55:40, sent 00:04:18 ago
  Ours
InterArPfx *0.0.0.2          10.255.4.93      0x80000002   205  0xd47e  44
  Prefix feee::10:255:4:93/128
  Prefix-options 0x0, Metric 0
  Gen timer 00:46:35
  Aging timer 00:56:35
  Installed 00:03:25 ago, expires in 00:56:35, sent 00:03:23 ago
  Ours
InterArPfx *0.0.0.3          10.255.4.93      0x80000001   1089 0x9bbb  44
  Prefix feee::10:255:4:85/128
  Prefix-options 0x0, Metric 10
  Gen timer 00:04:46
  Aging timer 00:41:51
  Installed 00:18:09 ago, expires in 00:41:51, sent 00:17:43 ago
  Ours
NSSA        0.0.0.1          10.255.4.97      0x80000002   505  0x45ee  44
  Prefix feee::200:200:1:0/124
  Prefix-options 0x8, Metric 10, Type 2,
  Aging timer 00:51:35
  Installed 00:08:22 ago, expires in 00:51:35, sent 02:37:54 ago
IntraArPfx  0.0.0.1          10.255.4.97      0x80000006   994  0x2f77  52
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:255:4:97/128
    Prefix-options 0x2, Metric 0
  Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
IntraArPfx  0.0.0.3          10.255.4.97      0x80000002   1059 0x4446  52
  Ref-lsa-type Network, Ref-lsa-id 0.0.0.2, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:10:2:0/126
    Prefix-options 0x0, Metric 0
  Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
  OSPF3 AS SCOPE link state database
  Type      ID              Adv Rtr          Seq            Age  Cksum  Len
Extern     0.0.0.1          10.255.4.85      0x80000002     206  0x9b86  44
  Prefix feee::100:100:1:0/124
  Prefix-options 0x0, Metric 20, Type 2,
  Aging timer 00:56:34
  Installed 00:03:23 ago, expires in 00:56:34, sent 02:37:54 ago
Extern     *0.0.0.1          10.255.4.93      0x80000001     1053 0x59c9  44
  Prefix feee::200:200:1:0/124

```

```

Prefix-options 0x0, Metric 10, Type 2,
Gen timer 00:25:12
Aging timer 00:42:26
Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago

```

```

OSPF3 Link-Local link state database, interface ge-1/3/0.0
Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link      *0.0.0.2      10.255.4.93  0x80000003  1059 0x4dab  64
fe80::290:69ff:fe39:1cdb
Options 0x11, priority 128
Prefix-count 1
Prefix feee::10:10:2:0/126 Prefix-options 0x0
Gen timer 00:12:56
Aging timer 00:42:20
Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Link      0.0.0.2      10.255.4.97  0x80000003  205  0xa87d  64
fe80::290:69ff:fe38:883e
Options 0x11, priority 128
Prefix-count 1
Prefix feee::10:10:2:0/126 Prefix-options 0x0
Aging timer 00:56:35
Installed 00:03:22 ago, expires in 00:56:35, sent 02:37:54 ago

```

```

OSPF3 Link-Local link state database, interface so-2/2/0.0
Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link      0.0.0.2      10.255.4.85  0x80000002  506  0x42bb  64
fe80::280:42ff:fe10:f169
Options 0x13, priority 128
Prefix-count 1
Prefix feee::10:10:1:0/126 Prefix-options 0x0
Aging timer 00:51:34
Installed 00:08:23 ago, expires in 00:51:34, sent 02:37:54 ago
Link      *0.0.0.3      10.255.4.93  0x80000002  505  0x6b7a  64
fe80::280:42ff:fe10:f177
Options 0x13, priority 128
Prefix-count 1
Prefix feee::10:10:1:0/126 Prefix-options 0x0
Gen timer 00:37:28
Aging timer 00:51:35
Installed 00:08:25 ago, expires in 00:51:35, sent 00:08:23 ago
Ours

```

### show ospf3 database summary

```

user@host> show ospf3 database summary
Area 0.0.0.0:
  2 Router LSAs
  1 InterArPfx LSAs
  1 InterArRtr LSAs
  1 IntraArPfx LSAs
Area 0.0.0.1:
  2 Router LSAs
  1 Network LSAs
  2 InterArPfx LSAs
  1 NSSA LSAs
  1 IntraArPfx LSAs
Externals:
  2 Extern LSAs
Interface ge-1/3/0.0:
  1 Link LSAs
Interface lo0.0:

```

Interface so-2/2/0.0:  
1 Link LSAs

## PART 18

# RIP

- [Overview on page 3343](#)
- [Configuration on page 3347](#)
- [Administration on page 3393](#)



## CHAPTER 57

# Overview

- [Layer 3 Protocols on page 3343](#)

## Layer 3 Protocols

- [Layer 3 Protocols Supported on EX Series Switches on page 3343](#)
- [Layer 3 Protocols Not Supported on EX Series Switches on page 3344](#)

## Layer 3 Protocols Supported on EX Series Switches

EX Series switches support the Junos OS Layer 3 features and configuration statements listed in [Table 314 on page 2939](#):

**Table 353: Supported Junos OS Layer 3 Protocol Statements and Features**

| Protocol           | Notes                                                                                                                                            | For More Information                                             |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| BGP                | Fully supported.                                                                                                                                 | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| BFD                | Fully supported.                                                                                                                                 | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| ICMP               | Fully supported.                                                                                                                                 | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| IGMPv1, v2, and v3 | Fully supported.                                                                                                                                 | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |
| IS-IS              | Supported, with the exceptions noted in " <a href="#">Layer 3 Protocols Not Supported on EX Series Switches</a> " on <a href="#">page 2940</a> . | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| MLD                | Fully supported (MLD versions 1 and 2).                                                                                                          | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |
| MPLS               | Supported, with the exceptions noted in " <a href="#">Layer 3 Protocols Not Supported on EX Series Switches</a> " on <a href="#">page 2940</a> . | <a href="#">Junos OS MPLS Applications Configuration Guide</a>   |
| OSPFv1, v2 and v3  | Supported, with the exceptions noted in " <a href="#">Layer 3 Protocols Not Supported on EX Series Switches</a> " on <a href="#">page 2940</a> . | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| PIM                | Fully supported.                                                                                                                                 | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |

Table 353: Supported Junos OS Layer 3 Protocol Statements and Features (*continued*)

| Protocol | Notes                                                                                                | For More Information                                                                                                     |
|----------|------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| PPM      | Supported. See <i>EX Series Switch Software Features Overview</i> for specific platform information. | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| RIP      | Fully supported.                                                                                     | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| RIPng    | Fully supported.                                                                                     | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| SNMP     | Fully supported.                                                                                     | <a href="#">Junos OS Network Management Configuration Guide</a>                                                          |
| VRRP     | Fully supported.                                                                                     | See “Understanding VRRP on EX Series Switches” on page 2501. See also <a href="#">Junos OS High Availability Guide</a> . |

- Related Documentation**
- [Layer 3 Protocols Not Supported on EX Series Switches](#) on page 2940
  - [EX Series Switch Software Features Overview](#)

## Layer 3 Protocols Not Supported on EX Series Switches

EX Series switches do not support the Junos OS Layer 3 protocols and features listed in [Table 315 on page 2940](#):

Table 354: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported

| Feature                                                                                                         | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DVMRP                                                                                                           | <ul style="list-style-type: none"> <li>• <b>dvmrp</b> and subordinate statements</li> </ul>                                                                                                                                                                                                          |
| Flow aggregation (cflowd)                                                                                       | <ul style="list-style-type: none"> <li>• <b>cflow</b> and subordinate statements</li> </ul>                                                                                                                                                                                                          |
| IPsec                                                                                                           | <ul style="list-style-type: none"> <li>• <b>[edit services]</b> statements related to IPsec</li> </ul>                                                                                                                                                                                               |
| IS-IS: <ul style="list-style-type: none"> <li>• ES-IS</li> <li>• IPv6 in multicast routing protocols</li> </ul> | <ul style="list-style-type: none"> <li>• <b>clns-routing</b> statement</li> <li>• <b>ipv6-multicast</b> statement</li> <li>• <b>lsp-interval</b> statement</li> <li>• <b>label-switched-path</b> statement</li> <li>• <b>lsp-lifetime</b> statement</li> <li>• <b>te-metric</b> statement</li> </ul> |
| Logical routers                                                                                                 | <ul style="list-style-type: none"> <li>• <b>logical-routers</b> and subordinate statements</li> </ul>                                                                                                                                                                                                |



**Table 354: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)**

| Feature                                                                                                                                                                                                                                                                                                                                                                                                                                      | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MPLS: <ul style="list-style-type: none"> <li>Fast Reroute (FRR)</li> <li>Label Distribution Protocol (LDP) (except on EX8200 switches)</li> <li>Layer 3 VPNs (except on EX8200 switches)</li> <li>Multiprotocol BGP (MP-BGP) for VPN-IPv4 family</li> <li>Pseudowire emulation (PWE3)</li> <li>Routing policy statements related to Layer 3 VPNs and MPLS (except on EX8200 switches)</li> <li>Virtual Private LAN Service (VPLS)</li> </ul> | <ul style="list-style-type: none"> <li><b>ldp</b> and all subordinate statements (except on EX8200 switches)</li> </ul>                                                                                                                                                                                                                                                                                           |
| Network Address Translation (NAT)                                                                                                                                                                                                                                                                                                                                                                                                            | <ul style="list-style-type: none"> <li><b>nat</b> and subordinate statements</li> <li>Policy statements related to NAT</li> </ul>                                                                                                                                                                                                                                                                                 |
| OSPF                                                                                                                                                                                                                                                                                                                                                                                                                                         | <ul style="list-style-type: none"> <li><b>demand-circuit</b> statement</li> <li><b>label-switched-path</b> and subordinate statements</li> <li><b>neighbor</b> statement within an OSPF area</li> <li><b>peer-interface</b> and subordinate statements within an OSPF area</li> <li><b>sham-link</b> statement</li> <li><b>te-metric</b> statement</li> </ul>                                                     |
| PPM                                                                                                                                                                                                                                                                                                                                                                                                                                          | <ul style="list-style-type: none"> <li>Not supported on EX2200 and EX3300 switches</li> </ul>                                                                                                                                                                                                                                                                                                                     |
| Routing instances: <ul style="list-style-type: none"> <li>Routing instance forwarding</li> </ul>                                                                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li><b>l2vpn</b> and subordinate statements (except on EX4500, EX4550, and EX8200 switches)</li> <li><b>ldp</b> and subordinate statements (except on EX8200 switches)</li> <li><b>vpls</b> and subordinate statements</li> </ul>                                                                                                                                              |
| Routed VLAN interfaces (RVIs)                                                                                                                                                                                                                                                                                                                                                                                                                | <ul style="list-style-type: none"> <li><b>family mpls</b> statement</li> </ul>                                                                                                                                                                                                                                                                                                                                    |
| SAP and SDP                                                                                                                                                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li><b>sap</b> and all subordinate statements</li> </ul>                                                                                                                                                                                                                                                                                                                       |
| General routing options in the <b>routing-options</b> hierarchy: <ul style="list-style-type: none"> <li>MPLS and label-switched-paths</li> </ul>                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li><b>auto-export</b> and subordinate statements</li> <li><b>dynamic-tunnels</b> and subordinate statements</li> <li><b>lsp-next-hop</b> and subordinate statements</li> <li><b>multicast</b> and subordinate statements</li> <li><b>p2mp-lsp-next-hop</b> and subordinate statements</li> <li><b>route-distinguisher-id</b> statement (except on EX8200 switches)</li> </ul> |

Table 354: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)

| Feature                                                                    | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Traffic sampling and forwarding in the <b>forwarding-options</b> hierarchy | <ul style="list-style-type: none"> <li>• <b>accounting</b> and subordinate statements</li> <li>• <b>family mpls</b> and <b>family multiservice</b> under <b>hash-key</b> hierarchy</li> <li>• Under <b>monitoring group-name</b> family <b>inet output</b> hierarchy: <ul style="list-style-type: none"> <li>• <b>cflowd</b> statement</li> <li>• <b>export-format-cflowd-version-5</b> statement</li> <li>• <b>flow-active-timeout</b> statement</li> <li>• <b>flow-export-destination</b> statement</li> <li>• <b>flow-inactive-timeout</b> statement</li> <li>• <b>interface</b> statement</li> </ul> </li> <li>• <b>port-mirroring</b> statement (On EX Series switches, port mirroring is implemented using the <b>analyzer</b> statement.)</li> <li>• <b>sampling</b> and subordinate statements</li> </ul> |

- Related Documentation**
- [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)
  - *EX Series Switch Software Features Overview*

# Configuration

- [Configuration Tasks on page 3347](#)
- [Configuration Statements: RIP on page 3350](#)
- [Configuration Statements: RIPng on page 3376](#)

## Configuration Tasks

---

- [Configuring a RIP Network \(J-Web Procedure\) on page 3347](#)

### Configuring a RIP Network (J-Web Procedure)



NOTE: This topic applies only to the J-Web Application package.

You can use the J-Web interface to create RIP networks.

To configure a RIP network:

1. Select **Configure** > **Routing** > **RIP**.



NOTE: After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:
  - **Add**—Configures a RIP instance. Enter information into the RIP Configuration page as described in [Table 355 on page 3348](#).
  - **Edit**—Modifies an existing RIP instance. Enter information into the configuration page for RIP as described in [Table 355 on page 3348](#).
  - **Delete**—Deletes an existing RIP instance.
4. To modify RIP global settings, click **Edit**. Enter information in the configuration as described in [Table 356 on page 3348](#).

Table 355: RIP Routing Configuration Summary

| Field                  | Function                                                                                                              | Your Action                                                                                                                                                                                                          |
|------------------------|-----------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>General tab</b>     |                                                                                                                       |                                                                                                                                                                                                                      |
| Routing instance name  | Specifies a name for the routing instance.                                                                            | Type or select and edit the name.                                                                                                                                                                                    |
| Preference             | Specifies the preference of external routes learned by RIP as compared to those learned from other routing protocols. | Type or select and edit the value.                                                                                                                                                                                   |
| Metric Out             | Specifies the metric value to add to routes transmitted to the neighbor.                                              | Type or select and edit the value.                                                                                                                                                                                   |
| Update interval        | Specifies an update time interval to periodically send out routes learned by RIP to neighbors.                        | Type or select and edit the value.                                                                                                                                                                                   |
| Route timeout          | Specifies the route timeout interval for RIP.                                                                         | Type or select and edit the value.                                                                                                                                                                                   |
| <b>Policies tab</b>    |                                                                                                                       |                                                                                                                                                                                                                      |
| Import Policy          | Applies one or more policies to routes being imported into the local routing device from the neighbors.               | <p>Click <b>Add</b> to add an import policy.</p> <p>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.</p> <p>Click <b>Remove</b> to remove an import policy.</p> |
| Export Policy          | Applies a policy to routes being exported to the neighbors.                                                           | <p>Click <b>Add</b> to add an export policy.</p> <p>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.</p> <p>Click <b>Remove</b> to remove an export policy.</p> |
| <b>Neighbors tab</b>   |                                                                                                                       |                                                                                                                                                                                                                      |
| RIP-Enabled Interfaces | Selects the interfaces to be associated with the RIP instance.                                                        | <p>To enable RIP on an interface, click the check box next to the interface name.</p> <p>Click Edit if you want to modify an interface's settings.</p>                                                               |

Table 356: Edit RIP Global Settings

| Field                 | Function                                                                                        | Your Action                        |
|-----------------------|-------------------------------------------------------------------------------------------------|------------------------------------|
| <b>General tab</b>    |                                                                                                 |                                    |
| Send                  | Specifies RIP send options.                                                                     | Select a value.                    |
| Receive               | Configure RIP receive options.                                                                  | Select a value.                    |
| Route timeout (sec)   | Specifies the route timeout interval for RIP.                                                   | Type a value.                      |
| Update interval (sec) | Specifies the update time interval to periodically send out routes learned by RIP to neighbors. | Type or select and edit the value. |

Table 356: Edit RIP Global Settings (*continued*)

| Field               | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Your Action                                                                                                                                                                                                          |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hold timeout (sec)  | Specifies the time period the expired route is retained in the routing table before being removed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Type or select and edit the value.                                                                                                                                                                                   |
| Metric in           | Specifies the metric to add to incoming routes when advertising into RIP routes that were learned from other protocols.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Type or select and edit the value.                                                                                                                                                                                   |
| RIB Group           | Specifies a routing table group to install RIP routes into multiple routing tables.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Select and edit the name of the routing table group.                                                                                                                                                                 |
| Message size        | Specifies the number of route entries to be included in every RIP update message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Type or select and edit the value.                                                                                                                                                                                   |
| Check Zero          | <p>Specifies whether the reserved fields in a RIP packet are zero. Options are:</p> <ul style="list-style-type: none"> <li>• <b>check-zero</b>—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.</li> <li>• <b>no-check-zero</b>—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.</li> </ul> | Select a value.                                                                                                                                                                                                      |
| Graceful switchover | Configures graceful switchover for OSPF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <p>To disable graceful restart, select <b>Disable</b>.</p> <p>Type or select and edit the estimated time for the restart to finish, in seconds.</p>                                                                  |
| Authentication Type | <p>Specifies the type of authentication for RIP route queries received on an interface. Options are:</p> <ul style="list-style-type: none"> <li>• None</li> <li>• MD5</li> <li>• Simple</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <p>Select the authentication type.</p> <p>Enter the authentication key for MD5.</p>                                                                                                                                  |
| <b>Policies tab</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                                                                                                      |
| Import Policy       | Applies one or more policies to routes being imported into the local routing device from the neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <p>Click <b>Add</b> to add an import policy.</p> <p>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.</p> <p>Click <b>Remove</b> to remove an import policy.</p> |

Table 356: Edit RIP Global Settings (*continued*)

| Field                    | Function                                                                       | Your Action                                                                                                                                                                                                   |
|--------------------------|--------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Export Policy            | Applies a policy to routes being exported to the neighbors.                    | Click <b>Add</b> to add an export policy.<br><br>Click <b>Move up</b> or <b>Move down</b> to move the selected policy up or down the list of policies.<br><br>Click <b>Remove</b> to remove an export policy. |
| <b>Trace Options tab</b> |                                                                                |                                                                                                                                                                                                               |
| File Name                | Specifies the name of the file to receive the output of the tracing operation. | Type or select and edit the name.                                                                                                                                                                             |
| Number of Files          | Specifies the maximum number of trace files.                                   | Type or select and edit the name.                                                                                                                                                                             |
| File Size                | Specifies the maximum size for each trace file.                                | Type or select and edit the name.                                                                                                                                                                             |
| World Readable           | Specifies whether the trace file can be read by any user or not.               | Select <b>True</b> to allow any user to read the file.<br><br>Select <b>False</b> to disallow all users being able to read the file.                                                                          |
| Flags                    | Specifies the tracing operation to perform.                                    | Select a value from the list.                                                                                                                                                                                 |


- Related Documentation**
- [Monitoring RIP Routing Information on page 3393](#)
  - [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)

## Configuration Statements: RIP

- [any-sender on page 3351](#)
- [authentication-key \(Protocols RIP\) on page 3352](#)
- [authentication-type \(Protocols RIP\) on page 3353](#)
- [bfd-liveness-detection \(Protocols RIP\) on page 3354](#)
- [check-zero on page 3357](#)
- [export \(Protocols RIP\) on page 3358](#)
- [graceful-restart \(Protocols RIP\) on page 3359](#)
- [group \(Protocols RIP\) on page 3360](#)
- [holddown \(Protocols RIP\) on page 3362](#)
- [import \(Protocols RIP\) on page 3363](#)
- [message-size on page 3364](#)
- [metric-in \(Protocols RIP\) on page 3365](#)
- [metric-out \(Protocols RIP\) on page 3366](#)
- [neighbor \(Protocols RIP\) on page 3367](#)

- [preference \(Protocols RIP\) on page 3368](#)
- [receive \(Protocols RIP\) on page 3369](#)
- [rib-group \(Protocols RIP\) on page 3370](#)
- [rip on page 3370](#)
- [route-timeout \(Protocols RIP\) on page 3371](#)
- [send \(Protocols RIP\) on page 3372](#)
- [traceoptions \(Protocols RIP\) on page 3373](#)
- [update-interval \(Protocols RIP\) on page 3376](#)

## any-sender

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>any-sender;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Disable strict sender address checks.</p> <p>If the sender of a RIP message does not belong to the subnet of the interface, the message is discarded. This situation might cause problems with dropped packets when RIP is running on point-to-point interfaces, or when the addresses on the interfaces do not fall in the same subnet. You can resolve this by disabling strict address checks on the RIP traffic.</p>                                                                                                                                |
|                                 | <div>  <p><b>NOTE:</b> The <code>any-sender</code> statement is supported only for peer-to-peer interfaces.</p> </div>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## authentication-key (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-key password;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <i>rip</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <i>neighbor</i></code><br><code>    <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <i>rip</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>],</code><br><code>[edit protocols <i>rip</i>],</code><br><code>[edit protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <i>rip</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor</i></code><br><code>    <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Require authentication for RIP route queries received on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <i>password</i> —Authentication password. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Route Authentication for RIP</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |



## authentication-type (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-type type;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the type of authentication for RIP route queries received on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | If you do not include this statement and the <b>authentication-key</b> statement, RIP authentication is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>type</b>—Authentication type:</p> <ul style="list-style-type: none"> <li>• <b>md5</b>—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing device uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.</li> <li>• <b>none</b>—Disable authentication. If <b>none</b> is configured, the configured authentication key is ignored.</li> <li>• <b>simple</b>—Use a simple password. The password is included in the transmitted packet, which makes this method of authentication relatively insecure. The password can be from 1 through 16 contiguous letters or digits long.</li> </ul>                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Route Authentication for RIP</i></li> <li>• <a href="#">authentication-key on page 3352</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## bfd-liveness-detection (Protocols RIP)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>bfd-liveness-detection {<br/>  authentication {<br/>    algorithm <i>algorithm-name</i>;<br/>    key-chain <i>key-chain-name</i>;<br/>    loose-check;<br/>  }<br/>  detection-time {<br/>    threshold <i>milliseconds</i>;<br/>  }<br/>  minimum-interval <i>milliseconds</i>;<br/>  minimum-receive-interval <i>milliseconds</i>;<br/>  multiplier <i>number</i>;<br/>  no-adaptation;<br/>  transmit-interval {<br/>    minimum-interval <i>milliseconds</i>;<br/>    threshold <i>milliseconds</i>;<br/>  }<br/>  version (1   automatic);<br/>}</pre>                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>  rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],<br/>[edit protocols rip <b>group</b> <i>group-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b><br/>  <i>neighbor-name</i>]</pre>                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.0.</p> <p>Options <b>detection-time threshold</b> and <b>transmit-interval threshold</b> introduced in Junos OS Release 8.2.</p> <p>Support for logical systems introduced in Junos OS Release 8.3.</p> <p>Option <b>no-adaptation</b> introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Options <b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> introduced in Junos OS Release 9.6.</p> <p>Options <b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> |
| <b>Description</b>         | Configure bidirectional failure detection timers and authentication.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>             | <p><b>authentication algorithm <i>algorithm-name</i></b>—Configure the algorithm used to authenticate the specified BFD session: <b>simple-password</b>, <b>keyed-md5</b>, <b>keyed-sha-1</b>, <b>meticulous-keyed-md5</b>, or <b>meticulous-keyed-sha-1</b>.</p> <p><b>authentication key-chain <i>key-chain-name</i></b>—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the <b>authentication-key-chains key-chain</b> statement at the <b>[edit security]</b> hierarchy level.</p>                                                                                                                                                                                                                     |

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication is not configured at both ends of the BFD session.

**detection-time threshold *milliseconds***—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**minimum-interval *milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

**Range:** 1 through 255,000 milliseconds

**minimum-receive-interval *milliseconds***—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

**Range:** 1 through 255,000 milliseconds

**multiplier *number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation**—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval *milliseconds***—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

**Default:** automatic

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |


- Related Documentation**
- *Example: Configuring BFD for RIP*
  - *Example: Configuring BFD Authentication for RIP*

## check-zero

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (check-zero   no-check-zero);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <i>rip</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <i>neighbor neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>rip</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor neighbor-name</i>],</p> <p>[edit protocols <i>rip</i>],</p> <p>[edit protocols rip group <i>group-name</i> <i>neighbor neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>rip</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor neighbor-name</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Some of the reserved fields in RIP version 1 packets must be zero, whereas in RIP version 2 packets, most of these reserved fields can contain nonzero values. By default, RIP discards version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.</p> <p>If you find that you are receiving RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero, you can configure RIP to receive these packets even though they are being sent in violation of the specifications in RFC 1058 and RFC 2453.</p> <p>Check whether the reserved fields in a RIP packet are zero:</p> <ul style="list-style-type: none"> <li>• <b>check-zero</b>—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.</li> <li>• <b>no-check-zero</b>—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.</li> </ul> |
| <b>Default</b>                  | check-zero                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring RIP</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## export (Protocols RIP)

---

|                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                           | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                  | <code>[edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <b>rip</b> <b>group</b> <i>group-name</i>],</code><br><code>[edit protocols rip <b>group</b> <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i>]</code>          |
| <b>Release Information</b>                                                                                                                                                                              | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                                                                                                                                                                      | <p>Apply a policy to routes being exported to the neighbors.</p> <p>By default, RIP does not export routes it has learned to its neighbors. To enable RIP to export routes, apply one or more export policies.</p> <p>If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the <b>metric-in</b> and <b>metric-out</b> statements.</p> |
| <div> <b>NOTE:</b> The export policy on RIP does not support manipulating routing information of the next hop.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                                                                                                                                                                          | <b><i>policy-names</i></b> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b>                                                                                                                                                                         | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                                            | <ul style="list-style-type: none"><li>• <i>Example: Configuring RIP</i></li><li>• <a href="#">import on page 3363</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                         |

---

## graceful-restart (Protocols RIP)

---

|                                 |                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>graceful-restart {<br/>  disable;<br/>  restart-time <i>seconds</i>;<br/>}</pre>                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>              | Configure graceful restart for RIP.                                                                                        |
| <b>Options</b>                  | <b>disable</b> —Disables graceful restart for RIP.<br><br>The remaining statement is explained separately.                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Junos OS High Availability Library for Routing Devices</i></li></ul>            |

## group (Protocols RIP)

---

```
Syntax  group group-name {
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            transmit-interval {
                threshold milliseconds;
                minimum-interval milliseconds;
            }
            multiplier number;
            version (0 | 1 | automatic);
        }
        demand-circuit;
        export policy;
        max-retrans-time seconds;
        metric-out metric;
        preference number;
        route-timeout seconds;
        update-interval seconds;
        neighbor neighbor-name {
            authentication-key password;
            authentication-type type;
            bfd-liveness-detection {
                authentication {
                    algorithm algorithm-name;
                    key-chain key-chain-name;
                    loose-check;
                }
                detection-time {
                    threshold milliseconds;
                }
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                transmit-interval {
                    threshold milliseconds;
                    minimum-interval milliseconds;
                }
                multiplier number;
                version (0 | 1 | automatic);
            }
            (check-zero | no-check-zero);
            demand-circuit;
            import policy-name;
            max-retrans-time seconds;
            message-size number;
```



```

metric-in metric;
metric-out metric;
receive receive-options;
route-timeout seconds;
send send-options;
update-interval seconds;
}
}

```

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                          |
| <b>Description</b>              | Configure a set of RIP neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group. Each group must contain at least one neighbor. You should create a group for every export policy.                                                                                                 |
| <b>Options</b>                  | <p><b><i>group-name</i></b>—Name of a group, up to 16 characters long.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring RIP</i></li> </ul>                                                                                                                                                                                                                                                                                    |

## holddown (Protocols RIP)

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>holddown seconds;</code>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ]                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Configure how long the expired route is retained in the routing table before being removed.</p> <p>When the hold-down timer runs on RIP demand circuits, routes are advertised as unreachable on other interfaces. When the hold-down timer expires, the route is removed from the routing table if all destinations detect that the route is unreachable or the remaining destinations are down.</p> |
| <b>Options</b>                  | <b>seconds</b> —Estimated time to wait before making updates to the routing table.<br><b>Range:</b> 10 through 180 seconds<br><b>Default:</b> 120 seconds                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring RIP Timers</i></li><li>• <i>RIP Demand Circuits Overview</i></li></ul>                                                                                                                                                                                                                                                                   |

## import (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Apply one or more policies to routes being imported by the local routing device from neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Applying Policies to RIP Routes Imported from Neighbors</i></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> <li>• <a href="#">export on page 3358</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## message-size

---

|                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                     | <code>message-size <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                            | <code>[edit logical-systems <i>logical-system-name</i> protocols <i>rip</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <i>neighbor</i></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <i>rip</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>],</code><br><code>[edit protocols <i>rip</i>],</code><br><code>[edit protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <i>rip</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor</i></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                        | Statement introduced before Junos OS Release 7.4.<br>Statement for SRX Series devices introduced in Junos OS Release 9.5.<br>Statement for J Series platform introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                                                                                                                                                                                                                                                                | Specify the number of route entries to be included in every RIP update message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <div> <b>TIP:</b> To ensure interoperability with other vendors' equipment, use the standard of 25 route entries per message. Do not change the default number of route entries in a RIP update message.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                                                                                                                                                                                                                                                                    | <b><i>number</i></b> —Number of route entries per update message.<br><b>Range:</b> 25 through 255 entries<br><b>Default:</b> 25 entries                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                   | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                      | <ul style="list-style-type: none"><li>• <i>Example: Configuring RIP</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## metric-in (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric-in <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify the metric to add to incoming routes when the routing device advertises into RIP routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIP routes learned through a specific neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><i>metric</i>—Metric value.</p> <p><b>Range:</b> 1 through 16</p> <p><b>Default:</b> 1</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the Metric Value Added to Imported RIP Routes</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## metric-out (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric-out <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIP routes sent from this neighbor.</p> <p>If you have included the <b>export</b> statement, RIP exports routes it has learned to the neighbors configured by including the <b>neighbor</b> statement.</p> <p>The metric associated with a RIP route (unless modified by an export policy) is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with a <b>metric-in</b> value of 2 is advertised with a combined metric of 7 when advertised to RIP neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured for that group with the <b>metric-out</b> statement.</p> <p>The metric for a route can be modified with an export policy. That metric is seen when the route is exported to the next hop.</p> <p>To increase the metric for routes advertised outside a group, include the <b>metric-out</b> statement.</p> |
| <b>Options</b>                  | <b><i>metric</i></b> —Metric value.<br><b>Range:</b> 1 through 16<br><b>Default:</b> 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Examples: Controlling Traffic with Metrics in a RIP Network</i>50</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## neighbor (Protocols RIP)

**Syntax** `neighbor neighbor-name {`  
     `authentication-key password;`  
     `authentication-type type;`  
     `bfd-liveness-detection {`  
         `authentication {`  
             `algorithm algorithm-name;`  
             `key-chain key-chain-name;`  
             `loose-check;`  
         `}`  
         `detection-time {`  
             `threshold milliseconds;`  
         `}`  
     `minimum-interval milliseconds;`  
     `minimum-receive-interval milliseconds;`  
     `transmit-interval {`  
         `threshold milliseconds;`  
         `minimum-interval milliseconds;`  
     `}`  
     `multiplier number;`  
     `version (0 | 1 | automatic);`  
     `}`  
     `(check-zero | no-check-zero);`  
     `demand-circuit;`  
     `import policy-name;`  
     `max-retrans-time seconds;`  
     `message-size number;`  
     `metric-in metric;`  
     `metric-out metric;`  
     `receive receive-options;`  
     `route-timeout seconds;`  
     `send send-options;`  
     `update-interval seconds;`  
     `}`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols rip **group** *group-name*],  
     [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
         rip **group** *group-name*],  
     [edit protocols rip **group** *group-name*],  
     [edit routing-instances *routing-instance-name* protocols rip **group** *group-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
     Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure neighbor-specific RIP parameters, thereby overriding the defaults set for the routing device.

**Options** *neighbor-name*—Name of an interface over which a routing device communicates to its neighbors.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring RIP*

---

## preference (Protocols RIP)

---

**Syntax** `preference preference;`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols rip **group** *group-name*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
rip **group** *group-name*],  
[edit protocols rip **group** *group-name*],  
[edit routing-instances *routing-instance-name* protocols rip **group** *group-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Specify the preference of external routes learned by RIP as compared to those learned from other routing protocols.

By default, Junos OS assigns a preference of 100 to routes that originate from RIP. When Junos OS determines a route's preference to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.

**Options** *preference*—Preference value. A lower value indicates a more preferred route.  
**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )  
**Default:** 100

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Route Preferences Overview*



## receive (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>receive receive-options;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure RIP receive options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><i>receive-options</i>—One of the following:</p> <ul style="list-style-type: none"> <li>• <b>both</b>—Accept both RIP version 1 and version 2 packets.</li> <li>• <b>none</b>—Do not receive RIP packets.</li> <li>• <b>version-1</b>—Accept only RIP version 1 packets.</li> <li>• <b>version-2</b>—Accept only RIP version 2 packets.</li> </ul> <p><b>Default:</b> <b>both</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets</i></li> <li>• <a href="#">send on page 3372</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## rib-group (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rib-group group-name;</code>                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Install RIP routes into multiple routing tables by configuring a routing table group.                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <i>group-name</i> —Name of the routing table group.                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Redistributing Routes Between Two RIP Instances</i></li></ul>                                                                                                                                                                                                                                       |

## rip

---

|                                 |                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rip {...}</code>                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                      |
| <b>Description</b>              | Enable RIP routing on the routing device.                                                                                                                                                                                                                           |
| <b>Default</b>                  | RIP is disabled on the routing device.                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring RIP</i></li></ul>                                                                                                                                                                                   |

## route-timeout (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>route-timeout seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip <a href="#">group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip <a href="#">group</a> <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group</a> <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip <a href="#">group</a> <i>group-name</i>],</p> <p>[edit protocols rip <a href="#">group</a> <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group</a> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group</a> <i>group-name</i> neighbor <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the route timeout interval for RIP. If a route is not refreshed after being installed in the routing table by the specified timeout interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>seconds</b>—Estimated time to wait before making updates to the routing table.</p> <p><b>Range:</b> 30 through 360 seconds</p> <p><b>Default:</b> 180 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring RIP Timers</i></li> <li>• <i>RIP Demand Circuits Overview</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## send (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>send <i>send-options</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>    <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>    <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure RIP send options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <i>send-options</i> —One of the following: <ul style="list-style-type: none"><li>• <b>broadcast</b>—Broadcast RIP version 2 packets (RIP version 1 compatible).</li><li>• <b>multicast</b>—Multicast RIP version 2 packets. This is the default.</li><li>• <b>none</b>—Do not send RIP updates.</li><li>• <b>version-1</b>—Broadcast RIP version 1 packets.</li></ul> <b>Default:</b> multicast                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets</i></li><li>• <a href="#">receive on page 3369</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## traceoptions (Protocols RIP)

|                            |                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>                                                                                                                                       |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>         | Set RIP protocol-level tracing options.                                                                                                                                                                                                                                                                                                                 |



**NOTE:** The `traceoptions` statement is not supported on QFabric systems.

**Default** The default RIP protocol-level trace options are inherited from the global `traceoptions` statement.

**Options** **disable**—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

**file *filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIP tracing output in the file `/var/log/rip-log`.

**files *number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

### RIP Tracing Options

- **auth**—RIP authentication
- **error**—RIP error packets

- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets such as request, poll, and poll entry packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

#### Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                              |                                                                                                  |
|------------------------------|--------------------------------------------------------------------------------------------------|
| <b>Required Privilege</b>    | routing—To view this statement in the configuration.                                             |
| <b>Level</b>                 | routing-control—To add this statement to the configuration.                                      |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Example: Tracing RIP Protocol Traffic</i></li> </ul> |

## update-interval (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>update-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> group <i>group-name</i> neighbor</code><br><code>    <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit protocols <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i> neighbor</code><br><code>    <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the interval at which routes learned by RIP are sent to neighbors. This timer controls the interval between routing updates. This timer is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can happen if all routing devices update their neighbors simultaneously.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>seconds</b> —Estimated time to wait before making updates to the routing table.<br><b>Range:</b> 10 through 60 seconds<br><b>Default:</b> 30 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring RIP Timers</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Configuration Statements: RIPng

---

- [export \(Protocols RIPng\)](#) on page 3377
- [graceful-restart \(Protocols RIPng\)](#) on page 3378
- [group \(Protocols RIPng\)](#) on page 3379
- [holddown \(Protocols RIPng\)](#) on page 3380
- [import \(Protocols RIPng\)](#) on page 3381



- [metric-in \(Protocols RIPng\) on page 3382](#)
- [metric-out \(Protocols RIPng\) on page 3383](#)
- [neighbor \(Protocols RIPng\) on page 3384](#)
- [preference \(Protocols RIPng\) on page 3385](#)
- [receive \(Protocols RIPng\) on page 3386](#)
- [ripng on page 3387](#)
- [route-timeout \(Protocols RIPng\) on page 3387](#)
- [send \(Protocols RIPng\) on page 3388](#)
- [traceoptions \(Protocols RIPng\) on page 3389](#)
- [update-interval \(Protocols RIPng\) on page 3391](#)

## export (Protocols RIPng)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols ripng <a href="#">group group-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i>],</p> <p>[edit protocols ripng <a href="#">group group-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i>]</p>                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Apply a policy or list of policies to routes being exported to the neighbors.</p> <p>By default, RIPng does not export routes it has learned to its neighbors. To have RIPng export routes, apply one or more export policies. To apply export policies and to filter routes being exported from the local routing device to its neighbors, include the <b>export</b> statement and list the name of the policy to be evaluated.</p> <p>You can define one or more export policies. If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the <b>metric-in</b> and <b>metric-out</b> statements.</p> |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIPng</a></li> <li>• <a href="#">import on page 3381</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## graceful-restart (Protocols RIPng)

---

|                                 |                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>graceful-restart {<br/>    disable;<br/>    restart-time <i>seconds</i>;<br/>}</pre>                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ripng],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng],<br>[edit protocols ripng],<br>[edit routing-instances <i>routing-instance-name</i> protocols ripng] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for routing instances introduced in Junos OS Release 9.0.                                                                                              |
| <b>Description</b>              | Configure graceful restart for RIPng.                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>disable</b> —Disables graceful restart for RIPng.<br><br>The remaining statement is explained separately.                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Junos OS High Availability Library for Routing Devices</i></li></ul>                                                                                                                                                                             |

## group (Protocols RIPvng)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> group <i>group-name</i> {   export [ <i>policy-names</i> ];   metric-out <i>metric</i>;   neighbor <i>neighbor-name</i> {     import <i>policy-name</i>;     metric-in <i>metric</i>;     receive &lt;none&gt;;     route-timeout <i>seconds</i>;     send &lt;none&gt;;     update-interval <i>seconds</i>;   }   preference <i>number</i>;   route-timeout <i>seconds</i>;   update-interval <i>seconds</i>; } </pre> |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols ripng],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng],</p> <p>[edit protocols ripng],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng]</p>                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Configure a set of RIPvng neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group.</p> <p>Each group must contain at least one neighbor. You should create a group for each export policy that you have.</p>                                                                                                                          |
| <b>Options</b>                  | <p><b><i>group-name</i></b>—Name of a group, up to 16 characters long.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring RIPvng</i></li> </ul>                                                                                                                                                                                                                                                                                                                                        |

## holddown (Protocols RIPng)

---

|                                 |                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>holddown seconds;</code>                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ripng],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng],<br>[edit protocols ripng],<br>[edit routing-instances <i>routing-instance-name</i> protocols ripng] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for routing instances introduced in Junos OS Release 9.0.                                                                                              |
| <b>Description</b>              | Configure how long the expired route is retained in the routing table before being removed.                                                                                                                                                                                                 |
| <b>Options</b>                  | <b>seconds</b> —Estimated time to wait before removing expired routes from the routing table.<br><b>Default:</b> 180 seconds<br><b>Range:</b> 10 through 180 seconds                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring RIPng Timers</i></li></ul>                                                                                                                                                                                                  |

## import (Protocols RIPng)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols ripng],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit protocols ripng],</p> <p>[edit protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more policies to routes being imported into the local routing device from its neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Applying Policies to RIPng Routes Imported from Neighbors</i></li> <li>• <a href="#">export on page 3377</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## metric-in (Protocols RIPng)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric-in <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols ripng],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols ripng group <i>group-name</i> <b>neighbor</b></code><br><code>    <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    ripng],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit protocols ripng],</code><br><code>[edit protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ripng],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> <b>neighbor</b></code><br><code>    <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for routing instances introduced in Junos OS Release 9.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify the metric to add to incoming routes when advertising into RIPng routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIPng routes learned through a specific neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <b><i>metric</i></b> —Metric value.<br><b>Range:</b> 1 through 16<br><b>Default:</b> 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring the Metric Value Added to Imported RIPng Routes</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## metric-out (Protocols RIPng)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric-out <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIPng routes sent from this neighbor.</p> <p>When an export policy is configured, RIPng exports all learned routes to neighbors configured with the <b>neighbor</b> statement.</p> <p>If a route being exported was learned from a member of the same RIPng group, the metric associated with that route (unless modified by an export policy) is the normal RIPng metric. For example, a RIPng route with a metric of 5 learned from a neighbor configured with a <b>metric-in</b> value of 2 is advertised with a combined metric of 7 when advertised to RIPng neighbors in the same group. However, if this route was learned from a RIPng neighbor in a different group or from a different protocol, the route is advertised with the metric value configured for that group with the <b>metric-out</b> statement. The default value for <b>metric-out</b> is 1.</p> <p>To modify the metric for routes advertised outside a group, include the <b>metric-out</b> statement.</p> |
| <b>Options</b>                  | <p><b><i>metric</i></b>—Metric value.</p> <p><b>Range:</b> 1 through 16</p> <p><b>Default:</b> 1</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring the Metric Value Added to Imported RIPng Routes</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## neighbor (Protocols RIPng)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>neighbor <i>neighbor-name</i> {<br/>    import [ <i>policy-names</i> ];<br/>    metric-in <i>metric</i>;<br/>    receive &lt;none&gt;;<br/>    route-timeout <i>seconds</i>;<br/>    send &lt;none&gt;;<br/>    update-interval <i>seconds</i>;<br/>}</pre>                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols ripng <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>  ripng group <i>group-name</i>],<br/>[edit protocols ripng <b>group</b> <i>group-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i>]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure neighbor-specific RIPng parameters, thereby overriding the defaults set for the routing device.                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b><i>neighbor-name</i></b>—Name of an interface over which a routing device communicates to its neighbors.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring RIPng</i></li></ul>                                                                                                                                                                                                                                                                                                                                           |



## preference (Protocols RIPng)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols ripng <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i>],</p> <p>[edit protocols ripng <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i>]</p>                                                                              |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Specify the preference of external routes learned by RIPng as compared to those learned from other routing protocols.</p> <p>By default, Junos OS assigns a preference of 100 to routes that originate from RIPng. When Junos OS determines that a route is to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.</p> <p>To modify the default RIPng preference value, include the <b>preference</b> statement.</p> |
| <b>Options</b>                  | <p><b>preference</b>—Preference value. A lower value indicates a more preferred route.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 100</p>                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring RIPng</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                     |

## receive (Protocols RIPng)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | receive <none>;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ripng],<br>[edit logical-systems <i>logical-system-name</i> protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i> ],<br>[edit protocols ripng],<br>[edit protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ripng],<br>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for routing instances introduced in Junos OS Release 9.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Enable or disable receiving of update messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>none</b> —(Optional) Disable receiving update messages.<br><b>Default:</b> Enabled                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">send on page 3388</a></li><li>• <i>Example: Configuring RIPng</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## ripng

|                                 |                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ripng {...}                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for routing instances introduced in Junos OS Release 9.0.                                                                      |
| <b>Description</b>              | Enable RIPng routing on the routing device.                                                                                                                                                                                                                         |
| <b>Default</b>                  | RIPng is disabled on the routing device.                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring RIPng</i></li> </ul>                                                                                                                                                                               |

## route-timeout (Protocols RIPng)

|                                 |                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | route-timeout <i>seconds</i> ;                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ripng],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng],<br>[edit protocols ripng],<br>[edit routing-instances <i>routing-instance-name</i> protocols ripng] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for routing instances introduced in Junos OS Release 9.0.                                                                                                  |
| <b>Description</b>              | Configure the route timeout interval for RIPng.                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>seconds</b> —Estimated time to wait before making updates to the routing table.<br><b>Range:</b> 30 through 360 seconds<br><b>Default:</b> 180 seconds                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring RIPng Timers</i></li> </ul>                                                                                                                                                                                                |

## send (Protocols RIPng)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | send <none>;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ripng],<br>[edit logical-systems <i>logical-system-name</i> protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances-name</i> protocols ripng],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i> ],<br>[edit protocols ripng],<br>[edit protocols ripng group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ripng],<br>[edit routing-instances <i>routing-instance-name</i> protocols ripng group <i>group-name</i> neighbor <i>neighbor-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for routing instances introduced in Junos OS Release 9.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Enable or disable sending of update messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b>none</b> —(Optional) Disable sending of update messages.<br><b>Default:</b> Enabled                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">receive on page 3386</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## traceoptions (Protocols RIPng)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;<i>flag-modifier</i>&gt; &lt;disable&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols ripng],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng],</p> <p>[edit protocols ripng],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ripng]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for routing instances introduced in Junos OS Release 9.0.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>         | Set RIPng protocol-level tracing options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default</b>             | The default RIPng protocol-level trace options are inherited from the global <b>traceoptions</b> statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIPng tracing output in the file <code>/var/log/ripng-log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>trace-file</i></b> reaches its maximum size, it is renamed <b><i>trace-file.0</i></b>, then <b><i>trace-file.1</i></b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 10 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>RIPng Tracing Options</b></p> <ul style="list-style-type: none"> <li><b>error</b>—RIPng error packets</li> <li><b>expiration</b>—RIPng route expiration processing</li> <li><b>holddown</b>—RIPng hold-down processing</li> <li><b>nsr-synchronization</b>—Nonstop routing synchronization events</li> <li><b>packets</b>—All RIPng packets</li> </ul> |

- **request**—RIPng information packets such as request, poll, and poll entry packets
- **trigger**—RIPng triggered updates
- **update**—RIPng update packets

#### Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

**no-world-readable**—(Optional) Do not allow any user to read the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|

**Related Documentation** • *Example: Tracing RIPng Protocol Traffic*

## update-interval (Protocols RIPng)

|                                 |                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | update-interval <i>seconds</i> ;                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ripng],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ripng],<br>[edit protocols ripng],<br>[edit routing-instances <i>routing-instance-name</i> protocols ripng] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for routing instances introduced in Junos OS Release 9.0.                                                                                                  |
| <b>Description</b>              | Configure the interval at which routes learned by RIPng are sent to neighbors.                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>seconds</b> —Estimated time to wait before making updates to the routing table.<br><b>Range:</b> 10 through 60 seconds<br><b>Default:</b> 30 seconds                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                         |
| <b>Related Documentation</b>    | • <i>Example: Configuring RIP Timers</i>                                                                                                                                                                                                                                                    |





# Administration

- [Routine Monitoring on page 3393](#)
- [Operational Commands: RIP on page 3394](#)
- [Operational Commands: RIPv6 on page 3403](#)

## Routine Monitoring

- [Monitoring RIP Routing Information on page 3393](#)

### Monitoring RIP Routing Information

Purpose



NOTE: This topic applies only to the J-Web Application package.

Use the monitoring functionality to monitor RIP routing on routing devices.

Action

To view RIP routing information in the J-Web interface, select **Monitor > Routing > RIP Information**.

To view RIP routing information in the CLI, enter the following CLI commands:

- `show rip statistics`
- `show rip neighbor`

Meaning

[Table 357 on page 3393](#) summarizes key output fields in the RIP routing display in the J-Web interface.

Table 357: Summary of Key RIP Routing Output Fields

| Field          | Values                            | Additional Information |
|----------------|-----------------------------------|------------------------|
| RIP Statistics |                                   |                        |
| Protocol Name  | The RIP protocol name.            |                        |
| Port number    | The port on which RIP is enabled. |                        |

Table 357: Summary of Key RIP Routing Output Fields (*continued*)

| Field                    | Values                                                                                 | Additional Information                                                                                                |
|--------------------------|----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| Hold down time           | The interval during which routes are neither advertised nor updated.                   |                                                                                                                       |
| Global routes learned    | Number of RIP routes learned on the logical interface.                                 |                                                                                                                       |
| Global routes held down  | Number of RIP routes that are not advertised or updated during the hold-down interval. |                                                                                                                       |
| Global request dropped   | Number of requests dropped.                                                            |                                                                                                                       |
| Global responses dropped | Number of responses dropped.                                                           |                                                                                                                       |
| <b>RIP Neighbors</b>     |                                                                                        |                                                                                                                       |
| Neighbor                 | Name of the RIP neighbor.                                                              | This value is the name of the interface on which RIP is enabled. Click the name to see the details for this neighbor. |
| State                    | State of the RIP connection: <b>Up</b> or <b>Dn</b> (Down).                            |                                                                                                                       |
| Source Address           | Local source address.                                                                  | This value is the configured address of the interface on which RIP is enabled.                                        |
| Destination Address      | Destination address.                                                                   | This value is the configured address of the immediate RIP adjacency.                                                  |
| Send Mode                | The mode of sending RIP messages.                                                      |                                                                                                                       |
| Receive Mode             | The mode in which messages are received.                                               |                                                                                                                       |
| In Metric                | Value of the incoming metric configured for the RIP neighbor.                          |                                                                                                                       |

**Related Documentation** • [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)

## Operational Commands: RIP

- [clear rip general-statistics](#)
- [clear rip statistics](#)
- [show rip general-statistics](#)
- [show rip neighbor](#)
- [show rip statistics](#)

## clear rip general-statistics

---

|                                                   |                                                                                                                                                                                                         |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3395</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3395</a>                                                                                          |
| <b>Syntax</b>                                     | clear rip general-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                    |
| <b>Syntax (EX Series Switches and QFX Series)</b> | clear rip general-statistics                                                                                                                                                                            |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                |
| <b>Description</b>                                | Clear RIP general statistics.                                                                                                                                                                           |
| <b>Options</b>                                    | <b>none</b> —Clear RIP general statistics.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | clear                                                                                                                                                                                                   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <a href="#">show rip general-statistics on page 3397</a></li> </ul>                                                                                            |
| <b>List of Sample Output</b>                      | <a href="#">clear rip general-statistics on page 3395</a>                                                                                                                                               |
| <b>Output Fields</b>                              | When you enter this command, you are provided feedback on the status of your request.                                                                                                                   |

## Sample Output

### clear rip general-statistics

```
user@host> clear rip general-statistics
```

## clear rip statistics

---

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3396</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3396</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                     | <code>clear rip statistics</code><br><code>&lt;instance (all   <i>instance-name</i>)&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;neighbor&gt;</code><br><code>&lt;peer (all   <i>address</i>)&gt;</code>                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <code>clear rip statistics</code><br><code>&lt;instance (all   <i>instance-name</i>)&gt;</code><br><code>&lt;neighbor&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                                | Clear RIP statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                                    | <b>none</b> —Reset RIP counters for all neighbors for all routing instances.<br><br><b>instance (all   <i>instance-name</i>)</b> —(Optional) Clear RIP statistics for all instances or for the specified routing instance only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>neighbor</b> —(Optional) Clear RIP statistics for the specified neighbor only.<br><br><b>peer (all   <i>address</i>)</b> —(Optional) Clear RIP statistics for a single peer or all peers. |
| <b>Required Privilege Level</b>                   | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">show rip statistics on page 3401</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>                      | <a href="#">clear rip statistics on page 3396</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>                              | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## Sample Output

### clear rip statistics

```
user@host> clear rip statistics
```

## show rip general-statistics

|                                                   |                                                                                                                                                                                          |
|---------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3397</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3397</a>                                                                           |
| <b>Syntax</b>                                     | show rip general-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                      |
| <b>Syntax (EX Series Switches and QFX Series)</b> | show rip general-statistics                                                                                                                                                              |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>                                | Display brief RIP statistics.                                                                                                                                                            |
| <b>Options</b>                                    | none—Display brief RIP statistics.<br><br>logical-system (all   <i>logical-system-name</i> )—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                     |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li><a href="#">clear rip general-statistics on page 3395</a></li> </ul>                                                                              |
| <b>List of Sample Output</b>                      | <a href="#">show rip general-statistics on page 3397</a>                                                                                                                                 |
| <b>Output Fields</b>                              | Table 358 on page 3397 lists the output fields for the <b>show rip general-statistics</b> command. Output fields are listed in the approximate order in which they appear.               |

**Table 358: show rip general-statistics Output Fields**

| Field Name  | Field Description                                      |
|-------------|--------------------------------------------------------|
| bad msgs    | Number of invalid messages received.                   |
| no rcv intf | Number of packets received with no matching interface. |
| curr memory | Amount of memory currently used by RIP.                |
| max memory  | Most memory used by RIP.                               |

## Sample Output

### show rip general-statistics

```
user@host> show rip general-statistics
```

```
RIPv2 I/O info:
  bad msgs      :      0
  no recv intf  :      0
  curr memory   :      0
  max memory    :      0
```

## show rip neighbor

---

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3399</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3399</a>                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax</b>                                     | <pre>show rip neighbor &lt;instance (all   <i>instance-name</i>)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show rip neighbor &lt;instance (all   <i>instance-name</i>)&gt; &lt;name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>                        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                                | Display information about RIP neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                    | <p><b>none</b>—Display information about all RIP neighbors for all instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name</b>—(Optional) Display detailed information about only the specified RIP neighbor.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>                      | <a href="#">show rip neighbor on page 3400</a><br><a href="#">show rip neighbor (With Demand Circuits Configured) on page 3400</a>                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>                              | <a href="#">Table 359 on page 3400</a> lists the output fields for the <b>show rip neighbor</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                              |

Table 359: show rip neighbor Output Fields

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Neighbor</b>            | Name of the RIP neighbor.<br><br><b>NOTE:</b> Beginning with Junos OS Release 11.1, when you configure demand circuits, the output displays a demand circuit (DC) flag next to neighbor interfaces configured for demand circuits.<br><br>If you configure demand circuits at the <b>[edit protocols rip group group-name neighbor neighbor-name]</b> hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the <b>[edit protocols rip group group-name]</b> hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits. |
| <b>State</b>               | State of the connection: <b>Up</b> or <b>Dn</b> (Down).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Source Address</b>      | Address of the port on the local router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Destination Address</b> | Address of the port on the remote router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Send Mode</b>           | Send options: <b>broadcast</b> , <b>multicast</b> , <b>none</b> , or <b>version 1</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Receive Mode</b>        | Type of packets to accept: <b>both</b> , <b>none</b> , <b>version 1</b> , or <b>version 2</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>In Met</b>              | Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Sample Output

## show rip neighbor

```

user@host> show rip neighbor
Neighbor      Local  Source  Destination  Send  Receive  In
-----      -
ge-2/3/0.0    Up    192.168.9.105  192.168.9.107  bcast  both      1
at-5/1/1.42    Dn    (null)      (null)      mcast  v2 only    3
at-5/1/0.42    Dn    (null)      (null)      mcast  both       3
at-5/1/0.0     Up    20.0.0.1     224.0.0.9     mcast  both       3
so-0/0/0.0     Up    192.168.9.97  224.0.0.9     mcast  both       3

```

## show rip neighbor (With Demand Circuits Configured)

```

user@host> show rip neighbor
Neighbor      Local  Source  Destination  Send  Receive  In
-----      -
so-0/1/0.0(DC) Up    10.10.10.2   224.0.0.9     mcast  both       1
so-0/2/0.0(DC) Up    13.13.13.2   224.0.0.9     mcast  both       1

```



## show rip statistics

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3401</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3401</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                                     | <pre>show rip statistics &lt;instance (all   <i>instance-name</i>)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>name</i>&gt; &lt;peer (all   <i>address</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show rip statistics &lt;instance (all   <i>instance-name</i>)&gt; &lt;<i>name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>                                | Display RIP statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                    | <p><b>none</b>—Display RIP statistics for all routing instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Display RIP statistics for all instances or for only the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>name</i></b>—(Optional) Display detailed information about only the specified RIP neighbor.</p> <p><b>peer (all   <i>address</i>)</b>—(Optional) Display RIP statistics for a single peer or all peers.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <a href="#">clear rip statistics on page 3396</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>                      | <a href="#">show rip statistics on page 3402</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>                              | <a href="#">Table 360 on page 3402</a> lists the output fields for the <b>show rip statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 360: show rip statistics Output Fields

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RIP info</b>          | <p>Information about RIP on the specified interface:</p> <ul style="list-style-type: none"> <li>• <b>port</b>—UDP port number used for RIP.</li> <li>• <b>update interval</b>—Interval between routing table updates, in seconds.</li> <li>• <b>holddown</b>—Hold-down interval, in seconds.</li> <li>• <b>timeout</b>—Timeout interval, in seconds.</li> <li>• <b>restart in progress</b>—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart.</li> <li>• <b>restart time</b>—Estimated time for the graceful restart to finish, in seconds.</li> <li>• <b>restart will complete in</b>—Remaining time for the graceful restart to finish, in seconds.</li> <li>• <b>rts learned</b>—Number of routes learned through RIP.</li> <li>• <b>rts held down</b>—Number of routes held down by RIP.</li> <li>• <b>rqsts dropped</b>—Number of received request packets that were dropped.</li> <li>• <b>resps dropped</b>—Number of received response packets that were dropped.</li> </ul>                                                                                                                                                   |
| <b>logical-interface</b> | <p>Name of the logical interface and its statistics:</p> <ul style="list-style-type: none"> <li>• <b>routes learned</b>—Number of routes learned on the logical interface.</li> <li>• <b>routes advertised</b>—Number of routes advertised by the logical interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Counter</b>           | <p>List of counter types:</p> <ul style="list-style-type: none"> <li>• <b>Updates Sent</b>—Number of update messages sent.</li> <li>• <b>Triggered Updates Sent</b>—Number of triggered update messages sent.</li> <li>• <b>Responses Sent</b>—Number of response messages sent.</li> <li>• <b>Bad Messages</b>—Number of invalid messages received.</li> <li>• <b>RIPv1 Updates Received</b>—Number of RIPv1 update messages received.</li> <li>• <b>RIPv1 Bad Route Entries</b>—Number of RIPv1 invalid route entry messages received.</li> <li>• <b>RIPv1 Updates Ignored</b>—Number of RIPv1 update messages ignored.</li> <li>• <b>RIPv2 Updates Received</b>—Number of RIPv2 update messages received.</li> <li>• <b>RIPv2 Bad Route Entries</b>—Number of RIPv2 invalid route entry messages received.</li> <li>• <b>RIPv2 Updates Ignored</b>—Number of RIPv2 update messages ignored.</li> <li>• <b>Authentication Failures</b>—Number of received update messages that failed authentication.</li> <li>• <b>RIP Requests Received</b>—Number of RIP request messages received.</li> <li>• <b>RIP Requests Ignored</b>—Number of RIP request messages ignored.</li> </ul> |
| <b>Total</b>             | Total number of packets for the selected counter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Last 5 min</b>        | Number of packets for the selected counter in the most recent 5-minute period.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Last minute</b>       | Number of packets for the selected counter in the most recent 1-minute period.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Sample Output

### show rip statistics

```
user@host> show rip statistics so-0/0/0.0
```

```

RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s
restart in progress: restart time 60s; restart will complete in 55s
    rts learned  rts held down  rqsts dropped  resps dropped
              0             0             0             0
so-0/0/0.0: 0 routes learned; 501 routes advertised
Counter      Total      Last 5 min  Last minute
-----
Updates Sent          0          0          0
Triggered Updates Sent  0          0          0
Responses Sent        0          0          0
Bad Messages          0          0          0
RIPv1 Updates Received  0          0          0
RIPv1 Bad Route Entries  0          0          0
RIPv1 Updates Ignored   0          0          0
RIPv2 Updates Received  0          0          0
RIPv2 Bad Route Entries  0          0          0
RIPv2 Updates Ignored   0          0          0
Authentication Failures  0          0          0
RIP Requests Received   0          0          0
RIP Requests Ignored    0          0          0

```

## Operational Commands: RIPvng

- `clear ripng general-statistics`
- `clear ripng statistics`
- `show ripng general-statistics`
- `show ripng neighbor`
- `show ripng statistics`

## clear ripng general-statistics

---

|                                    |                                                                                                                                                                                                           |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3404</a><br><a href="#">Syntax (EX Series Switches) on page 3404</a>                                                                                                           |
| <b>Syntax</b>                      | clear ripng general-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                    |
| <b>Syntax (EX Series Switches)</b> | clear ripng general-statistics                                                                                                                                                                            |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                     |
| <b>Description</b>                 | Clear RIP next generation (RIPng) general statistics.                                                                                                                                                     |
| <b>Options</b>                     | <b>none</b> —Clear RIPng general statistics.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | clear                                                                                                                                                                                                     |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"><li>• <a href="#">show ripng general-statistics on page 3406</a></li></ul>                                                                                              |
| <b>List of Sample Output</b>       | <a href="#">clear ripng general-statistics on page 3404</a>                                                                                                                                               |
| <b>Output Fields</b>               | When you enter this command, you are provided feedback on the status of your request.                                                                                                                     |

## Sample Output

### clear ripng general-statistics

```
user@host> clear ripng general-statistics
```

## clear ripng statistics

---

|                                  |                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>            | <a href="#">Syntax on page 3405</a><br><a href="#">Syntax (EX Series Switch) on page 3405</a>                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                    | clear ripng statistics<br>< <i>instance</i>   <i>name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                |
| <b>Syntax (EX Series Switch)</b> | clear ripng statistics<br>< <i>instance</i>   <i>name</i> >                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>       | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                              |
| <b>Description</b>               | Clear RIP next-generation (RIPng) statistics.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                   | <p><b>none</b>—Reset RIPng counters for all neighbors for all routing instances.</p> <p><b>instance</b>—(Optional) Reset RIPng counters for the specified instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name</b>—(Optional) Reset RIPng counters for the specified neighbor.</p> |
| <b>Required Privilege Level</b>  | clear                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>     | <ul style="list-style-type: none"> <li>• <a href="#">show ripng statistics on page 3410</a></li> </ul>                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>     | <a href="#">clear ripng statistics on page 3405</a>                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>             | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                              |

## Sample Output

### clear ripng statistics

```
user@host> clear ripng statistics
```

## show ripng general-statistics

|                                  |                                                                                                                                                                                              |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>            | <a href="#">Syntax on page 3406</a><br><a href="#">Syntax (EX Series Switch) on page 3406</a>                                                                                                |
| <b>Syntax</b>                    | show ripng general-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                        |
| <b>Syntax (EX Series Switch)</b> | show ripng general-statistics                                                                                                                                                                |
| <b>Release Information</b>       | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                        |
| <b>Description</b>               | Display general RIP next-generation (RIPng) statistics.                                                                                                                                      |
| <b>Options</b>                   | none—Display general RIPng statistics.<br><br>logical-system (all   <i>logical-system-name</i> )—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>  | view                                                                                                                                                                                         |
| <b>Related Documentation</b>     | <ul style="list-style-type: none"> <li><a href="#">clear ripng general-statistics on page 3404</a></li> </ul>                                                                                |
| <b>List of Sample Output</b>     | <a href="#">show ripng general-statistics on page 3406</a>                                                                                                                                   |
| <b>Output Fields</b>             | Table 361 on page 3406 lists the output fields for the <b>show ripng general-statistics</b> command. Output fields are listed in the approximate order in which they appear.                 |

**Table 361: show ripng general-statistics Output Fields**

| Field Name  | Field Description                                      |
|-------------|--------------------------------------------------------|
| bad msgs    | Number of invalid messages received.                   |
| no rcv intf | Number of packets received with no matching interface. |
| curr memory | Amount of memory currently used by RIPng.              |
| max memory  | Most memory used by RIPng.                             |

## Sample Output

### show ripng general-statistics

```

user@host> show ripng general-statistics
RIPng I/O info:
  bad msgs      :      0
  no rcv intf   :      0

```

```
curr memory : 0
max memory  : 0
```

## show ripng neighbor

|                                  |                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>            | <a href="#">Syntax on page 3408</a><br><a href="#">Syntax (EX Series Switch) on page 3408</a>                                                                                                                                                                                                                                    |
| <b>Syntax</b>                    | show ripng neighbor<br><logical-system (all   <i>logical-system-name</i> )><br>< <i>name</i> >                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switch)</b> | show ripng neighbor<br>< <i>name</i> >                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>       | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                            |
| <b>Description</b>               | Display information about RIP next-generation (RIPng) neighbors.                                                                                                                                                                                                                                                                 |
| <b>Options</b>                   | <p><b>none</b>—Display information about all RIPng neighbors.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>name</i></b>—(Optional) Display detailed information about a specific RIPng neighbor.</p> |
| <b>Required Privilege Level</b>  | view                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>     | <a href="#">show ripng neighbor on page 3409</a>                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>             | <a href="#">Table 362 on page 3408</a> lists the output fields for the <b>show ripng neighbor</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                               |

**Table 362: show ripng neighbor Output Fields**

| Field Name                 | Field Description                                                                                          |
|----------------------------|------------------------------------------------------------------------------------------------------------|
| <b>Neighbor</b>            | Name of RIPng neighbor.                                                                                    |
| <b>State</b>               | State of the connection: <b>Up</b> or <b>Dn</b> (Down).                                                    |
| <b>Source Address</b>      | Source address.                                                                                            |
| <b>Destination Address</b> | Destination address.                                                                                       |
| <b>Send</b>                | Send options: <b>broadcast</b> , <b>multicast</b> , <b>none</b> , <b>version 1</b> , or <b>yes</b> .       |
| <b>Recv</b>                | Type of packets to accept: <b>both</b> , <b>none</b> , <b>version 1</b> , or <b>yes</b> .                  |
| <b>In Met</b>              | Metric added to incoming routes when advertising into RIPng routes that were learned from other protocols. |



## Sample Output

show ripng neighbor

```
user@host> show ripng neighbor
```

| Neighbor   | State | Source<br>Address        | Dest<br>Address | Send  | Recv  | In<br>Met |
|------------|-------|--------------------------|-----------------|-------|-------|-----------|
| -----      | ----- | -----                    | -----           | ----- | ----- | -----     |
| fe-0/0/2.0 | Up    | fe80::290:69ff:fe68:b002 | ff02::9         | yes   | yes   | 1         |

## show ripng statistics

|                                  |                                                                                                                                                                                                                                                                                                                               |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>            | <a href="#">Syntax on page 3410</a><br><a href="#">Syntax (EX Series Switch) on page 3410</a>                                                                                                                                                                                                                                 |
| <b>Syntax</b>                    | show ripng statistics<br><logical-system (all   <i>logical-system-name</i> )><br>< <i>name</i> >                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switch)</b> | show ripng statistics<br>< <i>name</i> >                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>       | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                         |
| <b>Description</b>               | Display RIP next generation (RIPng) statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.                                                                                                                                              |
| <b>Options</b>                   | <p><b>none</b>—Display RIPng statistics for all neighbors.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>name</i></b>—(Optional) Display detailed information about a specific RIPng neighbor.</p> |
| <b>Required Privilege Level</b>  | view                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>     | <ul style="list-style-type: none"> <li>• <a href="#">clear ripng statistics on page 3405</a></li> </ul>                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>     | <a href="#">show ripng statistics on page 3411</a>                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>             | <a href="#">Table 363 on page 3410</a> lists the output fields for the <b>show ripng statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                          |

**Table 363: show ripng statistics Output Fields**

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RIPng Info</b> | Information about RIPng on the specified interface: <ul style="list-style-type: none"> <li>• <b>port</b>—UDP port number used for RIPng.</li> <li>• <b>holddown</b>—Hold-down interval, in seconds.</li> <li>• <b>rts learned</b>—Number of routes learned through RIPng.</li> <li>• <b>rts held down</b>—Number of routes held down by RIPng.</li> <li>• <b>rqsts dropped</b>—Number of received request packets that were dropped.</li> <li>• <b>resps dropped</b>—Number of received response packets that were dropped.</li> <li>• <b>restart</b>—Graceful restart status. Displayed when RIPng is or has been in the process of graceful restart.</li> </ul> |

Table 363: show ripng statistics Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>logical-interface</i> | Name of the logical interface and its statistics: <ul style="list-style-type: none"> <li><b>routes learned</b>—Number of routes learned on the logical interface.</li> <li><b>routes advertised</b>—Number of routes advertised by the logical interface.</li> <li><b>timeout</b>—Timeout interval, in seconds.</li> <li><b>update interval</b>—Interval between routing table updates, in seconds.</li> </ul>                                                                                                                                                                                                                                                                                                                                                    |
| <b>Counter</b>           | List of counter types: <ul style="list-style-type: none"> <li><b>Updates Sent</b>—Number of update messages sent.</li> <li><b>Triggered Updates Sent</b>—Number of triggered update messages sent.</li> <li><b>Responses Sent</b>—Number of response messages sent.</li> <li><b>Bad Messages</b>—Number of invalid messages received.</li> <li><b>Updates Received</b>—Number of RIPng update messages received.</li> <li><b>Bad Route Entries</b>—Number of RIPng invalid route entry messages received.</li> <li><b>Updates Ignored</b>—Number of RIPng update messages ignored.</li> <li><b>RIPng Requests Received</b>—Number of RIPng request messages received.</li> <li><b>RIPng Requests Ignored</b>—Number of RIPng request messages ignored.</li> </ul> |
| <b>Total</b>             | Total number of packets for the selected counter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Last 5 min</b>        | Number of packets for the selected counter in the most recent 5-minute period.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Last minute</b>       | Number of packets for the selected counter in the most recent 1-minute period.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Sample Output

### show ripng statistics

```

user@host> show ripng statistics
RIPng info: port 521; holddown 120s;
      rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0

so-0/1/3.0: 0 routes learned; 1 routes advertised; timeout 180s; update interval
20s
Counter              Total    Last 5 min  Last minute
-----
Updates Sent          934         16         4
Triggered Updates Sent    1         0         0
Responses Sent         0         0         0
Bad Messages          0         0         0
Updates Received       0         0         0
Bad Route Entries      0         0         0
Updates Ignored        0         0         0
RIPng Requests Received 0         0         0
RIPng Requests Ignored  0         0         0

```



## PART 19

# Routing Options

- [Overview on page 3415](#)
- [Configuration on page 3421](#)
- [Administration on page 3537](#)



## CHAPTER 60

# Overview

- [Layer 3 Protocols on page 3415](#)
- [Routing Options Overview on page 3418](#)

## Layer 3 Protocols

- [Layer 3 Protocols Supported on EX Series Switches on page 3415](#)
- [Layer 3 Protocols Not Supported on EX Series Switches on page 3416](#)

## Layer 3 Protocols Supported on EX Series Switches

EX Series switches support the Junos OS Layer 3 features and configuration statements listed in [Table 314 on page 2939](#):

**Table 364: Supported Junos OS Layer 3 Protocol Statements and Features**

| Protocol           | Notes                                                                                                         | For More Information                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| BGP                | Fully supported.                                                                                              | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| BFD                | Fully supported.                                                                                              | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| ICMP               | Fully supported.                                                                                              | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| IGMPv1, v2, and v3 | Fully supported.                                                                                              | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |
| IS-IS              | Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on EX Series Switches" on page 2940. | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |
| MLD                | Fully supported (MLD versions 1 and 2).                                                                       | <a href="#">Junos OS Multicast Protocols Configuration Guide</a> |
| MPLS               | Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on EX Series Switches" on page 2940. | <a href="#">Junos OS MPLS Applications Configuration Guide</a>   |
| OSPFv1, v2 and v3  | Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on EX Series Switches" on page 2940. | <a href="#">Junos OS Routing Protocols Configuration Guide</a>   |

Table 364: Supported Junos OS Layer 3 Protocol Statements and Features (*continued*)

| Protocol | Notes                                                                                                | For More Information                                                                                                     |
|----------|------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| PIM      | Fully supported.                                                                                     | <a href="#">Junos OS Multicast Protocols Configuration Guide</a>                                                         |
| PPM      | Supported. See <i>EX Series Switch Software Features Overview</i> for specific platform information. | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| RIP      | Fully supported.                                                                                     | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| RIPng    | Fully supported.                                                                                     | <a href="#">Junos OS Routing Protocols Configuration Guide</a>                                                           |
| SNMP     | Fully supported.                                                                                     | <a href="#">Junos OS Network Management Configuration Guide</a>                                                          |
| VRRP     | Fully supported.                                                                                     | See "Understanding VRRP on EX Series Switches" on page 2501. See also <a href="#">Junos OS High Availability Guide</a> . |

- Related Documentation**
- [Layer 3 Protocols Not Supported on EX Series Switches](#) on page 2940
  - [EX Series Switch Software Features Overview](#)

### Layer 3 Protocols Not Supported on EX Series Switches

EX Series switches do not support the Junos OS Layer 3 protocols and features listed in [Table 315 on page 2940](#):

Table 365: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported

| Feature                                                                                                         | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DVMRP                                                                                                           | <ul style="list-style-type: none"> <li>• <b>dvmrp</b> and subordinate statements</li> </ul>                                                                                                                                                                                                          |
| Flow aggregation (cflowd)                                                                                       | <ul style="list-style-type: none"> <li>• <b>cflow</b> and subordinate statements</li> </ul>                                                                                                                                                                                                          |
| IPsec                                                                                                           | <ul style="list-style-type: none"> <li>• <b>[edit services]</b> statements related to IPsec</li> </ul>                                                                                                                                                                                               |
| IS-IS: <ul style="list-style-type: none"> <li>• ES-IS</li> <li>• IPv6 in multicast routing protocols</li> </ul> | <ul style="list-style-type: none"> <li>• <b>clns-routing</b> statement</li> <li>• <b>ipv6-multicast</b> statement</li> <li>• <b>lsp-interval</b> statement</li> <li>• <b>label-switched-path</b> statement</li> <li>• <b>lsp-lifetime</b> statement</li> <li>• <b>te-metric</b> statement</li> </ul> |
| Logical routers                                                                                                 | <ul style="list-style-type: none"> <li>• <b>logical-routers</b> and subordinate statements</li> </ul>                                                                                                                                                                                                |



**Table 365: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)**

| Feature                                                                                                                                                                                                                                                                                                                                                                                                                                      | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MPLS: <ul style="list-style-type: none"> <li>Fast Reroute (FRR)</li> <li>Label Distribution Protocol (LDP) (except on EX8200 switches)</li> <li>Layer 3 VPNs (except on EX8200 switches)</li> <li>Multiprotocol BGP (MP-BGP) for VPN-IPv4 family</li> <li>Pseudowire emulation (PWE3)</li> <li>Routing policy statements related to Layer 3 VPNs and MPLS (except on EX8200 switches)</li> <li>Virtual Private LAN Service (VPLS)</li> </ul> | <ul style="list-style-type: none"> <li><b>ldp</b> and all subordinate statements (except on EX8200 switches)</li> </ul>                                                                                                                                                                                                                                                                                           |
| Network Address Translation (NAT)                                                                                                                                                                                                                                                                                                                                                                                                            | <ul style="list-style-type: none"> <li><b>nat</b> and subordinate statements</li> <li>Policy statements related to NAT</li> </ul>                                                                                                                                                                                                                                                                                 |
| OSPF                                                                                                                                                                                                                                                                                                                                                                                                                                         | <ul style="list-style-type: none"> <li><b>demand-circuit</b> statement</li> <li><b>label-switched-path</b> and subordinate statements</li> <li><b>neighbor</b> statement within an OSPF area</li> <li><b>peer-interface</b> and subordinate statements within an OSPF area</li> <li><b>sham-link</b> statement</li> <li><b>te-metric</b> statement</li> </ul>                                                     |
| PPM                                                                                                                                                                                                                                                                                                                                                                                                                                          | <ul style="list-style-type: none"> <li>Not supported on EX2200 and EX3300 switches</li> </ul>                                                                                                                                                                                                                                                                                                                     |
| Routing instances: <ul style="list-style-type: none"> <li>Routing instance forwarding</li> </ul>                                                                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li><b>l2vpn</b> and subordinate statements (except on EX4500, EX4550, and EX8200 switches)</li> <li><b>ldp</b> and subordinate statements (except on EX8200 switches)</li> <li><b>vpls</b> and subordinate statements</li> </ul>                                                                                                                                              |
| Routed VLAN interfaces (RVIs)                                                                                                                                                                                                                                                                                                                                                                                                                | <ul style="list-style-type: none"> <li><b>family mpls</b> statement</li> </ul>                                                                                                                                                                                                                                                                                                                                    |
| SAP and SDP                                                                                                                                                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li><b>sap</b> and all subordinate statements</li> </ul>                                                                                                                                                                                                                                                                                                                       |
| General routing options in the <b>routing-options</b> hierarchy: <ul style="list-style-type: none"> <li>MPLS and label-switched-paths</li> </ul>                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li><b>auto-export</b> and subordinate statements</li> <li><b>dynamic-tunnels</b> and subordinate statements</li> <li><b>lsp-next-hop</b> and subordinate statements</li> <li><b>multicast</b> and subordinate statements</li> <li><b>p2mp-lsp-next-hop</b> and subordinate statements</li> <li><b>route-distinguisher-id</b> statement (except on EX8200 switches)</li> </ul> |

Table 365: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported (*continued*)

| Feature                                                                    | Configuration Statements Not Supported on EX Series Switches                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Traffic sampling and forwarding in the <b>forwarding-options</b> hierarchy | <ul style="list-style-type: none"> <li>• <b>accounting</b> and subordinate statements</li> <li>• <b>family mpls</b> and <b>family multiservice</b> under <b>hash-key</b> hierarchy</li> <li>• Under <b>monitoring group-name</b> family <b>inet output</b> hierarchy: <ul style="list-style-type: none"> <li>• <b>cflowd</b> statement</li> <li>• <b>export-format-cflowd-version-5</b> statement</li> <li>• <b>flow-active-timeout</b> statement</li> <li>• <b>flow-export-destination</b> statement</li> <li>• <b>flow-inactive-timeout</b> statement</li> <li>• <b>interface</b> statement</li> </ul> </li> <li>• <b>port-mirroring</b> statement (On EX Series switches, port mirroring is implemented using the <b>analyzer</b> statement.)</li> <li>• <b>sampling</b> and subordinate statements</li> </ul> |

- Related Documentation**
- [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)
  - [EX Series Switch Software Features Overview](#)

## Routing Options Overview

- [Understanding Distributed Periodic Packet Management on EX Series Switches on page 3418](#)

### Understanding Distributed Periodic Packet Management on EX Series Switches

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks for particular processes so that other processes on the Juniper Networks EX Series Ethernet Switch can more optimally direct their resources. PPM is responsible for the periodic transmission of packets on behalf of its various client processes, which include the processes that control the Link Aggregation Control Protocol (LACP) and Bidirectional Forwarding Detection (BFD) protocols, and also for receiving packets on behalf of these client processes. PPM also gathers some statistics and sends process-specific packets. PPM cannot be disabled and is always running on any operational switch.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and either the access interfaces (on EX3200, EX4200, and EX4500 switches) or the line cards (on EX6200 and EX8200 switches) for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for LACP packets only.



**BEST PRACTICE:** We recommend that, generally, you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

**Related  
Documentation**

- [Configuring Distributed Periodic Packet Management on an EX Series Switch \(CLI Procedure\) on page 3424](#)



## CHAPTER 61

# Configuration

- [Configuration Tasks on page 3421](#)
- [Configuration Statements on page 3425](#)

### Configuration Tasks

---

- [Configuring Static Routing \(CLI Procedure\) on page 3422](#)
- [Configuring Static Routing \(J-Web Procedure\) on page 3422](#)
- [Configuring Distributed Periodic Packet Management on an EX Series Switch \(CLI Procedure\) on page 3424](#)

## Configuring Static Routing (CLI Procedure)

Static routes are routes that are manually configured and entered into the routing table. Dynamic routes, in contrast, are learned by the EX Series switch and added to the routing table using a protocol such as OSPF or RIP.

The switch uses static routes:

- When the switch does not have a route to a destination that has a better (lower) *preference* value. The preference is an arbitrary value in the range from 0 through 255 that the software uses to rank routes received from different protocols, interfaces, or remote systems. The routing protocol process generally determines the active route by selecting the route with the lowest preference value. In the given range, **0** is the lowest and **255** is the highest.
- When the switch cannot determine the route to a destination.
- When the switch is forwarding unroutable packets.

To configure basic static route options using the CLI:

- To configure the switch's default gateway:

```
[edit]
user@switch# set routing-options static route 0.0.0.0/0 next-hop 10.0.1.1
```

- To configure a static route and specify the next address to be used when routing traffic to the static route:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 next-hop 10.0.0.2.1
```

- To always keep the static route in the forwarding table:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 retain
```

- To prevent the static route from being readvertised:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 no-readvertise
```

- To remove inactive routes from the forwarding table:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 active
```

### Related Documentation

- [Configuring Static Routing \(J-Web Procedure\) on page 3422](#)
- [Monitoring Routing Information on page 3537](#)

## Configuring Static Routing (J-Web Procedure)

You can use the J-Web interface to configure static routes for EX Series switches.

To configure static routes:

1. Select **Configure > Routing > Static Routing**. The Static Routing page displays details of the configured routes.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—To configure a route. Enter information into the routing page as described in [Table 366 on page 3423](#).
- **Edit**—To modify an existing route. Enter information into the routing page as described in [Table 366 on page 3423](#).
- **Delete**—To delete an existing route.

**Table 366: Static Routing Configuration Summary**

| Field                | Function                                                                                                                      | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default Route</b> |                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Default Route        | <p>Specifies the default gateway for the switch.</p> <p><b>NOTE:</b> IPv6 is not supported on EX2200 and EX4500 switches.</p> | <p>To specify an IPv4 address:</p> <ol style="list-style-type: none"> <li>1. Select <b>IPv4</b>.</li> <li>2. Type an IP address—for example, <b>10.10.10.10</b>.</li> <li>3. Enter the subnet mask or address prefix. For example, 24 bits represents <b>255.255.255.0</b>.</li> </ol> <p>To specify an IPv6 address:</p> <ol style="list-style-type: none"> <li>1. Select <b>IPv6</b>.</li> <li>2. Type an IP address—for example, <b>2001:ab8:85a3::8a2e:370:7334</b>.</li> <li>3. Enter the subnet mask or address prefix.</li> </ol> |
| <b>Static Routes</b> |                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

Table 366: Static Routing Configuration Summary (*continued*)

| Field   | Function                                                                                         | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------|--------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nexthop | Specifies the next-hop address or addresses to be used when routing traffic to the static route. | <p>To add an address:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b>.</li> <li>2. In the IP address dialog, enter the IP address.</li> </ol> <p><b>NOTE:</b> If a route has multiple next-hop addresses, traffic is routed across each address in round-robin fashion.</p> <ol style="list-style-type: none"> <li>3. Click <b>OK</b>.</li> </ol> <p>To delete a next-hop address, select it from the list and click <b>Delete</b>.</p> |

**Related Documentation**

- [Configuring Static Routing \(CLI Procedure\) on page 3422](#)
- [Monitoring Routing Information on page 3537](#)
- [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)

## Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks so that other processes on the EX Series switch can more optimally direct their resources.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and either the access interfaces (on EX3200, EX4200, and EX4500 switches) or the line cards (on EX6200 and EX8200 switches) for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for LACP packets only.



**BEST PRACTICE:** We recommend that, generally, you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

This topic describes:

- [Disabling or Enabling Distributed Periodic Packet Management Globally on page 3425](#)
- [Disabling or Enabling Distributed Periodic Packet Management for LACP Packets on page 3425](#)



### Disabling or Enabling Distributed Periodic Packet Management Globally

Distributed PPM is enabled by default. Disable distributed PPM if you need to move all PPM processing to the Routing Engine. Enable distributed PPM if it was previously disabled and you need to run distributed PPM.

To disable distributed PPM:

```
[edit routing-options]
user@switch# set ppm no-delegate-processing
```

To enable distributed PPM if it was previously disabled:

```
[edit routing-options]
user@switch# delete ppm no-delegate-processing
```

### Disabling or Enabling Distributed Periodic Packet Management for LACP Packets

Distributed PPM is enabled by default. Disable distributed PPM for only LACP packets if you need to move all PPM processing for LACP packets to the Routing Engine.

To disable distributed PPM for LACP packets:

```
[edit protocols]
user@switch# set lacp ppm centralized
```

To enable distributed PPM for LACP packets if it was previously disabled:

```
[edit protocols]
user@switch# delete lacp ppm centralized
```

#### Related Documentation

- [Understanding Distributed Periodic Packet Management on EX Series Switches on page 3418](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2582](#)

## Configuration Statements

- [active on page 3429](#)
- [aggregate \(Routing\) on page 3430](#)
- [as-path \(Routing Options\) on page 3432](#)
- [asm-override-ssm on page 3434](#)
- [authentication-key-chains on page 3435](#)
- [autonomous-system on page 3436](#)
- [backup-pe-group on page 3438](#)
- [backups on page 3439](#)
- [bandwidth \(Multicast Flow Map\) on page 3440](#)
- [bfd-liveness-detection \(Routing Options Static Route\) on page 3441](#)
- [bgp-orf-cisco-mode on page 3445](#)
- [bmp on page 3447](#)
- [brief on page 3449](#)

- [centralized](#) on page 3450
- [community \(Routing Options\)](#) on page 3451
- [confederation](#) on page 3453
- [disable \(Routing Options\)](#) on page 3454
- [discard](#) on page 3455
- [export \(Routing Options\)](#) on page 3456
- [export-rib](#) on page 3457
- [fate-sharing](#) on page 3459
- [flow](#) on page 3460
- [flow-map](#) on page 3461
- [forwarding-cache \(Flow Maps\)](#) on page 3462
- [forwarding-cache \(Multicast\)](#) on page 3463
- [forwarding-table](#) on page 3464
- [generate](#) on page 3465
- [import \(Routing Options\)](#) on page 3466
- [import-policy](#) on page 3467
- [import-rib](#) on page 3468
- [indirect-next-hop](#) on page 3469
- [install \(Routing Options\)](#) on page 3470
- [instance-export](#) on page 3471
- [instance-import](#) on page 3471
- [interface \(Multicast Static Routes\)](#) on page 3472
- [interface \(Routing Options\)](#) on page 3473
- [interface-routes](#) on page 3474
- [local-address \(Routing Options\)](#) on page 3475
- [martians](#) on page 3476
- [maximum-bandwidth \(Routing Options\)](#) on page 3477
- [maximum-paths](#) on page 3478
- [maximum-prefixes](#) on page 3480
- [med-igp-update-interval](#) on page 3481
- [metric \(Aggregate, Generated, or Static Route\)](#) on page 3482
- [multicast \(Routing Options\)](#) on page 3483
- [no-qos-adjust](#) on page 3484
- [options \(Routing Options\)](#) on page 3485
- [pim-to-igmp-proxy](#) on page 3486
- [pim-to-mld-proxy](#) on page 3487
- [policy \(Aggregate and Generated Routes\)](#) on page 3488

- [policy \(Flow Maps\) on page 3489](#)
- [policy \(SSM Maps\) on page 3490](#)
- [ppm on page 3491](#)
- [ppm \(Ethernet Switching\) on page 3492](#)
- [preference \(Routing Options\) on page 3493](#)
- [prefix on page 3494](#)
- [qualified-next-hop \(Static Routes\) on page 3495](#)
- [readvertise on page 3497](#)
- [redundant-sources on page 3498](#)
- [resolution on page 3499](#)
- [resolution-ribs on page 3500](#)
- [resolve on page 3501](#)
- [restart-duration \(Routing Options\) on page 3502](#)
- [retain on page 3503](#)
- [reverse-oif-mapping on page 3504](#)
- [rpf-check-policy \(Routing Options RPF\) on page 3505](#)
- [rib \(General\) on page 3506](#)
- [rib \(Route Resolution\) on page 3508](#)
- [rib-group \(Routing Options\) on page 3509](#)
- [rib-groups on page 3510](#)
- [route-distinguisher-id on page 3512](#)
- [route-record on page 3513](#)
- [router-id on page 3514](#)
- [routing-options on page 3515](#)
- [scope on page 3516](#)
- [scope-policy on page 3517](#)
- [source \(Source-Specific Multicast\) on page 3518](#)
- [source-routing on page 3519](#)
- [ssm-groups on page 3520](#)
- [ssm-map \(Routing Options Multicast\) on page 3521](#)
- [static \(Routing Options\) on page 3522](#)
- [subscriber-leave-timer on page 3528](#)
- [tag \(Routing Options\) on page 3529](#)
- [threshold \(Multicast Forwarding Cache\) on page 3530](#)
- [timeout \(Flow Maps\) on page 3531](#)
- [timeout \(Multicast\) on page 3532](#)

- [traceoptions \(Routing Options\) on page 3533](#)
- [upstream-interface on page 3536](#)

## active

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (active   passive);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Determine whether static, aggregate, or generated routes are removed from the routing and forwarding tables when they become inactive. Static routes are only removed from the routing table if the next hop becomes unreachable. This can occur if the local or neighbor interface goes down. Routes that have been configured to remain continually installed in the routing and forwarding tables are marked with <b>reject</b> next hops when they are inactive.</p> <ul style="list-style-type: none"> <li>• <b>active</b>—Remove a route from the routing and forwarding tables when it becomes inactive.</li> <li>• <b>passive</b>—Have a route remain continually installed in the routing and forwarding tables even when it becomes inactive.</li> </ul> <p>Include the <b>active</b> statement when configuring an individual route in the <b>route</b> portion of the <b>static</b> statement to override a <b>passive</b> option specified in the <b>defaults</b> portion of the statement.</p>                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>                  | <b>active</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## aggregate (Routing)

---

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax              | <pre>aggregate {   defaults {     ... aggregate-options ...   }   route destination-prefix {     policy policy-name;     ... aggregate-options ...   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Hierarchy Level     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>rib</b> <i>routing-table-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>rib</b> <i>routing-table-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>rib</b> <i>routing-table-name</i>],</p> <p>[edit routing-options],</p> <p>[edit routing-options <b>rib</b> <i>routing-table-name</i>]</p>                                                                                                                                                                                                                                                                                                                                                                           |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Description         | Configure aggregate routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Options             | <p><b>aggregate-options</b>—Additional information about aggregate routes that is included with the route when it is installed in the routing table. Specify zero or more of the following options in <b>aggregate-options</b>. Each option is explained separately.</p> <ul style="list-style-type: none"><li>• (<b>active</b>   <b>passive</b>);</li><li>• <b>as-path</b> &lt;<i>as-path</i>&gt; &lt;origin (egp   igp   incomplete)&gt; &lt;atomic-aggregate&gt; &lt;aggregator <i>as-number</i> <i>ip-address</i>&gt;;</li><li>• (<b>brief</b>   <b>full</b>);</li><li>• <b>community</b> [ <i>community-ids</i> ];</li><li>• <b>discard</b>;</li><li>• (<b>metric</b>   <b>metric2</b>   <b>metric3</b>   <b>metric4</b>) <i>value</i> &lt;type <i>type</i>&gt;;</li><li>• (<b>preference</b>   <b>preference2</b>   <b>color</b>   <b>color2</b>) <i>preference</i> &lt;type <i>type</i>&gt;;</li><li>• <b>tag</b> <i>metric type number</i>;</li></ul> <p><b>defaults</b>—Specify global aggregate route options. These options only set default attributes inherited by all newly created aggregate routes. These are treated as global defaults</p> |

and apply to all the aggregate routes you configure in the **aggregate** statement. This part of the **aggregate** statement is optional.

**route *destination-prefix***—Configure a nondefault aggregate route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

The **policy** statement is explained separately.

|                                 |                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.                                                           |
|                                 | routing-control—To add this statement to the configuration.                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li></ul> |

## as-path (Routing Options)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>as-path &lt;as-path&gt; &lt;aggregator as-number ip-address&gt; &lt;atomic-aggregate&gt; &lt;origin (egp   igp   incomplete)&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>         | <p>Associate BGP autonomous system (AS) path information with a static, aggregate, or generated route.</p> <p>In Junos OS Release 9.1 and later, the numeric range for the AS number is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i>&lt;16-bit high-order value in decimal&gt;.&lt;16-bit low-order value in decimal&gt;</i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value in the range from 0.0 through 65535.65535 in AS-dot notation format.</p>                                                                                                                  |
| <b>Default</b>             | No AS path information is associated with static routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>             | <p><b>aggregator</b>—(Optional) Attach the BGP <b>aggregator</b> path attribute to the aggregate route. You must specify the last AS number that formed the aggregate route (encoded as two octets) for <b>as-number</b>, followed by the IP address of the BGP system that formed the aggregate route for <b>ip-address</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



**as-path**—(Optional) AS path to include with the route. It can include a combination of individual AS path numbers and AS sets. Enclose sets in brackets ( [ ] ). The first AS number in the path represents the AS immediately adjacent to the local AS. Each subsequent number represents an AS that is progressively farther from the local AS, heading toward the origin of the path. You cannot specify a regular expression for **as-path**. You must use a complete, valid AS path.

**atomic-aggregate**—(Optional) Attach the BGP **atomic-aggregate** path attribute to the aggregate route. This path attribute indicates that the local system selected a less specific route instead of a more specific route.

**origin egp**—(Optional) BGP origin attribute that indicates that the path information originated in another AS.

**origin igp**—(Optional) BGP origin attribute that indicates that the path information originated within the local AS.

**origin incomplete**—(Optional) BGP origin attribute that indicates that the path information was learned by some other means.

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                                       |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> <li>• <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i></li> </ul> |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## asm-override-ssm

---

|                                 |                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | asm-override-ssm;                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                 |
| <b>Description</b>              | Enable the routing device to accept any-source multicast join messages (*G) for group addresses that are within the default or configured range of source-specific multicast groups.                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Source-Specific Multicast Groups with Any-Source Override</i></li></ul>                                                                                                                                                                                             |

## authentication-key-chains

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> authentication-key-chains {   key-chain <i>key-chain-name</i> {     description <i>text-string</i>;     key <i>key</i> {       algorithm (md5   hmac-sha-1);       options (basic   isis-enhanced);       secret <i>secret-data</i>;       start-time <i>yyyy-mm-dd.hh:mm:ss</i>;     }     tolerance <i>seconds</i>;   } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit security]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Configure authentication key updates for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol. When the <b>authentication-key-chains</b> statement is configured at the <b>[edit security]</b> hierarchy level, and is associated with the BGP, LDP, or IS-IS protocols at the <b>[edit protocols]</b> hierarchy level or with the BFD protocol using the <b>bfd-liveness-detection</b> statement, authentication key updates can occur without interrupting routing and signaling protocols such as Open Shortest Path First (OSPF) and Resource Reservation Setup Protocol (RSVP).</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i></li> <li>• <i>Example: Configuring BFD Authentication for Static Routes</i></li> <li>• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## autonomous-system

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>autonomous-system <i>autonomous-system</i> &lt;asdot-notation&gt; &lt;loops <i>number</i>&gt; {<br/>    independent-domain &lt;no-attrset&gt;;<br/>}</code>                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options]                                                                                                                                                                                                             |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br><b>asdot-notation</b> option introduced in Junos OS Release 9.3.<br><b>asdot-notation</b> option introduced in Junos OS Release 9.3 for EX Series switches.<br><b>no-attrset</b> option introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers. |
| <b>Description</b>         | Specify the routing device's AS number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

An autonomous system (AS) is a set of routing devices that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it. ASs are identified by a number that is assigned by the Network Information Center (NIC) in the United States (<http://www.isi.edu>).

If you are using BGP on the routing device, you must configure an AS number.

The AS path attribute is modified when a route is advertised to an EBGP peer. Each time a route is advertised to an EBGP peer, the local routing device prepends its AS number to the existing path attribute, and a value of 1 is added to the AS number.

In Junos OS Release 9.1 and later, the numeric range is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. RFC 4893 introduces two new optional transitive BGP attributes, AS4\_PATH and AS4\_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS\_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.

**Options**    ***autonomous-system***—AS number. Use a number assigned to you by the NIC.

**Range:** 1 through 4,294,967,295 ( $2^{32} - 1$ ) in plain-number format for 4-byte AS numbers

In this example, the 4-byte AS number 65,546 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 65546;
}
```

**Range:** 0.0 through 65535.65535 in AS-dot notation format for 4-byte numbers

In this example, 1.10 is the AS-dot notation format for 65,546:

```
[edit]
routing-options {
  autonomous-system 1.10;
}
```

**Range:** 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

In this example, the 2-byte AS number 60,000 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 60000;
}
```

**asdot-notation**—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

**Default:** Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

**loops number**—(Optional) Specify the number of times detection of the AS number in the AS\_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

**Range:** 1 through 10

**Default:** 1



**NOTE:** When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the loops statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.

Use the **independent-domain** option if the loops statement must be enabled only on a subset of routing instances.

The remaining statement is explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Examples: Configuring External BGP Peering*
- *Examples: Configuring Internal BGP Peering*

---

## backup-pe-group

---

**Syntax** `backup-pe-group group-name {  
    backups [ addresses ];  
    local-address address;  
}`

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],  
[edit logical-systems *logical-system-name* routing-options multicast],  
[edit routing-instances *routing-instance-name* routing-options multicast],  
[edit routing-options multicast]

**Release Information** Statement introduced in Junos OS Release 9.0.  
Statement introduced in Junos OS Release 9.5 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

**Description** Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

**Options** **backups *addresses***—Specify the address of backup PE routers for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

**local-address *address***—Specify the address of the local PE router for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

**pe-group-name**—Specify the name for the group of PE routers that provide ingress PE router redundancy for point-to-multipoint LSPs.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Ingress PE Redundancy*
- *Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs*

## backups

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>backups [ <i>addresses</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>backup-pe-group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>backup-pe-group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>backup-pe-group</b> <i>group-name</i>],</p> <p>[edit routing-options multicast <b>backup-pe-group</b> <i>group-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the address of backup PEs for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b><i>addresses</i></b> —Addresses of other PEs in the backup group.                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Ingress PE Redundancy</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                      |

## bandwidth (Multicast Flow Map)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bandwidth ( <i>bps</i>   <i>adaptive</i> );</code>                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">flow-map</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">flow-map</a> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">flow-map</a> ],<br>[edit routing-options multicast <a href="#">flow-map</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                         |
| <b>Description</b>              | Configure the bandwidth property for multicast flow maps.                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>adaptive</b> —Specify that the bandwidth is measured for the flows that are matched by the flow map.<br><br><b>bps</b> —Bandwidth, in bits per second, for the flow map.<br><b>Range:</b> 0 through any amount of bandwidth<br><b>Default:</b> 2 Mbps                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring a Multicast Flow Map</i></li></ul>                                                                                                                                                                                                                                                                                                                                          |



## bfd-liveness-detection (Routing Options Static Route)

**Syntax**

```

bfd-liveness-detection {
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    holddown-interval milliseconds;
    local-address ip-address;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-receive-ttl number;
    multiplier number;
    neighbor address;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}

```

**Hierarchy Level**

```

[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options rib routing-table-name static route destination-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options rib routing-table-name static route destination-prefix qualified-next-hop
 (interface-name | address)],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options static route destination-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options static route destination-prefix qualified-next-hop (interface-name |
 address)],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
 route destination-prefix],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
 route destination-prefix qualified-next-hop (interface-name | address)],
[edit logical-systems logical-system-name routing-options static route destination-prefix],
[edit logical-systems logical-system-name routing-options static route destination-prefix
 qualified-next-hop (interface-name | address)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
 route destination-prefix],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
 route destination-prefix qualified-next-hop (interface-name | address)],
[edit routing-instances routing-instance-name routing-options static route destination-prefix],
[edit routing-instances routing-instance-name routing-options static route destination-prefix
 qualified-next-hop (interface-name | address)],
[edit routing-options rib routing-table-name static route destination-prefix],
[edit routing-options rib routing-table-name static route destination-prefix qualified-next-hop
 (interface-name | address)],
[edit routing-options static route destination-prefix],

```

[edit routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)]

- Release Information** Statement introduced before Junos OS Release 7.4.  
**detection-time threshold** and **transmit-interval threshold** options introduced in Junos OS Release 8.2.  
**local-address** statement introduced in Junos OS Release 8.2.  
**minimum-receive-ttl** statement introduced in Junos OS Release 8.2.  
Support for logical routers introduced in Junos OS Release 8.3.  
**holddown-interval** statement introduced in Junos OS Release 8.5.  
**no-adaptation** statement introduced in Junos OS Release 9.0.  
Support for IPv6 static routes introduced in Junos OS Release 9.1.  
**authentication algorithm**, **authentication key-chain**, and **authentication loose-check** statements introduced in Junos OS Release 9.6.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.  
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
- Description** Configure bidirectional failure detection timers and authentication criteria for static routes.

**Options** **authentication algorithm** *algorithm-name*—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

**authentication key-chain** *key-chain-name*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

**detection-time threshold** *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**holddown-interval** *milliseconds*—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

**Range:** 0 through 255,000

**Default:** 0

**local-address** *ip-address*—Enable a multihop BFD session and configure the source address for the BFD session.

**minimum-interval** *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.

**Range:** 1 through 255,000

**minimum-receive-interval** *milliseconds*—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

**Range:** 1 through 255,000

**minimum-receive-ttl** *number*—Configure the time to live (TTL) for the multihop BFD session.

**Range:** 1 through 255

**Default:** 255

**multiplier** *number*—Configure number of hello packets not received by the neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**neighbor address**—Configure a next-hop address for the BFD session for a next hop specified as an interface name.

**no-adaptation**—Specify for BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold milliseconds**—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295

**transmit-interval minimum-interval milliseconds**—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

**Range:** 1 through 255,000


**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

**Default:** automatic

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                                                                                                |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Example: Configuring BFD for Static Routes</i></li><li>• <i>Example: Configuring BFD Authentication for Static Routes</i></li></ul> |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## bgp-orf-cisco-mode

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bgp-orf-cisco-mode;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit routing-options <b>outbound-route-filter</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|                                 | <p> <b>NOTE:</b> To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Level</b>                    | routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

- Related Documentation**
- *Example: Configuring BGP Prefix-Based Outbound Route Filtering*

## bmp

```
Syntax  bmp {
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
    authentication-key-chain authentication-key-chain;
    connection-mode (active | passive);
    hold-down {
        seconds;
        flaps flaps;
        period seconds;
    }
    initiation-message text;
    local-address address;
    local-port port;
    monitor (disable | enable);
    priority (high | low | medium);
    route-monitoring {
        none;
        post-policy {
            exclude-non-eligible;
        }
        pre-policy {
            exclude-non-feasible;
        }
    }
}
station station-name {
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
    authentication-key-chain authentication-key-chain;
    connection-mode (active | passive);
    hold-down {
        seconds;
        flaps flaps;
        period seconds;
    }
    initiation-message text;
    local-address address;
    local-port port;
    monitor (disable | enable);
    priority (high | low | medium);
    route-monitoring {
        none;
        post-policy {
            exclude-non-eligible;
        }
        pre-policy {
            exclude-non-feasible;
        }
    }
}
station-address (ip-address | name);
station-port port-number;
statistics-timeout seconds;
traceoptions {
```

```
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier>;
    }
}
station-address (ip-address | name);
station-port port-number;
statistics-timeout seconds;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier>;
}
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [bgp](#)],  
[edit logical-systems *logical-system-name* protocols bgp [group](#) *group-name*],  
[edit logical-systems *logical-system-name* protocols bgp group *group-name* [neighbor](#) *address*],  
[edit logical-systems *logical-system-name* [routing-options](#)],  
[edit protocols [bgp](#)],  
[edit protocols bgp [group](#) *group-name*],  
[edit protocols bgp group *group-name* [neighbor](#) *address*],  
[edit [routing-options](#)]

**Release Information** Statement introduced in Junos OS Release 9.5.  
Statement introduced in Junos OS Release 9.5 for EX Series switches.  
Statement introduced in Junos OS Release 12.3 for ACX Series routers.  
Support for BMP version 3 introduced in Junos OS Release 13.3.

**Description** Configure the BGP Monitoring Protocol (BMP), which enables the routing device to collect data from the BGP Adjacency-RIB-In routing tables and periodically send that data to a monitoring station.

**Options** The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring the BGP Monitoring Protocol*




## brief

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (brief   full);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure all AS numbers from all contributing paths to be included in the aggregate or generated route's path.</p> <ul style="list-style-type: none"> <li>• <b>brief</b>—Include only the longest common leading sequences from the contributing AS paths. If this results in AS numbers being omitted from the aggregate route, the BGP <b>ATOMIC_ATTRIBUTE</b> path attribute is included with the aggregate route.</li> <li>• <b>full</b>—Include all AS numbers from all contributing paths in the aggregate or generated route's path. Include this option when configuring an individual route in the <b>route</b> portion of the <b>generate</b> statement to override a <b>retain</b> option specified in the <b>defaults</b> portion of the statement.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>                  | full                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> <li>• <a href="#">aggregate on page 3430</a></li> <li>• <a href="#">generate on page 3465</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## centralized

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | centralized;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols lacp ppm]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Disable distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all PPM processing for LACP packets on the Routing Engine.</p> <p>This statement disables distributed PPM processing for only LACP packets. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the <b>no-delegate-processing</b> statement in the [edit routing-options ppm] hierarchy.</p> |
|                                 | <div><b>BEST PRACTICE:</b> We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.</div>                                                                                                                                                                                 |
| <b>Default</b>                  | Distributed PPM processing is enabled for all packets that use PPM.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure) on page 3424</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP (CLI Procedure) on page 2671</a></li><li>• <i>Configuring Distributed Periodic Packet Management</i></li><li>• <i>Configuring Link Aggregation</i></li></ul>                                                                                                                           |

## community (Routing Options)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>community ([ <i>community-ids</i> ]   no-advertise   no-export   no-export-subconfed   none);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-options (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   <b>static</b>) (defaults   route)]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>         | Associate BGP community information with a static, aggregate, or generated route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>             | No BGP community information is associated with static routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>             | <p><b><i>community-ids</i></b>—One or more community identifiers. The <b><i>community-ids</i></b> format varies according to the type of attribute that you use.</p> <p>The BGP community attribute format is <b><i>as-number:community-value</i></b>:</p> <ul style="list-style-type: none"> <li>• <b><i>as-number</i></b>—AS number of the community member. It can be a value from 1 through 65,535. The AS number can be a decimal or hexadecimal value.</li> <li>• <b><i>community-value</i></b>—Identifier of the community member. It can be a number from 0 through 65,535.</li> </ul> <p>For more information about BGP community attributes, see the “Configuring the Extended Communities Attribute” section in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>.</p> <p>For specifying the BGP community attribute only, you also can specify <b><i>community-ids</i></b> as one of the following well-known community names defined in RFC 1997:</p> <ul style="list-style-type: none"> <li>• <b>no-advertise</b>—Routes containing this community name are not advertised to other BGP peers.</li> <li>• <b>no-export</b>—Routes containing this community name are not advertised outside a BGP confederation boundary.</li> </ul>                                                 |

- **no-export-subconfed**—Routes containing this community are advertised to IBGP peers with the same AS number, but not to members of other confederations.



**NOTE:** Extended community attributes are not supported at the [edit routing-options] hierarchy level. You must configure extended communities at the [edit policy-options] hierarchy level. For information about configuring extended communities, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

|                                 |                                                                |
|---------------------------------|----------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.           |
|                                 | routing-control—To add this statement to the configuration.    |
| <b>Related Documentation</b>    | • <i>Examples: Configuring Static Routes</i>                   |
|                                 | • <i>Example: Summarizing Routes Through Route Aggregation</i> |
|                                 | • <i>Example: Conditionally Generating Static Routes</i>       |
|                                 | • <a href="#">aggregate on page 3430</a>                       |
|                                 | • <a href="#">generate on page 3465</a>                        |
|                                 | • <a href="#">static on page 3522</a>                          |

## confederation

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>confederation <i>confederation-autonomous-system</i> members [ <i>autonomous-systems</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Specify the routing device's confederation AS number.</p> <p>If you administer multiple ASs that contain a very large number of BGP systems, you can group them into one or more <i>confederations</i>. Each confederation is identified by its own AS number, which is called a <i>confederation AS number</i>. To external ASs, a confederation appears to be a single AS. Thus, the internal topology of the ASs making up the confederation is hidden.</p> <p>The BGP path attributes <b>NEXT_HOP</b>, <b>LOCAL_PREF</b>, and <b>MULTI_EXIT_DISC</b>, which normally are restricted to a single AS, are allowed to be propagated throughout the ASs that are members of the same confederation.</p> <p>Because each confederation is treated as if it were a single AS, you can apply the same routing policy to all the ASs that make up the confederation.</p> <p>Grouping ASs into confederations reduces the number of BGP connections required to interconnect ASs.</p> <p>If you are using BGP, you can enable the local routing device to participate as a member of an AS confederation. To do this, include the <b>confederation</b> statement.</p> <p>Specify the AS confederation identifier, along with the peer AS numbers that are members of the confederation.</p> <p>Note that peer adjacencies do not form if two BGP neighbors disagree about whether an adjacency falls within a particular confederation.</p> |
| <b>Options</b>                  | <p><b><i>autonomous-systems</i></b>—AS numbers of the confederation members.<br/><b>Range:</b> 1 through 65,535</p> <p><b><i>confederation-autonomous-system</i></b>—Confederation AS number. Use one of the numbers assigned to you by the NIC.<br/><b>Range:</b> 1 through 65,535</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |


**Related Documentation**    • *Example: Configuring BGP Confederations*

## **disable (Routing Options)**

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <b>routing-options</b> graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> <b>routing-options</b> graceful-restart],<br>[edit routing-instances <i>routing-instance-name</i> <b>routing-options</b> graceful-restart],<br>[edit <b>routing-options</b> graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                          |
| <b>Description</b>              | Disable graceful restart.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | • <i>Junos OS High Availability Library for Routing Devices</i>                                                                                                                                                                                                                                                                                                                             |

## discard

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | discard;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.</p> <p>To propagate static routes into the routing protocols, include the <b>discard</b> statement when you define the route, along with a routing policy.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|                                 | <p> <b>NOTE:</b> In other vendors' software, a common way to propagate static routes into routing protocols is to configure the routes so that the next-hop routing device is the loopback address (commonly, 127.0.0.1). However, configuring static routes in this way (by including a statement such as <b>route address/mask-length next-hop 127.0.0.1</b>) does not propagate the static routes, because the forwarding table ignores static routes whose next-hop routing device is the loopback address.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>                  | When an aggregate route becomes active, it is installed in the routing table with a reject next hop, which means that ICMP unreachable messages are sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- [aggregate on page 3430](#)
- [generate on page 3465](#)

## export (Routing Options)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-name</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <a href="#">forwarding-table</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options <a href="#">forwarding-table</a> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options <a href="#">forwarding-table</a> ],<br>[edit routing-options <a href="#">forwarding-table</a> ]                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Apply one or more policies to routes being exported from the routing table into the forwarding table.</p> <p>In the <b>export</b> statement, list the name of the routing policy to be evaluated when routes are being exported from the routing table into the forwarding table. Only active routes are exported from the routing table.</p> <p>You can reference the same routing policy one or more times in the same or a different <b>export</b> statement.</p> <p>You can apply export policies to routes being exported from the routing table into the forwarding table for the following features:</p> <ul style="list-style-type: none"><li>• Per-packet load balancing</li><li>• Class of service (CoS)</li></ul> |
| <b>Options</b>                  | <i>policy-name</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Load Balancing BGP Traffic</i></li><li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li><li>• <i>How a Routing Policy Is Evaluated</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |



## export-rib

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>export-rib routing-table-name;</code>                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit routing-options <b>rib-groups</b> <i>group-name</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                 |
| <b>Description</b>         | <p>Specify the name of the routing table from which Junos OS should export routing information. For any individual RIB group, only one table can be specified in the <b>export-rib</b> statement.</p>                                                                                                                                                                                                                                                          |

The **export-rib** statement specifies the source table from which routing information is advertised.

One common use of the **export-rib** statement is interdomain routing. The export RIB is the table used when BGP extracts routes to advertise to peers. In multicast interdomain routing, for example, the export RIB is likely to be inet.2.

Another use of **export-rib** is dynamic route leaking between the global routing table (inet.0) and a VRF routing table (*instance.inet.0*). For example, you can use a RIB group to copy routes learned in the VRF into the global routing table, inet.0, or copy routes learned in inet.0 into a VRF. You define the use of this RIB group in the VRF's BGP configuration. In a routing policy you can do dynamic filtering of routes. For instance, you can use an import policy to only copy routes with certain communities into the global routing table.

For example:

```
rib-groups {
  rib-interface-routes-v4 {
    import-rib [ inet.0 VRF.inet.0 ];
  }
  rib-import-VRF-routes-to-inet0-v4 {
    export-rib VRF.inet.0;
    import-rib [ VRF.inet.0 inet.0 ];
    import-policy rib-import-VRF-routes-to-inet0-v4;
  }
  rib-import-inet0-routes-to-VRF-v4 {
    export-rib inet.0;
    import-rib [ inet.0 VRF.inet.0 ];
    import-policy rib-import-inet0-routes-to-VRF-v4;
  }
}
routing-options {
  interface-routes {
    rib-group {
```

```

        inet rib-interface-routes-v4;
    }
}
protocols {
    bgp {
        group iBGP-peers {
            type internal;
            family inet {
                unicast {
                    rib-group rib-import-inet0-routes-to-VRF-v4;
                }
            }
        }
    }
}
routing-instances {
    VRF {
        routing-options {
            interface-routes {
                rib-group {
                    inet rib-interface-routes-v4;
                }
            }
        }
        protocols {
            bgp {
                group peersin-VRF {
                    family inet {
                        unicast {
                            rib-group rib-import-VRF-routes-to-inet0-v4;
                        }
                    }
                }
            }
        }
    }
}
}

```

**Options** *routing-table-name*—Routing table group name.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- *Example: Configuring a PIM RPF Routing Table*
- *Example: Configuring DVMRP to Announce Unicast Routes*
- *Example: Configuring a Dedicated PIM RPF Routing Table*
- *Example: Configuring Any-Source Multicast for Draft-Rosen VPNs*
- [import-rib on page 3468](#)
- *passive*

## fate-sharing

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>fate-sharing {   group <i>group-name</i> {     cost <i>value</i>;     from <i>address</i> &lt;to <i>address</i>&gt;;   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Specify a backup path in case the primary path becomes unusable.</p> <p>You specify one or more objects with common characteristics within a group. All objects are treated as /32 host addresses. The objects can be a LAN interface, a router ID, or a point-to-point link. Sequence is insignificant.</p> <p>Changing the fate-sharing database does not affect existing established LSPs until the next CSPF reoptimization. The fate-sharing database does affect fast-reroute detour path computations.</p>                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>cost <i>value</i></b>—Cost assigned to the group.<br/> <b>Range:</b> 1 through 65,535<br/> <b>Default:</b> 1</p> <p><b>from <i>address</i></b>—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate <b>from</b> addresses in the group.</p> <p><b>group <i>group-name</i></b>—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.</p> <p><b>to <i>address</i></b>—(Optional) Address of egress router. For point-to-point link objects, you must specify both a <b>from</b> and a <b>to</b> address.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Ingress Router for MPLS-Signaled LSPs</i></li> <li>• <i>Junos OS MPLS Applications Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## flow

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> flow {     route <i>name</i> {         match {             <i>match-conditions</i>;         }         term-order (legacy   standard);         then {             <i>actions</i>;         }     }     firewall-install-disable;     term-order (legacy   standard);     validation {         traceoptions {             file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;             flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;;         }     } } </pre>                                              |
| <b>Hierarchy Level</b>     | <p>[edit routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options]</p>                                                                                                                                                                                                                                                                               |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>term-order</b> statement introduced in Junos OS Release 10.0</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>firewall-install-disable</b> statement introduced in Junos OS Releases 12.1X48 and 12.3 for PTX Series routers.</p>                                                                                                                                                                               |
| <b>Description</b>         | Configure a flow route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>             | legacy                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>             | <p><b>actions</b>—An action to take if conditions match.</p> <p><b>firewall-install-disable</b>—(PTX Series routers only) Disable installing flow-specification firewall filters in the firewall process (dfwd).</p> <p><b>Default:</b> If you omit the <b>firewall-install-disable</b> statement, the default behavior is <b>firewall-install-disable</b> mode.</p> <p><b>match-conditions</b>—Match packets to these conditions.</p> <p><b>route <i>name</i></b>—Name of the flow route.</p> <p><b>standard</b>—Specify to use version 7 or later of the flow-specification algorithm.</p> |

**term-order (legacy | standard)**—Specify the version of the flow-specification algorithm.

- **legacy**—Use version 6 of the flow-specification algorithm.
- **standard**—Use version 7 of the flow-specification algorithm.

**then**—Actions to take on matching packets.

The remaining statements are explained separately.

|                                 |                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.                                        |
|                                 | routing-control—To add this statement to the configuration.                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Flow Routes</i></li> </ul> |

## flow-map

|                                 |                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>flow-map flow-map-name {   bandwidth (bps   adaptive);   forwarding-cache {     timeout (never non-discard-entry-only   minutes);   }   policy [ policy-names ];   redundant-sources [ addresses ]; }</pre>                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                              |
| <b>Description</b>              | Configure multicast flow maps.                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>flow-map-name</b>—Name of the flow-map.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring a Multicast Flow Map</i></li> </ul>                                                                                                                                                                                                                                                   |

## forwarding-cache (Flow Maps)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | forwarding-cache {<br>timeout (minutes   never non-discard-entry-only );<br>}                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure multicast forwarding cache properties for the flow map.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                         |

## forwarding-cache (Multicast)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> forwarding-cache {   allow-maximum;   family (inet   inet6) {     threshold {       log-warning value;       suppress value &lt;reuse value&gt;;     }   }   threshold {     log-warning value;     suppress value &lt;reuse value&gt;;   }   timeout minutes; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>         | <p>Configure multicast forwarding cache properties. These properties include threshold suppression and reuse limits, the threshold at which a warning message is logged, and timeout values.</p> <p>Specify a value for the threshold at which to suppress new multicast forwarding cache entries and an optional reuse value for the threshold at which the router begins to create new multicast forwarding cache entries. The range for both is from 1 through 200,000. If configured, the reuse value should be less than the suppression threshold value. The suppression value is mandatory. If you do not specify the optional reuse value, then the number of multicast forwarding cache entries is limited to the suppression value. A new entry is created as soon as the number of multicast forwarding cache entries falls below the suppression value.</p> <p>You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the <b>threshold</b> statement globally for the multicast forwarding cache or including the <b>family</b> statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p> |
| <b>Default</b>             | By default, there are no limits on the number of multicast forwarding cache entries.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>             | <p><b>family (inet   inet6)</b>—(Optional) Apply the configured thresholds to either IPv4 or IPv6 multicast forwarding cache entries.</p> <p><b>Default:</b> By default, the configured thresholds are applied to both IPv4 and IPv6 multicast forwarding cache entries.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring the Multicast Forwarding Cache*

---

## forwarding-table

---

**Syntax** forwarding-table {  
    chained-composite-next-hop;  
    export [ *policy-name* ];  
    (indirect-next-hop | no-indirect-next-hop);  
    (indirect-next-hop-change-acknowledgements |  
        no-indirect-next-hop-change-acknowledgements);  
    krt-nexthop-ack-timeout *interval*;  
    unicast-reverse-path (active-paths | feasible-paths);  
}

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-options],  
[edit routing-options]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure information about the routing device's forwarding table.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Load Balancing BGP Traffic*



## generate

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> generate {   defaults {     generate-options;   }   route destination-prefix {     policy policy-name;     generate-options;   } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>rib</b> <i>routing-table-name</i>],</p> <p>[edit routing-options],</p> <p>[edit routing-options <b>rib</b> <i>routing-table-name</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | Configure generated routes, which are used as routes of last resort.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>             | <p><b>defaults</b>—(Optional) Specify global generated route options. These options only set default attributes inherited by all newly created generated routes. These are treated as global defaults and apply to all the generated routes you configure in the <b>generate</b> statement.</p> <p><b>generate-options</b>—Additional information about generated routes, which is included with the route when it is installed in the routing table. Specify zero or more of the following options in <b>generate-options</b>. Each option is explained separately.</p> <ul style="list-style-type: none"> <li>• (<b>active</b>   <b>passive</b>);</li> <li>• <b>as-path</b> <i>&lt;as-path&gt;</i> <i>&lt;origin (egp   igp   incomplete)&gt;</i> <i>&lt;atomic-aggregate&gt;</i> <i>&lt;aggregator as-number in-address&gt;</i>;</li> <li>• (<b>brief</b>   <b>full</b>);</li> <li>• <b>community</b> [ <i>community-ids</i> ];</li> <li>• <b>discard</b>;</li> <li>• (<b>metric</b>   <i>metric2</i>   <i>metric3</i>   <i>metric4</i>) <i>value</i> <i>&lt;type type&gt;</i>;</li> <li>• (<b>preference</b>   <i>preference2</i>   <b>color</b>   <i>color2</i>) <i>preference</i> <i>&lt;type type&gt;</i>;</li> <li>• <b>tag</b> <i>metric type number</i>;</li> </ul> <p><b>route destination-prefix</b>—Configure a non-default generated route:</p> <ul style="list-style-type: none"> <li>• <b>default</b>—For the default route to the destination. This is equivalent to specifying an IP address of 0.0.0.0/0.</li> </ul> |

- *destination-prefix/prefix-length—/destination-prefix* is the network portion of the IP address, and *prefix-length* is the destination prefix length.

The [policy](#) statement is explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Conditionally Generating Static Routes*

---

## import (Routing Options)

---

**Syntax** `import [ policy-names ];`

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options resolution [rib](#)],  
[edit logical-systems *logical-system-name* routing-options resolution [rib](#)],  
[edit routing-instances *routing-instance-name* routing-options resolution [rib](#)],  
[edit routing-options resolution [rib](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Specify one or more import policies to use for route resolution.


**Options** *policy-names*—Name of one or more import policies.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.


**Related Documentation**

- *Example: Configuring Route Resolution on PE Routers*


## import-policy

|                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                      | <code>import-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                                             | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit routing-options <b>rib-groups</b> <i>group-name</i>]</p> |
| <b>Release Information</b>                                                                                                                                                                                                         | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                  |
| <b>Description</b>                                                                                                                                                                                                                 | <p>Apply one or more policies to routes imported into the routing table group. The <b>import-policy</b> statement complements the <b>import-rib</b> statement and cannot be used unless you first specify the routing tables to which routes are being imported.</p>                                                                                                                                                                                           |
| <div>  <p><b>NOTE:</b> On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                                                                                                                                                                                                     | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b>                                                                                                                                                                                                    | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                                                                                       | <ul style="list-style-type: none"> <li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li> <li>• <a href="#">export-rib on page 3457</a></li> <li>• <i>passive</i></li> </ul>                                                                                                                                                                                                                                   |

## import-rib

|                                                                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                        | <code>import-rib [ <i>routing-table-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                                               | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit routing-options <b>rib-groups</b> <i>group-name</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                                                                                                                                                                                                           | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                                                   | <p>Specify the name of the routing table into which Junos OS should import routing information. The first routing table name you enter is the primary routing table. Any additional names you enter identify secondary routing tables. When a protocol imports routes, it imports them into the primary and any secondary routing tables. If the primary route is deleted, the secondary route also is deleted. For IPv4 import routing tables, the primary routing table must be <b>inet.0</b> or <b>routing-instance-name.inet.0</b>. For IPv6 import routing tables, the primary routing table must be <b>inet6.0</b>.</p> <p>In Junos OS Release 9.5 and later, you can configure an IPv4 import routing table that includes both IPv4 and IPv6 routing tables. Including both types of routing tables permits you, for example, to populate an IPv6 routing table with IPv6 addresses that are compatible with IPv4. In releases prior to Junos OS Release 9.5, you could configure an import routing table with only either IPv4 or IPv6 routing tables.</p> |
| <div>  <p><b>NOTE:</b> On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                                                                                                                                       | <i>routing-table-names</i> —Name of one or more routing tables.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b>                                                                                                                                                                                                      | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                                                                         | <ul style="list-style-type: none"> <li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li> <li>• <a href="#">export-rib on page 3457</a></li> <li>• <i>passive</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## indirect-next-hop

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | (indirect-next-hop   no-indirect-next-hop);                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | [edit logical-systems <i>logical-system-name</i> routing-options <a href="#">forwarding-table</a> ],<br>[edit routing-options <a href="#">forwarding-table</a> ]                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                 |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Enable indirectly connected next hops for route convergence. This statement is implemented on the Packet Forward Engine to speed up forwarding information base (FIB) updates. Configuring this statement significantly speeds convergence times. The only downside of configuring this statement is that some additional FIB memory overhead is required. Unless routes have an extremely high number of next hops, this increased memory usage should not be noticeable. |
| <div>  <b>NOTE:</b> <ul style="list-style-type: none"> <li>When virtual private LAN service (VPLS) is configured on the routing device, the <code>indirect-next-hop</code> statement is configurable at the [edit routing-options <code>forwarding-table</code>] hierarchy level. However, this configuration is not applicable to indirect nexthops specific to VPLS routing instances.</li> <li>By default, the Junos Trio Modular Port Concentrator (MPC) chipset on MX Series routers is enabled with indirectly connected next hops, and this cannot be disabled using the <code>no-indirect-next-hop</code> statement.</li> <li>By default, indirectly connected next hops are enabled on PTX Series routers.</li> </ul> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <code>indirect-next-hop</code> —Enable indirectly connected next hops.<br><code>no-indirect-next-hop</code> —Explicitly disable indirect next hops.                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <ul style="list-style-type: none"> <li><i>Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine</i></li> </ul>                                                                                                                                                                                                                                                                                                            |

## install (Routing Options)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (install   no-install);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>static</b> (defaults   route)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],<br>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],<br>[edit logical-systems <i>logical-system-name</i> routing-options <b>static</b> (defaults   route)],<br>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],<br>[edit routing-instances <i>routing-instance-name</i> routing-options <b>static</b> (defaults   route)],<br>[edit routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)]<br>[edit routing-options <b>static</b> (defaults   route)] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure whether Junos OS installs all static routes into the forwarding table. Even if you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>install</b> —Explicitly install all static routes into the forwarding table. Include this statement when configuring an individual route in the <b>route</b> portion of the <b>static</b> statement to override a <b>no-install</b> option specified in the <b>defaults</b> portion of the statement.<br><br><b>no-install</b> —Do not install the route into the forwarding table, even if it is the route with the lowest preference.<br><br><b>Default:</b> install                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Examples: Configuring Static Routes</i></li><li>• <a href="#">static on page 3522</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## instance-export

|                                 |                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>instance-export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                   |
| <b>Description</b>              | Apply one or more policies to routes being exported from a routing instance.                                                                                                                                                                                                                |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more export policies.                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>                                                                                                                                       |

## instance-import

|                                 |                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>instance-import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                   |
| <b>Description</b>              | Apply one or more policies to routes being imported into a routing instance.                                                                                                                                                                                                                |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>                                                                                                                                       |

## interface (Multicast Static Routes)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface <i>interface-names</i> {<br/>    disable;<br/>    maximum-bandwidth <i>bps</i>;<br/>    no-qos-adjust;<br/>    reverse-oif-mapping {<br/>        no-qos-adjust;<br/>    }<br/>    subscriber-leave-timer <i>seconds</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i><br/>    routing-options <a href="#">multicast</a>],<br/>[edit logical-systems <i>logical-system-name</i> routing-options <a href="#">multicast</a>],<br/>[edit routing-instances <i>routing-instance-name</i> routing-options <a href="#">multicast</a>],<br/>[edit routing-options <a href="#">multicast</a>]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Enable multicast traffic on an interface.</p> <p>By default, multicast packets are forwarded by enabling Protocol Independent Multicast (PIM) on an interface. PIM adds multicast routes into the routing table.</p> <p>You can also configure multicast packets to be forwarded over a static route, such as a static route associated with an LSP next hop. Multicast packets are accepted on an interface and forwarded over a static route in the forwarding table. This is useful when you want to enable multicast traffic on a specific interface without configuring PIM on the interface.</p> <p>You cannot enable multicast traffic on an interface and configure PIM on the same interface simultaneously.</p> <p>Static routes must be configured before you can enable multicast on an interface. Configuring the <b>interface</b> statement alone does not install any routes into the routing table. This feature relies on the static route configuration.</p> |
| <b>Options</b>                  | <p><b><i>interface-names</i></b>—Name of one or more interfaces on which to enable multicast traffic.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Defining Interface Bandwidth Maximums</i></li><li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |



## interface (Routing Options)

|                            |                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>interface <i>interface-names</i> {     maximum-bandwidth <i>bps</i>;     no-qos-adjust;     reverse-oif-mapping {         no-qos-adjust;     }     subscriber-leave-timer <i>seconds</i>; }</pre>                                                                                                                                                 |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                              |
| <b>Description</b>         | Enable multicast traffic on an interface.                                                                                                                                                                                                                                                                                                              |



**TIP:** You cannot enable multicast traffic on an interface by using the `routing-options multicast interface` statement and configure PIM on the interface.

|                                 |                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Names of the physical or logical interface.</p> <p>The remaining statements are explained separately.</p>                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Defining Interface Bandwidth Maximums</i></li> <li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li> </ul> |

## interface-routes

---

**Syntax**

```
interface-routes {  
    family (inet | inet6) {  
        export {  
            lan;  
            point-to-point;  
        }  
    }  
    rib-group group-name;  
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],  
[edit logical-systems *logical-system-name* routing-options],  
[edit routing-instances *routing-instance-name* routing-options],  
[edit routing-options]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.



**NOTE:** On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

---

**Description** Associate a routing table group with the routing device's interfaces, and specify routing table groups into which interface routes are imported.

By default, IPv4 interface routes (also called direct routes) are imported into routing table **inet.0**, and IPv6 interface routes are imported into routing table **inet6.0**. If you are configuring alternate routing tables for use by some routing protocols, it might be necessary to import the interface routes into the alternate routing tables. To define the routing tables into which interface routes are imported, you create a routing table group and associate it with the routing device's interfaces.

To create the routing table groups, include the **passive** statement at the **[edit routing-options]** hierarchy level.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group.

To export local routes, include the **export** statement.

To export LAN routes, include the **lan** option. To export point-to-point routes, include the **point-to-point** option.

Only local routes on point-to-point interfaces configured with a destination address are exportable.

- Options**
- inet**—Specify the IPv4 address family.
  - inet6**—Specify the IPv6 address family.
  - lan**—Export LAN routes.
  - point-to-point**—Export point-to-point routes.

The remaining statements are explained separately.

- Required Privilege Level**
- routing—To view this statement in the configuration.
  - routing-control—To add this statement to the configuration.

- Related Documentation**
- *Example: Importing Direct and Static Routes Into a Routing Instance*
  - *Example: Configuring Multiple Routing Instances of OSPF*
  - *passive*

## local-address (Routing Options)

**Syntax** `local-address address;`

**Hierarchy Level**

```
[edit logical-systems logical-system-name routing-instances routing-instance-name
  routing-options multicast backup-pe-group group-name],
[edit logical-systems logical-system-name routing-options multicast backup-pe-group
  group-name],
[edit routing-instances routing-instance-name routing-options multicast backup-pe-group
  group-name],
[edit routing-options multicast backup-pe-group group-name]
```

**Release Information**

Statement introduced in Junos OS Release 9.0.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure the address of the local PE for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

**Options** **address**—Address of local PEs in the backup group.

**Required Privilege Level**

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

- Related Documentation**
- *Example: Configuring Ingress PE Redundancy*

## martians

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>martians {<br/>    destination-prefix match-type &lt;allow&gt;;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems logical-system-name routing-instances routing-instance-name<br/>    routing-options],<br/>[edit logical-systems logical-system-name routing-instances routing-instance-name<br/>    routing-options rib routing-table-name],<br/>[edit logical-systems logical-system-name routing-options],<br/>[edit logical-systems logical-system-name routing-options rib routing-table-name],<br/>[edit routing-instances routing-instance-name routing-options],<br/>[edit routing-instances routing-instance-name routing-options rib routing-table-name],<br/>[edit routing-options],<br/>[edit routing-options rib routing-table-name]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure martian addresses.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>allow</b>—(Optional) Explicitly allow a subset of a range of addresses that has been disallowed. The <b>allow</b> option is the only supported action.</p> <p><b>destination-prefix</b>—Destination route you are configuring:</p> <ul style="list-style-type: none"><li>• <b>destination-prefix/prefix-length—destination-prefix</b> is the network portion of the IP address, and <b>prefix-length</b> is the destination prefix length.</li><li>• <b>default</b>—Default route to use when routing packets do not match a network or host in the routing table. This is equivalent to specifying the IP address <b>0.0.0.0/0</b>.</li></ul> <p><b>match-type</b>—Criteria that the destination must match:</p> <ul style="list-style-type: none"><li>• <b>exact</b>—Exactly match the route's mask length.</li><li>• <b>longer</b>—The route's mask length is greater than the specified mask length.</li><li>• <b>orlonger</b>—The route's mask length is equal to or greater than the specified mask length.</li><li>• <b>through destination-prefix</b>—The route matches the first prefix, the route matches the second prefix for the number of bits in the route, and the number of bits in the route is less than or equal to the number of bits in the second prefix.</li><li>• <b>upto prefix-length</b>—The route's mask length falls between the two destination prefix lengths, inclusive.</li></ul> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**Related Documentation** • *Example: Configuring Martian Addresses*

## maximum-bandwidth (Routing Options)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-bandwidth <i>bps</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit dynamic-profiles <i>profile-name</i> routing-instances <i>instance-name</i> routing-options multicast interface <i>interface-name</i>],</p> <p>[edit dynamic-profiles <i>profile-name</i> routing-options multicast interface <i>interface-name</i>]</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-options multicast <a href="#">interface</a> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>dynamic-profiles</b> hierarchy level added in Junos OS Release 11.2.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure the multicast bandwidth for the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b><i>bps</i></b>—Bandwidth rate, in bits per second, for the multicast interface.</p> <p><b>Range:</b> 0 through any amount of bandwidth</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | • <i>Example: Defining Interface Bandwidth Maximums</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## maximum-paths

|                            |                                                                                                                                                                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>maximum-paths <i>path-limit</i> &lt;log-interval <i>seconds</i>&gt; &lt;log-only   threshold <i>value</i>&gt;;</code>                                                                                                                                                                 |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                  |
| <b>Description</b>         | Configure a limit for the number of routes installed in a routing table based upon the route path.                                                                                                                                                                                          |



**NOTE:** The `maximum-paths` statement is similar to the `maximum-prefixes` statement. The `maximum-prefixes` statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The `maximum-paths` limit applies the total number of routes (two). The `maximum-prefixes` limit applies to the total number of unique prefixes (one).

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b> | <p><code>log-interval <i>seconds</i></code>—(Optional) Minimum time interval (in seconds) between log messages.</p> <p><b>Range:</b> 5 through 86,400</p> <p><code>log-only</code>—(Optional) Sets the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.</p> <p><code><i>path-limit</i></code>—Maximum number of routes. If this limit is reached, a warning is triggered and additional routes are rejected.</p> <p><b>Range:</b> 1 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> No default</p> <p><code>threshold <i>value</i></code>—(Optional) Percentage of the maximum number of routes that starts triggering a warning. You can configure a percentage of the <code><i>path-limit</i></code> value that starts triggering the warnings.</p> <p><b>Range:</b> 1 through 100</p> |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|




**NOTE:** When the number of routes reaches the **threshold** value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *path-limit* value, then additional routes are rejected.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs*

## maximum-prefixes

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <code>maximum-prefixes <i>prefix-limit</i> &lt;log-interval <i>seconds</i>&gt; &lt;log-only   threshold <i>percentage</i>&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <p>Configure a limit for the number of routes installed in a routing table based upon the route prefix.</p> <p>Using a prefix limit, you can curtail the number of prefixes received from a CE router in a VPN. Prefix limits apply only to dynamic routing protocols and are not applicable to static or interface routes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <div>  <p><b>NOTE:</b> The <code>maximum-prefixes</code> statement is similar to the <code>maximum-paths</code> statement. The <code>maximum-prefixes</code> statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:</p> <pre> OSPF 10.10.10.0/24 ISIS 10.10.10.0/24 </pre> <p>These are two routes, but only one destination (prefix). The <code>maximum-paths</code> limit applies the total number of routes (two). The <code>maximum-prefixes</code> limit applies to the total number of unique prefixes (one).</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <p><b>log-interval <i>seconds</i></b>—(Optional) Minimum time interval (in seconds) between log messages.</p> <p><b>Range:</b> 5 through 86,400</p> <p><b>log-only</b>—(Optional) Sets the prefix limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.</p> <p><b><i>prefix-limit</i></b>—Maximum number of route prefixes. If this limit is reached, a warning is triggered and any additional routes are rejected.</p> <p><b>Range:</b> 1 through 4,294,967,295</p> <p><b>Default:</b> No default</p> <p><b>threshold <i>value</i></b>—(Optional) Percentage of the maximum number of prefixes that starts triggering a warning. You can configure a percentage of the <i>prefix-limit</i> value that starts triggering the warnings.</p> |



**Range:** 1 through 100



**NOTE:** When the number of routes reaches the threshold value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *prefix-limit* value, then additional routes are rejected.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs*

## med-igp-update-interval

**Syntax** med-igp-update-interval *minutes*;

**Hierarchy Level** [edit routing-options]

**Release Information** Statement introduced in Junos OS Release 9.0  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure a timer for how long to delay updates for the multiple exit discriminator (MED) path attribute for BGP groups and peers configured with the **metric-out igp offset delay-med-update** statement. The timer delays MED updates for the interval configured unless the MED is lower than the previously advertised attribute or another attribute associated with the route has changed or if the BGP peer is responding to a refresh route request.

**Options** *minutes*—Interval to delay MED updates.  
**Range:** 10 through 600  
**Default:** 10 minutes

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates*
- [metric-out on page 3004](#)

## metric (Aggregate, Generated, or Static Route)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (metric   metric2   metric3   metric4) <i>metric</i> <type type>;                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options ( <a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a> ) (defaults   route)],<br>[edit routing-options ( <a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a> ) (defaults   route)]                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                  |
| <b>Description</b>              | Specify the metric value for an aggregate, generated, or static route. You can specify up to four metric values, starting with <b>metric</b> (for the first metric value) and continuing with <b>metric2</b> , <b>metric3</b> , and <b>metric4</b> .                                                                                                                                                                                     |
| <b>Options</b>                  | <b>metric</b> —Metric value.<br><b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ )<br><br><b>type type</b> —(Optional) Type of route.<br><br>When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric.<br><b>Range:</b> 1 through 16 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Summarizing Static Routes Through Route Aggregation</i></li><li>• <i>Example: Conditionally Generating Static Routes</i></li><li>• <a href="#">aggregate on page 3430</a></li><li>• <a href="#">generate on page 3465</a></li><li>• <a href="#">static on page 3522</a></li></ul>                                                                                                    |

## multicast (Routing Options)

**Syntax**

```
multicast {
  forwarding-cache {
    threshold suppress value <reuse value>;
  }
  interface interface-name {
    enable;
  }
  scope scope-name {
    interface [ interface-names ];
    prefix destination-prefix;
  }
  ssm-groups {
    address;
  }
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],  
[edit logical-systems *logical-system-name* routing-options],  
[edit routing-instances *routing-instance-name* routing-options],  
[edit routing-options]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

**Description** Configure generic multicast properties.



**NOTE:** You cannot apply a scoping policy to a specific routing instance. All scoping policies are applied to all routing instances. However, you can apply the `scope` statement to a specific routing instance.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Examples: Configuring Administrative Scoping*
- *Example: Configuring Source-Specific Multicast Groups with Any-Source Override*
- *Examples: Configuring the Multicast Forwarding Cache*
- *Multicast Protocols Feature Guide for Routing Devices*
- ([indirect-next-hop on page 3469](#) | no-indirect-next-hop)

## no-qos-adjust

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-qos-adjust;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface interface-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface interface-name</a> <a href="#">reverse-oif-mapping</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">interface interface-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">interface interface-name</a> <a href="#">reverse-oif-mapping</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface interface-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface interface-name</a> <a href="#">reverse-oif-mapping</a>],</p> <p>[edit routing-options multicast <a href="#">interface interface-name</a>],</p> <p>[edit routing-options multicast <a href="#">interface interface-name</a> <a href="#">reverse-oif-mapping</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Statement added to [edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], and [edit routing-options multicast interface <i>interface-name</i>] hierarchy levels in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Disable hierarchical bandwidth adjustment for all subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## options (Routing Options)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>options {   syslog (level <i>level</i>   upto level <i>level</i>); }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Configure the types of system logging messages sent about the routing protocols process to the system message logging file. These messages are also displayed on the system console. You can log messages at a particular level, or up to and including a particular level.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>level <i>level</i></b>—Severity of the message. It can be one or more of the following levels, in order of decreasing urgency:</p> <ul style="list-style-type: none"> <li>• <b>alert</b>—Conditions that should be corrected immediately, such as a corrupted system database.</li> <li>• <b>critical</b>—Critical conditions, such as hard drive errors.</li> <li>• <b>debug</b>—Software debugging messages.</li> <li>• <b>emergency</b>—Panic or other conditions that cause the system to become unusable.</li> <li>• <b>error</b>—Standard error conditions.</li> <li>• <b>info</b>—Informational messages.</li> <li>• <b>notice</b>—Conditions that are not error conditions, but might warrant special handling.</li> <li>• <b>warning</b>—System warning messages.</li> </ul> <p><b>upto level <i>level</i></b>—Log all messages up to a particular level.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <b>syslog</b> in the <i>Junos OS Administration Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## pim-to-igmp-proxy

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>pim-to-igmp-proxy {<br/>  upstream-interface [ interface-names ];<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast]                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Configure the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Internet Group Management Protocol (IGMP) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the <b>pim-to-igmp-proxy</b> statement enables you to use IGMP to forward IPv4 multicast traffic across the PIM sparse mode domains.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM-to-IGMP Message Translation</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## pim-to-mld-proxy

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>pim-to-mld-proxy {<br/>    upstream-interface [ interface-names ];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast]                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Multicast Listener Discovery (MLD) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the <b>pim-to-mld-proxy</b> statement enables you to use MLD to forward IPv6 multicast traffic across the PIM sparse mode domains.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Configuring PIM-to-MLD Message Translation</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## policy (Aggregate and Generated Routes)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>policy <i>policy-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | <p>Associate a routing policy when configuring an aggregate or generated route's destination prefix in the <b>routes</b> part of the <b>aggregate</b> or <b>generate</b> statement. This provides the equivalent of an import routing policy filter for the destination prefix. That is, each potential contributor to an aggregate route, along with any aggregate options, is passed through the policy filter. The policy then can accept or reject the route as a contributor to the aggregate route.</p> <p>If the contributor is accepted, the policy can modify the default preferences. The contributor with the numerically smallest prefix becomes the most preferred, or <i>primary</i>, contributor. A rejected contributor still can contribute to a less specific aggregate route. If you do not specify a policy filter, all candidate routes contribute to an aggregate route.</p> <p>The following algorithm is used to compare two generated contributing routes in order to determine which one is the primary or preferred contributor:</p> <ol style="list-style-type: none"> <li>1. Compare the protocol's <b>preference</b> of the contributing routes. The lower the preference, the better the route. This is similar to the comparison that is done while determining the best route for the routing table.</li> <li>2. Compare the protocol's <b>preference2</b> of the contributing routes. The lower <b>preference2</b> value is better. If only one route has <b>preference2</b>, then this route is preferred.</li> <li>3. The preference values are the same. Proceed with a numerical comparison of the prefixes' values. <ol style="list-style-type: none"> <li>a. The primary contributor is the numerically smallest prefix value.</li> <li>b. If the two prefixes are numerically equal, the primary contributor is the route that has the smallest prefix length value.</li> </ol> </li> </ol> |



At this point, the two routes are the same. The primary contributor does not change. An additional next hop is available for the existing primary contributor.

A rejected contributor still can contribute to less specific generated route. If you do not specify a policy filter, all candidate routes contribute to a generated route.

|                                 |                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <i>policy-name</i> —Name of a routing policy.                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> <li>• <a href="#">aggregate on page 3430</a></li> <li>• <a href="#">generate on page 3465</a></li> </ul> |

## policy (Flow Maps)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">flow-map</a> <i>flow-map-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">flow-map</a> <i>flow-map-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">flow-map</a> <i>flow-map-name</i> ],<br>[edit routing-options multicast <a href="#">flow-map</a> <i>flow-map-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure a flow map policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies for flow mapping.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## policy (SSM Maps)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>],</code><br><code>[edit routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Apply one or more policies to an SSM map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies for SSM mapping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To view this statement in the configuration.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring SSM Mapping</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## ppm

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ppm {     no-delegate-processing; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.4.</p> <p>Statement introduced in Junos OS Release 10.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>(M120, M320, MX Series, T Series, TX Matrix routers, M7i and M10i routers with Enhanced CFEB [CFEB-E], EX Series switches, and QFX Series only) Disable distributed periodic packet management (PPM) to the Packet Forwarding Engine (on routers), to access ports (on EX3200 and EX4200 switches, and QFX Series), or to line cards (on EX6200 and EX8200 switches).</p> <p>After you disable PPM, PPM processing continues to run on the Routing Engine.</p> <p>In Junos OS Release 8.2, PPM was moved from the Routing Engine to the Packet Forwarding Engine, access ports, or line cards. The <b>no-delegate-processing</b> statement disables the default behavior and restores the legacy behavior.</p> |
| <b>Default</b>                  | Distributed PPM processing is enabled for all protocols that use PPM.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>no-delegate-processing</b> —Disable PPM to the Packet Forwarding Engine, access ports, or line cards. Distributed PPM is enabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure) on page 3424</a></li> <li>• <i>Ensuring That Distributed ppm Is Not Disabled</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## ppm (Ethernet Switching)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ppm {<br><a href="#">centralized</a> ;<br>}                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols lacp]                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.1 for T Series devices.                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure PPM processing options for Link Aggregation Control Protocol (LACP) packets.</p> <p>This command configures the PPM processing options for LACP packets only. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the <b>no-delegate-processing</b> configuration statement in the <b>[edit routing-options ppm]</b> statement hierarchy.</p> |
| <b>Default</b>                  | Distributed PPM processing is enabled for all packets that use PPM.                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure) on page 3424</a></li><li>• <i>Configuring Distributed Periodic Packet Management</i></li></ul>                                                                                                                                                                                                   |

## preference (Routing Options)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>(preference   preference2   color   color2) preference &lt;type type&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Preference value for a static, aggregate, or generated route. You also can specify a secondary preference value (<b>preference2</b>), as well as colors, which are even finer-grained preference values (<b>color</b> and <b>color2</b>).</p> <p>If the Junos OS routing table contains a dynamic route to a destination that has a better (lower) preference value than the static, aggregate, or generated route, the dynamic route is chosen as the active route and is installed in the forwarding table.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>preference</b>—Preference value. A lower number indicates a more preferred route.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 5 (for static routes), 130 (for aggregate and generated routes)</p> <p><b>type type</b>—(Optional) Type of route.</p> <p><b>Range:</b> 1 through 16</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> <li>• <a href="#">aggregate on page 3430</a></li> <li>• <a href="#">generate on page 3465</a></li> <li>• <a href="#">static on page 3522</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## prefix

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>prefix destination-prefix;</code>                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>scope</b> <i>scope-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>scope</b> <i>scope-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>scope</b> <i>scope-name</i> ],<br>[edit routing-options multicast <b>scope</b> <i>scope-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                             |
| <b>Description</b>              | Configure the prefix for multicast scopes.                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b>destination-prefix</b> —Address range for the multicast scope.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Examples: Configuring Administrative Scoping</i></li><li>• <i>Example: Creating a Named Scope for Multicast Scoping</i></li><li>• <i>multicast</i></li></ul>                                                                                                                                                                                                                                                             |

## qualified-next-hop (Static Routes)

**Syntax** `qualified-next-hop (address | interface-name) {  
     bfd-liveness-detection {  
         authentication {  
             algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 | meticulous-keyed-sha-1 |  
                 simple-password);  
             key-chain key-chain-name;  
             loose-check;  
         }  
         detection-time {  
             threshold milliseconds;  
         }  
         holddown-interval milliseconds;  
         minimum-interval milliseconds;  
         minimum-receive-interval milliseconds;  
         multiplier number;  
         no-adaptation;  
         transmit-interval {  
             minimum-interval milliseconds;  
             threshold milliseconds;  
         }  
         version (1 | automatic);  
     }  
     interface interface-name;  
     metric metric;  
     preference preference;  
}`

**Hierarchy Level** `[edit logical-systems logical-system-name routing-instances routing-instance-name  
     routing-options static route destination-prefix],  
     [edit logical-systems logical-system-name routing-options rib inet6.0 static route  
         destination-prefix],  
     [edit logical-systems logical-system-name routing-options static route destination-prefix],  
     [edit routing-instances routing-instance-name routing-options static route destination-prefix],  
     [edit routing-options rib inet6.0 static route destination-prefix],  
     [edit routing-options static route destination-prefix]`

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Statement introduced in Junos OS Release 11.3 for the QFX Series.  
 Statement introduced in Junos OS Release 12.3 for ACX Series routers.

**Description** Configure a static route with multiple possible next hops, each of which can have its own preference value, IGP metric that is used when the route is exported into an IGP, and Bidirectional Forwarding Detection (BFD) settings. If multiple links are operational, the one with the most preferred next hop is used. The most preferred next hop is the one with the lowest preference value.

**Options** *address*—IPv4, IPv6, or ISO network address of the next hop.  
*interface-name*—Name of the interface on which to configure an independent metric or preference for a static route. To configure an unnumbered interface as the next-hop

interface for a static route, specify **qualified-next-hop *interface-name***, where *interface-name* is the name of the IPv4 or IPv6 unnumbered interface.



**NOTE:** For an Ethernet interface to be configured as the qualified next hop for a static route, it must be an unnumbered interface.

To configure an Ethernet interface as an unnumbered interface, configure the `unnumbered-address <interface-name>` statement at the [edit interfaces <interface-name> unit <logical-unit-number> family <family-name>] hierarchy level as described in *Configuring an Unnumbered Interface*.

The remaining statements are explained separately.

|                                 |                                                                                |
|---------------------------------|--------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.                           |
|                                 | routing-control—To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | • <i>Example: Configuring Static Route Preferences and Qualified Next Hops</i> |
|                                 | • <i>Example: Enabling BFD on Qualified Next Hops in Static Routes</i>         |



## readvertise


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (readvertise   no-readvertise);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>static</b> (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>static</b> (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>static</b> (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit routing-options <b>static</b> (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure whether static routes are eligible to be readvertised by routing protocols:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>                  | Static routes are eligible to be readvertised (that is, exported from the routing table into dynamic routing protocols) if a policy to do so is configured. To mark an IPv4 static route as being ineligible for readvertisement, include the <b>no-readvertise</b> statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>readvertise</b>—Readvertise static routes. Include the <b>readvertise</b> statement when configuring an individual route in the <b>route</b> portion of the <b>static</b> statement to override a <b>no-readvertise</b> option specified in the <b>defaults</b> portion of the statement.</p> <p><b>no-readvertise</b>—Mark a static route as being ineligible for readvertisement. Include the <b>no-readvertise</b> option when configuring the route.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Controlling Static Routes in Routing and Forwarding Tables</i></li> <li>• <a href="#">static on page 3522</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## redundant-sources

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>redundant-sources [ <i>addresses</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i>],</code><br><code>[edit routing-options multicast <b>flow-map</b> <i>flow-map-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure a list of redundant sources for multicast flows defined by a flow map.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <b><i>addresses</i></b> —List of IPv4 or IPv6 addresses for use as redundant (backup) sources for multicast flows defined by a flow map.                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring a Multicast Flow Map</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## resolution

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> resolution {     rib routing-table-name {         import [ policy-names ];         resolution-ribs [ routing-table-names ];     } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure the router to perform custom route resolution on protocol next hops of routes in a certain routing table. The protocol next hop is used to determine the forwarding next-hop.</p> <p>For example, you might want to direct <b>inet.2</b> route resolution to use topology routing tables <b>:red.inet.0</b> and <b>:blue.inet.0</b> for protocol next-hop IP address lookups. Or you might want to direct <b>bgp.l3vpn.0</b> to use the information in <b>inet.0</b> to resolve routes, thus overriding the default behavior, which is to use <b>inet.3</b>.</p> <p>You can specify up to two routing tables in the <b>resolution-ribs</b> statement. The route resolution scheme first checks the first-listed routing table for the protocol next-hop address. If the address is found, it uses this entry. If it is not found, the resolution scheme checks second-listed routing table. Hence, only one routing table is used for each protocol nexthop address. For example, if you configure <b>resolution rib bgp.l3vpn.0 resolution-ribs [inet.0 inet.3]</b>, <b>inet.0</b> is checked first and then <b>inet.3</b> is checked.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 20px;"> <p> <b>NOTE:</b> Customizing route resolution might cause the routing protocol process (rpd) to consume more memory resources than it ordinarily would. When you customize route resolution, we recommend that you check the memory resources by running the <b>show system processes</b> and the <b>show task memory</b> commands. For more information, see <i>Routing Protocol Process Memory FAQs</i>.</p> </div> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- Related Documentation**
- *Example: Configuring Route Resolution on PE Routers*
  - *Example: Configuring Route Resolution on Route Reflectors*
  - *Example: Configuring Multitopology Routing Based on a Multicast Source*

## resolution-ribs

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>resolution-ribs [ <i>routing-table-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution <a href="#">rib</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options resolution <a href="#">rib</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options resolution <a href="#">rib</a>],</code><br><code>[edit routing-options resolution <a href="#">rib</a>]</code>                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Specify one or more routing tables to use for route resolution.</p> <p>This statement enables you to override the default routing tables that Junos OS uses for route resolution. For example, suppose that the resolution routing table is <b>inet.3</b>, but you want to allow fallback resolution through <b>inet.0</b>. One example use case is overriding the <b>bgp.rtarget.0 (family route-target)</b> routing table resolution from using only <b>inet.3</b> to using both <b>inet.3</b> and <b>inet.0</b>.</p> |
| <b>Options</b>                  | <b><i>routing-table-names</i></b> —Name of one or more routing tables.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Route Resolution on PE Routers</i></li><li>• <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i></li></ul>                                                                                                                                                                                                                                                                                                                       |

## resolve

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | resolve;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>static</b> (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>static</b> (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>static</b> (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit routing-options <b>static</b> (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Statically configure routes to be resolved to a next hop that is not directly connected. The route is resolved through the <b>inet.0</b> and <b>inet.3</b> routing tables.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>                  | Static routes can point only to a directly connected next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">static on page 3522</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## restart-duration (Routing Options)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>restart-duration seconds;</code>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <b>routing-options</b> graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> <b>routing-options</b> graceful-restart],<br>[edit routing-instances <i>routing-instance-name</i> <b>routing-options</b> graceful-restart],<br>[edit <b>routing-options</b> graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure the restart timer for graceful restart.                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>seconds</b> —Configure the time period for the restart to last.<br><b>Range:</b> 120 through 900 seconds<br><b>Default:</b> 300 seconds                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Junos OS High Availability Library for Routing Devices</i></li></ul>                                                                                                                                                                                                                                                                             |

## retain

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (no-retain   retain);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>static</b> (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>static</b> (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>static</b> (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> <b>static</b> (defaults   route)],</p> <p>[edit routing-options <b>static</b> (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure statically configured routes to be deleted from or retained in the forwarding table when the routing protocol process shuts down normally:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | Statically configured routes are deleted from the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>no-retain</b>—Delete statically configured routes from the forwarding table when the routing protocol process shuts down normally. To explicitly specify that routes be deleted from the forwarding table, include the <b>no-retain</b> statement. Include this statement when configuring an individual route in the <b>route</b> portion of the <b>static</b> statement to override a <b>retain</b> option specified in the <b>defaults</b> portion of the statement.</p> <p><b>retain</b>—Have a static route remain in the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.</p>                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <a href="#">static on page 3522</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## reverse-oif-mapping

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>reverse-oif-mapping {<br/>    no-qos-adjust;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-options multicast <a href="#">interface</a> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>The <b>no-qos-adjust</b> statement added in Junos OS Release 9.5.</p> <p>The <b>no-qos-adjust</b> statement introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                |
| <b>Description</b>              | <p>Enable the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD request it receives over the multicast VLAN.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                          |



## rpf-check-policy (Routing Options RPF)

|                                 |                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rpf-check-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                 |
| <b>Description</b>              | Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more multicast RPF check policies.                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring RPF Policies</i></li> </ul>                                                                                                                                                                                                                                                           |

## rib (General)

---

```
Syntax  rib routing-table-name {  
        aggregate {  
            defaults {  
                ... aggregate-options ...  
            }  
            route destination-prefix {  
                policy policy-name;  
                ... aggregate-options ...  
            }  
        }  
        generate {  
            defaults {  
                generate-options;  
            }  
            route destination-prefix {  
                policy policy-name;  
                generate-options;  
            }  
        }  
        martians {  
            destination-prefix match-type <allow>;  
        }  
    }  
    static {  
        defaults {  
            static-options;  
        }  
        rib-group group-name;  
        route destination-prefix {  
            next-hop;  
            static-options;  
        }  
    }  
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],  
[edit logical-systems *logical-system-name* routing-options],  
[edit routing-instances *routing-instance-name* routing-options],  
[edit routing-options]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Create a routing table.

Explicitly creating a routing table with ***routing-table-name*** is optional if you are not adding any static, martian, aggregate, or generated routes to the routing table and if you also are creating a routing table group.



**NOTE:** The IPv4 multicast routing table (`inet.1`) and the IPv6 multicast routing table (`inet6.1`) are not supported for this statement.

**Default** If you do not specify a routing table name with the *routing-table-name* option, the software uses the default routing tables, which are `inet.0` for unicast routes and `inet.1` for the multicast cache.

**Options** *routing-table-name*—Name of the routing table, in the following format:  
*protocol [.identifier]*.

In a routing instance, the routing table name must include the routing instance name.

For example, if the routing instance name is `link0`, the routing table name might be `link0.inet6.0`.

- *protocol* is the protocol family. It can be `inet6` for the IPv6 family, `inet` for the IPv4 family, `iso` for the ISO protocol family, or *instance-name.iso.0* for an ISO routing instance.
- *identifier* is a positive integer that specifies the instance of the routing table.

**Default:** `inet.0`

The remaining statements are explained separately.

**Required Privilege Level** `routing`—To view this statement in the configuration.  
`routing-control`—To add this statement to the configuration.

**Related Documentation**

- *Example: Creating Routing Tables*
- *passive*

## rib (Route Resolution)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rib <i>routing-table-name</i> {<br/>    <b>import</b> [ <i>policy-names</i> ];<br/>    <b>resolution-ribs</b> [ <i>routing-table-names</i> ];<br/>}</pre>                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>resolution</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options <b>resolution</b> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options <b>resolution</b> ],<br>[edit routing-options <b>resolution</b> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                          |
| <b>Description</b>              | Specify a routing table name for route resolution.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Route Resolution on PE Routers</i></li></ul>                                                                                                                                                                                                                                                            |

## rib-group (Routing Options)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rib-group group-name;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <a href="#">interface-routes</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <a href="#">interface-routes</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options static],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <a href="#">interface-routes</a>],</p> <p>[edit routing-options <a href="#">interface-routes</a>],</p> <p>[edit routing-options rib <i>routing-table-name</i> static],</p> <p>[edit routing-options static]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Configure which routing table groups interface routes are imported into.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b><i>group-name</i></b>—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. It generally does not make sense to specify more than a single routing table group.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Importing Direct and Static Routes Into a Routing Instance</i></li> <li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li> <li>• <a href="#">interface-routes on page 3474</a></li> <li>• <a href="#">rib-groups on page 3510</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                            |

## rib-groups

---

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax              | <pre>rib-groups {<br/>    group-name {<br/>        export-rib group-name;<br/>        import-policy [ policy-names ];<br/>        import-rib [ group-names ];<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Hierarchy Level     | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Release Information | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Description         | <p>Group one or more routing tables to form a routing table group. A routing protocol can import routes into all the routing tables in the group and can export routes from a single routing table.</p> <p>Each routing table group must contain one or more routing tables that Junos OS uses when importing routes (specified in the <b>import-rib</b> statement) and optionally can contain one routing table group that Junos OS uses when exporting routes to the routing protocols (specified in the <b>export-rib</b> statement).</p> <p>The first routing table you specify is the <i>primary routing table</i>, and any additional routing tables are the <i>secondary routing tables</i>.</p> <p>The primary routing table determines the address family of the routing table group. To configure an IP version 4 (IPv4) routing table group, specify <b>inet.0</b> as the primary routing table. To configure an IP version 6 (IPv6) routing table group, specify <b>inet6.0</b> as the primary routing table. If you configure an IPv6 routing table group, the primary and all secondary routing tables must be IPv6 routing tables (<b>inet6.x</b>).</p> <p>In Junos OS Release 9.5 and later, you can include both IPv4 and IPv6 routing tables in an IPv4 import routing table group using the <b>import-rib</b> statement. In releases prior to Junos OS Release 9.5, you can only include either IPv4 or IPv6 routing tables in the same <b>import-rib</b> statement. The ability to configure an import routing table group with both IPv4 and IPv6 routing tables enables you, for example, to populate the <b>inet6.3</b> routing table with IPv6 addresses that are compatible with IPv4. Specify <b>inet.0</b> as the primary routing table, and specify <b>inet6.3</b> as a secondary routing table.</p> |



**NOTE:** On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

---



**NOTE:** If you configure an import routing table group that includes both IPv4 and IPv6 routing tables, any corresponding export routing table group must include only IPv4 routing tables.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group. For more information, see *Example: Configuring Multiple Routing Instances of OSPF*.

After specifying the routing table from which to import routes, you can apply one or more policies to control which routes are installed in the routing table group. To apply a policy to routes being imported into the routing table group, include the **import-policy** statement.

**Options** *group-name*—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens.

The remaining statements are explained separately.


**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- [rib-group on page 3509](#)

## route-distinguisher-id

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>route-distinguisher-id <i>ip-address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Automatically assign a route distinguisher to the routing instance.</p> <p>If you configure the <b>route-distinguisher</b> statement in addition to the <b>route-distinguisher-id</b> statement, the value configured for <b>route-distinguisher</b> supersedes the value generated from <b>route-distinguisher-id</b>.</p> <div><p><b>NOTE:</b> To avoid a conflict in the two route distinguisher values, it is recommended to ensure that the first half of the route distinguisher obtained by configuring the <b>route-distinguisher</b> statement is different from the first half of the route distinguisher obtained by configuring the <b>route-distinguisher-id</b> statement.</p></div> |
| <b>Options</b>                  | <i>ip-address</i> —Address for routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring BGP Route Target Filtering for VPNs</i></li><li>• <i>Configuring Routing Instances on PE Routers in VPNs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



---

## route-record

---

|                                 |                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | route-record;                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-options]                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                    |
| <b>Description</b>              | Export the AS path and routing information to the traffic sampling process.<br><br>Before you can perform flow aggregation, the routing protocol process must export the AS path and routing information to the sampling process. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling Flow Aggregation</i></li><li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li></ul>                                                                      |

## router-id

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>router-id address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>     | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options],</code><br><code>[edit routing-options]</code>                                                                                                                         |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                 |
| <b>Description</b>         | <p>Specify the routing device's IP address.</p> <p>The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier usually is the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.</p> |



**NOTE:** We strongly recommend that you configure the router identifier under the `[edit routing-options]` hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

For more information about the router identifier in OSPF, see *Example: Configuring an OSPF Router Identifier*.

You must configure a router-id in order for BGP and OSPF to function in a routing instance. Use the **show route instance detail** command to display the router-id value for a routing instance. If the router-id is **0.0.0.0**, then the routing instance has no router-id.

For more information about the router identifier in OSPF, see *Example: Configuring an OSPF Router Identifier*.



**NOTE:** If you run OSPF for IPv6 or BGP for IPv6 in a routing instance, you must configure an IPv4 router identifier (**router-id**) in the routing instance itself. In other words, the IPv4 **router-id** in the main routing instance is not inherited by other routing instances. Even if you run *only* IPv6 OSPF or BGP in a routing instance, the IPv4 **router-id** must be configured because OSPF and BGP, even when used exclusively with IPv6, use the IPv4 **router-id** for handshaking. If you do not configure the IPv4 **router-id** in the IPv6 OSPF or BGP routing instance, then the IPv6 protocols will use invalid IPv4 address 0.0.0.0 and the adjacencies and connections will fail.

|                                 |                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <b>address</b> —IP address of the routing device.<br><b>Default:</b> Address of the first interface encountered by Junos OS                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring External BGP Peering</i></li> <li>• <i>Examples: Configuring Internal BGP Peering</i></li> </ul> |

## routing-options

|                                 |                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | routing-options { ... }                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit],<br>[edit logical-systems <i>logical-system-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                      |
| <b>Description</b>              | Configure protocol-independent routing properties.                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Protocol-Independent Routing Properties Feature Guide for Routing Devices</i></li> </ul>                                                                                           |

## scope

---

|                                 |                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>scope scope-name {<br/>    interface [ <i>interface-names</i> ];<br/>    <b>prefix</b> <i>destination-prefix</i>;<br/>}</pre>                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                             |
| <b>Description</b>              | Configure multicast scoping.                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b><i>scope-name</i></b> —Name of the multicast scope.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Creating a Named Scope for Multicast Scoping</i></li></ul>                                                                                                                                                                                                                      |

## scope-policy

**Syntax** `scope-policy [ policy-names ];`

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-options multicast],  
[edit routing-options multicast]



**NOTE:** You can configure a scope policy at these two hierarchy levels only. You cannot apply a scope policy to a specific routing instance, because all scoping policies are applied to all routing instances. However, you can apply the `scope` statement to a specific routing instance at the [edit routing-instances *routing-instance-name* routing-options multicast] or [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast] hierarchy level.

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

**Description** Apply policies for scoping. The policy must be correctly configured at the **edit policy-options policy-statement** hierarchy level.

**Options** *policy-names*—Name of one or more multicast scope policies.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**


- [scope on page 3516](#)
- *Example: Using a Scope Policy for Multicast Scoping*

## source (Source-Specific Multicast)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source [ <i>addresses</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>],</code><br><code>[edit routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Specify IPv4 or IPv6 source addresses for an SSM map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b><i>addresses</i></b> —IPv4 or IPv6 source addresses.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To view this statement in the configuration.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring SSM Mapping</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## source-routing

|                                 |                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | source-routing {<br>(ip   ipv6)<br>}                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit routing-options]                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement for IPv6 introduced in Junos OS Release 8.2.<br>Statement for IPv4 introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                               |
| <b>Description</b>              | <p>Enable source routing.</p> <p>Source routing allows a sender of a packet to partially or completely specify the route the packet takes through the network. In contrast, in non-source routing protocols, routers in the network determine the path based on the packet's destination.</p>   |
|                                 | <div>  <p><b>NOTE:</b> We recommend that you not use source routing. Instead, we recommend that you use policy-based routing or filter-based forwarding to route packets based on source addresses.</p> </div> |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring Filter-Based Forwarding on the Source Address</i></li> </ul>                                                                                                                                                                     |

## ssm-groups

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-groups [ <i>ip-addresses</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</code><br><code>[edit routing-options multicast]</code>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure source-specific multicast (SSM) groups.</p> <p>By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the <b>ssm-groups</b> statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the <b>ssm-groups</b> statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.</p> <p>IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.</p> |
| <b>Options</b>                  | <i>ip-addresses</i> —List of one or more additional SSM group addresses separated by a space.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Source-Specific Multicast Groups with Any-Source Override</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



## ssm-map (Routing Options Multicast)

|                                 |                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ssm-map <i>ssm-map-name</i> {     policy [ <i>policy-names</i> ];     source [ <i>addresses</i> ]; }</pre>                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                 |
| <b>Description</b>              | Configure SSM mapping.                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b><i>ssm-map-name</i></b>—Name of the SSM map.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring SSM Mapping</i></li> </ul>                                                                                                                                                                                                                                                            |

## static (Routing Options)

---

```
Syntax static {
    defaults {
        static-options;
    }
    rib-group group-name;
    route destination-prefix {
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            local-address ip-address;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-receive-ttl number;
            multiplier number;
            neighbor address;
            no-adaptation;
            transmit-interval {
                threshold milliseconds;
                minimum-interval milliseconds;
            }
            version (1 | automatic);
        }
        next-hop address;
        next-hop options;
        qualified-next-hop address {
            bfd-liveness-detection {
                authentication {
                    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                        meticulous-keyed-sha-1 | simple-password);
                    key-chain key-chain-name;
                    loose-check;
                }
                detection-time {
                    threshold milliseconds;
                }
                holddown-interval milliseconds;
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                version (1 | automatic);
            }
        }
    }
}
```

```

        metric metric;
        preference preference;
    }
    static-options;
}

```

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options],</p> <p>[edit routing-options rib <i>routing-table-name</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos 9.6.</p> <p>Support for BFD authentication introduced in Junos 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                              |
| <b>Description</b>         | <p>Configure static routes to be installed in the routing table. You can specify any number of routes within a single <b>static</b> statement, and you can specify any number of <b>static</b> options in the configuration.</p>                                                                                                                                                                                                                                                    |

**Options defaults**—(Optional) Specify global static route options. These options only set default attributes inherited by all newly created static routes. These are treated as global defaults and apply to all the static routes you configure in the **static** statement.



**NOTE:** Specifying the global static route options does not create default routes. These options only set default attributes inherited by all newly created static routes.

**route**—Configure individual static routes. In this part of the **static** statement, you optionally can configure static route options. These options apply to the individual destination only and override any options you configured in the **defaults** part of the **static** statement.

- **destination-prefix/prefix-length—destination-prefix** is the network portion of the IP address, and **prefix-length** is the destination prefix length.

When you configure an individual static route in the **route** part of the **static** statement, specify the destination of the route (in **route destination-prefix**) in one of the following ways:

- **network/mask-length**, where **network** is the network portion of the IP address and **mask-length** is the destination prefix length.
- **default** if this is the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.



**NOTE:** IPv4 packets with a destination of 0.0.0.0 (the obsoleted limited broadcast address) and IPv6 packets with a destination of 0::0 are discarded by default. To forward traffic destined to these addresses, you can add a static route to 0.0.0.0/32 for IPv4 or 0::0/128 for IPv6.

- **nsap-prefix—nsap-prefix** is the network service access point (NSAP) address for ISO.
- **next-hop address**—Reach the next-hop routing device by specifying an IP address, an interface name, or an ISO network entity title (NET).

IPv4 or IPv6 address of the next hop to the destination, specified as:

- IPv4 or IPv6 address of the next hop
- Interface name (for point-to-point interfaces only)
- **address** or **interface-name** to specify an IP address of a multipoint interface or an interface name of a point-to-point interface.



**NOTE:** If an interface becomes unavailable, all configured static routes on that interface are withdrawn from the routing table.



**NOTE:** Load balancing is not supported on management and internal Ethernet (fxo) interfaces because this type of interface cannot handle the routing process. On fxp interfaces, you cannot configure multiple next hops and enable load balancing.

**next-hop options**—Additional information for how to manage forwarding of packets to the next hop.

- **discard**—Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP (or ICMPv6) unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.
- **iso-net**—Reach the next-hop routing device by specifying an ISO NSAP.

- **next-table *routing-table-name***—Name of the next routing table to the destination.

If you use the **next-table** action, the configuration must include a term qualifier that specifies a different table than the one specified in the **next-table** action. In other words, the term qualifier in the **from** statement must exclude the table in the **next-table** action. In the following example, the first term contains **rib vrf-customer2.inet.0** as a matching condition. The action specifies a next-hop in a different routing table, **vrf-customer1.inet.0**. The second term does the opposite by using **rib vrf-customer1.inet.0** in the match condition and **vrf-customer2.inet.0** in the **next-table** action.

```
term 1 {
  from {
    protocol bgp;
    rib vrf-customer2.inet.0;
    community customer;
  }
  then {
    next-hop next-table vrf-customer1.inet.0;
  }
}
term 2 {
  from {
    protocol bgp;
    rib vrf-customer1.inet.0;
    community customer;
  }
  then {
    next-hop next-table vrf-customer2.inet.0;
  }
}
```



**NOTE:** Within a routing instance, you cannot configure a static route with the **next-table inet.0** statement if any static route in the main routing instance is already configured with the **next-table** statement to point to the **inet.0** routing table of the routing instance. For example, if you configure on the main routing instance a static route 192.168.88.88/32 with the **next-table test.inet.0** statement and the routing instance **test** is also configured with a static route 192.168.88.88/32 with the **next-table inet.0** statement, the commit operation fails. Instead, you must configure a routing table group both on the main instance and on the routing instance, which enables you to install the static route into both routing tables.

- **receive**—Install a route for this next-hop destination into the routing table.

The **receive** option forces the packet to be sent to the Routing Engine.

The **receive** option can be useful in the following cases:

- For receiving MPLS packets destined to a VRF instance's loopback address
- For receiving packets on a link's subnet address, with zeros in the host portion of the address
- **reject**—Do not forward packets addressed to this destination. Instead, drop the packets, send ICMP (or ICMPv6) unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

**static-options**—(Optional under **route**) Additional information about static routes, which is included with the route when it is installed in the routing table.

You can specify one or more of the following in **static-options**. Each of the options is explained separately.

- (**active** | **passive**);
- **as-path** <as-path> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator as-number in-address>;
- **community** [ *community-ids* ];
- (**install** | **no-install**);
- (**metric** | **metric2** | **metric3** | **metric4**) *value* <type type>;
- (**preference** | **preference2** | **color** | **color2**) *preference* <type type>;
- (**readvertise** | **no-readvertise**);
- (**resolve** | **no-resolve**);
- (**retain** | **no-retain**);
- **tag** *metric type number*;

The remaining statements are explained separately.

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> </ul>                      |

## subscriber-leave-timer

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>subscriber-leave-timer seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface interface-name</a> ],<br>[edit routing-options multicast <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2.<br>Statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                         |
| <b>Description</b>              | Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message.                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>seconds</b> —Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message. Specifying a value of 0 results in an immediate update. This is the same as if the statement were not configured.<br><b>Range:</b> 0 through 30<br><b>Default:</b> 0 seconds                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                               |



## tag (Routing Options)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>tag metric type number;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   <a href="#">static</a>) (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Associate a tag with a static, aggregate, or generated route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>                  | No tag strings are associated with routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><i>metric</i>—Tag metric.</p> <p><b>Range:</b> 0 through 4,294,967,295</p> <p><i>type number</i>—Tag type.</p> <p><b>Range:</b> 1 through 16</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> <li>• <a href="#">aggregate on page 3430</a></li> <li>• <a href="#">generate on page 3465</a></li> <li>• <a href="#">static on page 3522</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## threshold (Multicast Forwarding Cache)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>threshold {<br/>    log-warning <i>value</i>;<br/>    suppress <i>value</i> &lt;reuse <i>value</i>&gt;;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Hierarchy Level          | <pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i><br/>    routing-options multicast forwarding-cache],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i><br/>    routing-options multicast forwarding-cache family (inet   inet6)],<br/>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache],<br/>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache<br/>    family (inet   inet6)],<br/>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],<br/>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache<br/>    (inet   inet6)],<br/>[edit routing-options multicast forwarding-cache],<br/>[edit routing-options multicast forwarding-cache family (inet   inet6)]</pre> |
| Release Information      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Description              | <p>Configure the global suppression, reuse, and warning log message thresholds for multicast forwarding cache limits. You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the <b>threshold</b> statement globally for the multicast forwarding cache or including the <b>family</b> statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p> <p>To confirm the configured threshold values, use the <b>show multicast forwarding-cache statistics</b> command.</p>                                                                                                                                                                                                                                                                                                                  |
| Options                  | <p><b>reuse <i>value</i></b>—(Optional) Value at which to begin creating new multicast forwarding cache entries. If configured, this number should be less than the <b>suppress</b> value.</p> <p><b>Range:</b> 1 through 200,000</p> <p><b>suppress <i>value</i></b>—Value at which to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number should be greater than the <b>reuse</b> value.</p> <p><b>Range:</b> 1 through 200,000</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                               |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Related Documentation    | <ul style="list-style-type: none"><li>• <i>Examples: Configuring the Multicast Forwarding Cache</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## timeout (Flow Maps)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | timeout (never non-discard-entry-only   <i>minutes</i> );                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ]                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure the timeout value for multicast forwarding cache entries associated with the flow map.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>minutes</b> —Length of time that the forwarding cache entry remains active.<br><b>Range:</b> 1 through 720<br><br><b>never non-discard-entry-only</b> —Specify that the forwarding cache entry always remain active. If you omit the <b>non-discard-entry-only</b> option, all multicast forwarding entries, including those in forwarding and pruned states, are kept forever. If you include the <b>non-discard-entry-only</b> option, entries with forwarding states are kept forever, and entries with pruned states time out. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                   |

## timeout (Multicast)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>timeout <i>minutes</i> &lt;family (inet   inet6)&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">forwarding-cache</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">forwarding-cache</a> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">forwarding-cache</a> ],<br>[edit routing-options multicast <a href="#">forwarding-cache</a> ]                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure the timeout value for multicast forwarding cache entries.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b><i>minutes</i></b> —Length of time that the forwarding cache limit remains active.<br><b>Range:</b> 1 through 720<br><br><b><i>family (inet   inet6)</i></b> —(Optional) Apply the configured timeout to either IPv4 or IPv6 multicast forwarding cache entries. Configuring the <b>timeout</b> statement globally for the multicast forwarding cache or including the <b>family</b> statement to configure the timeout value for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.<br><br><b>Default:</b> By default, the configured timeout applies to both IPv4 and IPv6 multicast forwarding cache entries. |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring the Multicast Forwarding Cache</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## traceoptions (Routing Options)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>                                                                                                                                                                                                                                         |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>nsr-synchronization</b> flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p><b>nsr-synchronization</b> and <b>nsr-packet</b> flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>nsr-synchronization</b> flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p><b>nsr-synchronization</b> flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p><b>nsr-synchronization</b> flag for PIM added in Junos OS Release 9.3.</p> <p><b>nsr-synchronization</b> flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>nsr-synchronization</b> flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> |
| <b>Description</b>         | <p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>             | If you do not include this statement, no global tracing operations are performed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>             | <p><b>Values:</b></p> <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place global routing protocol tracing output in the file <b>routing-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and</p>                                                                                                                                                                                                 |

so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization-detail**—(MPLS only) Tracing operations for nonstop active routing in detail
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Tracing Global Routing Protocol Operations*
- [Tracing Nonstop Active Routing Synchronization Events on page 2573](#)

## upstream-interface

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>upstream-interface [ <i>interface-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>pim-to-igmp-proxy</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>pim-to-mld-proxy</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>pim-to-igmp-proxy</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>pim-to-mld-proxy</b>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>pim-to-igmp-proxy</b>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>pim-to-mld-proxy</b>],</code><br><code>[edit routing-options multicast <b>pim-to-igmp-proxy</b>],</code><br><code>[edit routing-options multicast <b>pim-to-mld-proxy</b>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure at least one, but not more than two, upstream interfaces on the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain. The RP routing device translates PIM join or prune messages into corresponding IGMP report or leave messages (if you include the <b>pim-to-igmp-proxy</b> statement), or into corresponding MLD report or leave messages (if you include the <b>pim-to-mld-proxy</b> statement). The routing device then proxies the IGMP or MLD report or leave messages to one or both upstream interfaces to forward IPv4 multicast traffic (for IGMP) or IPv6 multicast traffic (for MLD) across the PIM domains.                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b><i>interface-names</i></b> —Names of one or two upstream interfaces to which the RP routing device proxies IGMP or MLD report or leave messages for transmission of multicast traffic across PIM domains. You can specify a maximum of two upstream interfaces on the RP routing device. To configure a set of two upstream interfaces, specify the full interface names, including all physical and logical address components, within square brackets ( [ ] ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM-to-IGMP Message Translation</i></li><li>• <i>Configuring PIM-to-MLD Message Translation</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



# Administration

- [Routine Monitoring on page 3537](#)
- [Operational Commands on page 3539](#)

## Routine Monitoring

- [Monitoring Routing Information on page 3537](#)

### Monitoring Routing Information

#### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to view the **inet.0** routing table on the routing device.

#### Action

To view the routing tables in the J-Web interface, select **Monitor > Routing > Route Information**. Apply a filter or a combination of filters to view messages. You can use filters to display relevant events.

To view the routing table in the CLI, enter the following commands in the CLI interface:

- **show route terse**
- **show route detail**

#### Meaning

[Table 367 on page 3537](#) describes the different filters, their functions, and the associated actions.

[Table 368 on page 3538](#) summarizes key output fields in the routing information display.

**Table 367: Filtering Route Messages**

| Field               | Function                                                 | Your Action                    |
|---------------------|----------------------------------------------------------|--------------------------------|
| Destination Address | Specifies the destination address of the route.          | Enter the destination address. |
| Protocol            | Specifies the protocol from which the route was learned. | Enter the protocol name.       |

Table 367: Filtering Route Messages (*continued*)

| Field            | Function                                                                                                                             | Your Action                                                     |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|
| Next hop address | Specifies the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it. | Enter the next hop address.                                     |
| Receive protocol | Specifies the dynamic routing protocol using which the routing information was received through a particular neighbor.               | Enter the routing protocol.                                     |
| Best route       | Specifies only the best route available.                                                                                             | Select the view details of the best route.                      |
| Inactive routes  | Specifies the inactive routes.                                                                                                       | Select the view details of inactive routes.                     |
| Exact route      | Specifies the exact route.                                                                                                           | Select the view details of the exact route.                     |
| Hidden routes    | Specifies the hidden routes.                                                                                                         | Select the view details of hidden routes.                       |
| Search           | Applies the specified filter and displays the matching messages.                                                                     | To apply the filter and display messages, click <b>Search</b> . |

Table 368: Summary of Key Routing Information Output Fields

| Field                  | Values                                                                                                                          | Additional Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Static Route Addresses | The list of static route addresses.                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Protocol               | Protocol from which the route was learned: <b>Static</b> , <b>Direct</b> , <b>Local</b> , or the name of a particular protocol. |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Preference             | The preference is the individual preference value for the route.                                                                | The route preference is used as one of the route selection criteria.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Next-Hop               | Network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.          | <p>If a next hop is listed as <b>Discard</b>, all traffic with that destination address is discarded rather than routed. This value generally means that the route is a static route for which the <b>discard</b> attribute has been set.</p> <p>If a next hop is listed as <b>Reject</b>, all traffic with that destination address is rejected. This value generally means that the address is unreachable. For example, if the address is a configured interface address and the interface is unavailable, traffic bound for that address is rejected.</p> <p>If a next hop is listed as <b>Local</b>, the destination is an address on the host (either the loopback address or Ethernet management port 0 address, for example).</p> |
| Age                    | How long the route has been active.                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 368: Summary of Key Routing Information Output Fields (*continued*)

| Field   | Values                                                                                                                                                                                                                                                                     | Additional Information         |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
| State   | Flags for this route.                                                                                                                                                                                                                                                      | There are many possible flags. |
| AS Path | <p>AS path through which the route was learned. The letters of the AS path indicate the path origin:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>?</b>—Incomplete. Typically, the AS path was aggregated.</li> </ul> |                                |

- Related Documentation**
- [Configuring Static Routing \(J-Web Procedure\) on page 3422](#)
  - [Configuring Static Routing \(CLI Procedure\) on page 3422](#)
  - [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)

## Operational Commands

- `clear ipv6 neighbors`
- `show as-path`
- `show as-path domain`
- `show as-path summary`
- `show ipv6 neighbors`
- `show route`
- `show route active-path`
- `show route all`
- `show route aspath-regex`
- `show route best`
- `show route brief`
- `show route community`
- `show route community-name`
- `show route damping`
- `show route detail`
- `show route exact`
- `show route export`
- `show route extensive`
- `show route flow validation`
- `show route inactive-path`
- `show route inactive-prefix`

- [show route instance](#)
- [show route label](#)
- [show route label-switched-path](#)
- [show route martians](#)
- [show route next-hop](#)
- [show route no-community](#)
- [show route protocol](#)
- [show route range](#)
- [show route receive-protocol](#)
- [show route resolution](#)
- [show route snooping](#)
- [show route source-gateway](#)
- [show route summary](#)
- [show route table](#)
- [show route terse](#)

---

## clear ipv6 neighbors

---

|                                 |                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear ipv6 neighbors<br><all   host <i>hostname</i> >                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.3 for EX Series switches.<br>Command introduced in Junos OS Release 12.2 for the QFX Series.                                                |
| <b>Description</b>              | Clear IPv6 neighbor cache information.                                                                                                                                                                                                  |
| <b>Options</b>                  | <b>none</b> —Clear all IPv6 neighbor cache information.<br><br><b>all</b> —(Optional) Clear all IPv6 neighbor cache information.<br><br><b>host <i>hostname</i></b> —(Optional) Clear the information for the specified IPv6 neighbors. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show ipv6 neighbors on page 3551</a></li></ul>                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">clear ipv6 neighbors on page 3541</a>                                                                                                                                                                                       |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                   |

### Sample Output

#### clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
```

## show as-path

---

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax              | <a href="#">Syntax on page 3542</a><br><a href="#">Syntax (EX Series Switches) on page 3542</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Syntax                      | <code>show as-path</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Syntax (EX Series Switches) | <code>show as-path</code><br><code>&lt;brief   detail&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Release Information         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Description                 | <p>Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the <b>as-path-prepend</b> action) or through aggregation.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p> |
| Options                     | <p><b>none</b>—Display basic information about AS paths that the local routing device is using (same as brief).</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Required Privilege Level    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Related Documentation       | <ul style="list-style-type: none"><li>• <a href="#">show as-path summary on page 3549</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| List of Sample Output       | <a href="#">show as-path on page 3543</a><br><a href="#">show as-path detail on page 3544</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Output Fields               | <a href="#">Table 369 on page 3543</a> lists the output fields for the <b>show as-path</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 369: show as-path Output Fields

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output   |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| <b>Total AS paths</b> | Total number of AS paths.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>brief none</b> |
| <b>Bucket</b>         | Bucket number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels        |
| <b>Count</b>          | Number of AS path entries in this bucket.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels        |
| <b>AS path</b>        | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> <li>• <b>Atomic</b>—Route is an aggregate of several route prefixes.</li> <li>• <b>Aggregater</b>—Routing device has summarized a range of prefixes.</li> </ul> | All levels        |
| <b>domain</b>         | Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.                                                                                                                                                                                                                                                                                                                                              | <b>detail</b>     |
| <b>neighbor as</b>    | AS peer address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail</b>     |
| <b>length</b>         | Length of the AS path.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail</b>     |
| <b>segments</b>       | Length of the AS segment descriptor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail</b>     |
| <b>references</b>     | Path reference count.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>     |

## Sample Output

### show as-path

```

user@host> show as-path
Total AS paths: 30382
  Bucket 0      Count: 36
    I
    14203 2914 174 31752 I
    14203 2914 701 21512 I
    14203 2914 1239 26632 I
    14203 2914 1239 29704 I
    14203 2914 4323 10248 I
    14203 2914 4766 23560 I
    14203 2914 6395 32776 I
    14203 2914 7911 11272 I
    14203 2914 12180 18440 I
    14203 2914 17408 17416 I
    14203 2914 701 702 24586 I
    14203 2914 1239 4657 9226 I
    14203 2914 1239 7132 16394 I
    14203 2914 1299 8308 34826 I
    14203 2914 3320 5603 28682 I

```

```

14203 2914 3491 1680 33802 I
14203 2914 3549 7908 27658 I
14203 2914 3549 20804 30730 I
14203 2914 7018 2687 9226 I
14203 2914 174 9318 9318 23564 I
14203 2914 701 3786 3786 23564 I
14203 2914 701 4761 4795 9228 I
14203 2914 1239 7132 5673 18444 I
14203 2914 3491 20485 24588 24588 I
14203 2914 5511 2200 1945 2060 I
14203 2914 7911 14325 14325 14348 I
14203 2914 701 4637 9230 9230 9230 I
14203 2914 6395 14 14 14 14 I
14203 2914 9299 6163 6163 6163 9232 I
14203 2914 3356 3356 3356 3356 11955 21522 I
14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
14203 2914 174 30209 30222 30222 30222 ?
14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
Bucket 1    Count: 28
14203 2914 35847 I
14203 2914 174 19465 I
14203 2914 174 35849 I
14203 2914 2828 32777 I
14203 2914 4323 14345 I
14203 2914 4323 29705 I
14203 2914 6395 32777 I

```

...

## show as-path detail

```

user@host> show as-path detail
Total AS paths: 30410
Bucket 0    Count: 36
AS path: I
  domain 0, length 0, segments 0, references 54
AS path: 14203 2914 174 31752 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 701 21512 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 26632 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 29704 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4323 10248 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4766 23560 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 6395 32776 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 7911 11272 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 12180 18440 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 17408 17416 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 701 702 24586 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 1239 4657 9226 I

```



```

    domain 1, neighbor as: 14203, length 5, segments 1, references 7
AS path: 14203 2914 1239 7132 16394 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 1299 8308 34826 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3320 5603 28682 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3491 1680 33802 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 7908 27658 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 20804 30730 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 7018 2687 9226 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 174 9318 9318 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 3786 3786 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4761 4795 9228 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 14
AS path: 14203 2914 1239 7132 5673 18444 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 3491 20485 24588 24588 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 4
AS path: 14203 2914 5511 2200 1945 2060 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 7911 14325 14325 14348 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4637 9230 9230 9230 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 3
AS path: 14203 2914 6395 14 14 14 14 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 10
...

```

## show as-path domain

|                                    |                                                                                                                                                                                                                                                    |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3546</a><br><a href="#">Syntax (EX Series Switches) on page 3546</a>                                                                                                                                                    |
| <b>Syntax</b>                      | show as-path domain<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                        |
| <b>Syntax (EX Series Switches)</b> | show as-path domain                                                                                                                                                                                                                                |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                              |
| <b>Description</b>                 | Display autonomous system (AS) path domain information.                                                                                                                                                                                            |
| <b>Options</b>                     | <b>none</b> —(Optional) Display AS path domain information for all routing instances.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>       | <a href="#">show as-path domain on page 3548</a>                                                                                                                                                                                                   |
| <b>Output Fields</b>               | <a href="#">Table 370 on page 3546</a> lists the output fields for the <b>show as-path domain</b> command. Output fields are listed in the approximate order in which they appear                                                                  |

**Table 370: show as-path domain Output Fields**

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                               |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Domain</b>       | Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.                                                                                                                                              |
| <b>Primary</b>      | Primary AS number.                                                                                                                                                                                                                                                                                                                              |
| <b>References</b>   | Path reference count.                                                                                                                                                                                                                                                                                                                           |
| <b>Number Paths</b> | Number of known AS paths.                                                                                                                                                                                                                                                                                                                       |
| <b>Flags</b>        | Information about the AS path: <ul style="list-style-type: none"> <li>• <b>ASLoop</b>—Path contains an AS loop.</li> <li>• <b>Atomic</b>—Path includes the ATOMIC_AGGREGATE path attribute.</li> <li>• <b>Local</b>—Path was created by local aggregation.</li> <li>• <b>Master</b>—Path was created by the master routing instance.</li> </ul> |
| <b>Local AS</b>     | AS number of the local routing device.                                                                                                                                                                                                                                                                                                          |

Table 370: show as-path domain Output Fields (*continued*)

| Field Name   | Field Description                                       |
|--------------|---------------------------------------------------------|
| <b>Loops</b> | How many times this AS number can appear in an AS path. |

## Sample Output

show as-path domain

```
user@host> show as-path domain
Domain: 1          Primary: 10458
References:        3 Paths:      30383
Flags: Master
Local AS: 10458   Loops: 1
```

## show as-path summary

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3549</a><br><a href="#">Syntax (EX Series Switches) on page 3549</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                      | <pre>show as-path summary &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switches)</b> | show as-path summary                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                 | <p>Display autonomous system (AS) path summary information.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p> |
| <b>Options</b>                     | <p><b>none</b>—(Optional) Display AS path summary information for all routing instances.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li>• <a href="#">show as-path on page 3542</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>       | <a href="#">show as-path summary on page 3550</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>               | <p><a href="#">Table 371 on page 3549</a> lists the output fields for the <b>show as-path summary</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                        |

**Table 371: show as-path summary Output Fields**

| Field Name      | Field Description                             |
|-----------------|-----------------------------------------------|
| <b>AS Paths</b> | Number of AS paths.                           |
| <b>Buckets</b>  | Number of hash buckets in use.                |
| <b>Max</b>      | Maximum number of AS path entries per bucket. |
| <b>Min</b>      | Minimum number of AS path entries per bucket. |
| <b>Avg</b>      | Average number of AS path entries per bucket. |

Table 371: show as-path summary Output Fields (*continued*)

| Field Name    | Field Description                                 |
|---------------|---------------------------------------------------|
| Std deviation | Standard deviation of AS path entries per bucket. |

## Sample Output

show as-path summary

```
user@host> show as-path summary
AS Paths  Buckets  Max   Min   Avg   Std deviation
30425     1024     95    12    29    6.481419
```

## show ipv6 neighbors

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show ipv6 neighbors                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.3 for EX Series switches.<br>Command introduced in Junos OS Release 12.2 for the QFX Series. |
| <b>Description</b>              | Display information about the IPv6 neighbor cache.                                                                                                                                       |
| <b>Options</b>                  | This command has no options.                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear ipv6 neighbors on page 3541</a></li> </ul>                                                                                    |
| <b>List of Sample Output</b>    | <a href="#">show ipv6 neighbors on page 3551</a>                                                                                                                                         |
| <b>Output Fields</b>            | <a href="#">Table 372 on page 3551</a> describes the output fields for the <b>show ipv6 neighbors</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 372: show ipv6 neighbors Output Fields**

| Field Name        | Field Description                                                                                                          |
|-------------------|----------------------------------------------------------------------------------------------------------------------------|
| IPv6 Address      | Name of the IPv6 interface.                                                                                                |
| Linklayer Address | Link-layer address.                                                                                                        |
| State             | State of the link: <b>up</b> , <b>down</b> , <b>incomplete</b> , <b>reachable</b> , <b>stale</b> , or <b>unreachable</b> . |
| Exp               | Number of seconds until the entry expires.                                                                                 |
| Rtr               | Whether the neighbor is a routing device: <b>yes</b> or <b>no</b> .                                                        |
| Secure            | Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: <b>yes</b> or <b>no</b> .              |
| Interface         | Name of the interface.                                                                                                     |

## Sample Output

### show ipv6 neighbors

```

user@host> show ipv6 neighbors
IPv6 Address          Linklayer Address  State      Exp Rtr Secure
Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale      546 yes no
fe-1/2/0.1

```

|                                    |                   |       |     |     |    |
|------------------------------------|-------------------|-------|-----|-----|----|
| fe80::2a0:a514:0:24c<br>fe-1/2/0.1 | 00:05:85:8f:c8:bd | stale | 258 | yes | no |
| fe80::2a0:a514:0:64c<br>fe-1/2/1.5 | 00:05:85:8f:c8:bd | stale | 111 | yes | no |
| fe80::2a0:a514:0:a4c<br>fe-1/2/2.9 | 00:05:85:8f:c8:bd | stale | 327 | yes | no |



## show route

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3553</a><br><a href="#">Syntax (EX Series Switches) on page 3553</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                      | <pre>show route &lt;all&gt; &lt;destination-prefix&gt; &lt;logical-system (all   logical-system-name)&gt; &lt;private&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switches)</b> | <pre>show route &lt;all&gt; &lt;destination-prefix&gt; &lt;private&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>private</b> introduced in Junos OS Release 9.5.</p> <p>Option <b>private</b> introduced in Junos OS Release 9.5 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                 | Display the active entries in the routing tables.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                     | <p><b>none</b>—Display brief information about all active entries in the routing tables.</p> <p><b>all</b>—(Optional) Display information about all routing tables, including private, or internal, routing tables.</p> <p><b>destination-prefix</b>—(Optional) Display active entries for the specified address or range of addresses.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>private</b>—(Optional) Display information only about all private, or internal, routing tables.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li>• <i>Example: Configuring RIP</i></li> <li>• <i>Example: Configuring RIPng</i></li> <li>• <i>Example: Configuring IS-IS</i></li> <li>• <i>Examples: Configuring Internal BGP Peering</i></li> <li>• <i>Examples: Configuring External BGP Peering</i></li> <li>• <i>Examples: Configuring OSPF Routing Policy</i></li> </ul>                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>       | <a href="#">show route on page 3556</a><br><a href="#">show route on page 3557</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

[show route destination-prefix on page 3557](#)

[show route extensive on page 3557](#)

**Output Fields** [Table 373 on page 3554](#) describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

**Table 373: show route Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <i>number routes</i>       | <p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active).</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly.</li> </ul> <p>However, if you have configured advertisement of multiple routes (with the <b>add-path</b> or <b>advertise-inactive</b> statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong.</p> <ul style="list-style-type: none"> <li>• <b>hidden</b> (routes that are not used because of a routing policy).</li> </ul> |
| <i>destination-prefix</i>  | <p>Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul>                                                                                                                                                                      |

Table 373: show route Output Fields (*continued*)

| Field Name                                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [ <i>protocol, preference</i> ]                   | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• - —A hyphen indicates the last active route.</li> <li>• *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>                                                                                                                                                                                                                                   |
| <i>weeks:days</i><br><i>hours:minutes:seconds</i> | How long the route been known (for example, <b>2w4d 13:11:14</b> , or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| metric                                            | Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| localpref                                         | Local preference value included in the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| from                                              | Interface from which the route was received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| AS path                                           | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• I—IGP.</li> <li>• E—EGP.</li> <li>• ?—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• [ ]—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured.</li> <li>• { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• ( )—Parentheses enclose a confederation.</li> <li>• ( [ ] )—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |

Table 373: show route Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>validation-state</b> | <p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Unverified</b>—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>to</b>               | <p>Next hop to the destination. An angle bracket (&gt;) indicates that the route is the selected route.</p> <p>If the destination is <b>Discard</b>, traffic is dropped.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>via</b>              | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> <li>• <b>lsp-path-name</b>—Name of the LSP used to reach the next hop.</li> <li>• <b>label-action</b>—MPLS label and operation occurring at the next hop. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label). For VPNs, expect to see multiple <b>push</b> operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).</li> </ul> |

## Sample Output

### show route

```

user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
    *[MVPN/70] 19:53:41, metric2 1
    Indirect
1:65500:1:10.0.0.40/240
    *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
    *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30

```

```

AS path: I
> to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
[BGP/170] 19:53:25, localpref 100, from 10.0.0.33
AS path: I
> to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF

```

## show route

The following sample output shows a VPN route with composite next hops enabled. The first **Push** operation corresponds to the outer label. The second **Push** operation corresponds to the inner label.

```
user@host> show route 70.0.0.0
```

```

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

```

```

70.0.0.0/24      @[BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                  AS path: 13980 ?, validation-state: unverified
                  > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                  [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                  AS path: 13980 ?, validation-state: unverified
                  > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                  #[Multipath/255] 00:28:28, metric2 102
                  > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                  to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)

```

## show route destination-prefix

```
user@host> show route 172.16.0.0/12
```

```

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

172.16.0.0/12    *[Static/5] 2w4d 12:54:27
                  > to 192.168.167.254 via fxp0.0

```

## show route extensive

```
user@host> show route extensive
```

```

v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
        PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group 225.1.1.1

        Next hop type: Indirect
        Address: 0x92455b8
        Next-hop reference count: 2
        Source: 10.0.0.30
        Protocol next hop: 10.0.0.40
        Indirect next hop: 2 no-forward
        State: <Active Int Ext>
              Local AS: 65500 Peer AS: 65500
        Age: 3 Metric2: 1
        Validation State: unverified
        Task: BGP_65500.10.0.0.30+179
        Announcement bits (2): 0-PIM.v1 1-mvpn global task
        AS path: I (Originator) Cluster list: 10.0.0.30
        AS path: Originator ID: 10.0.0.40
        Communities: target:65520:100

```

```
Import Accepted
Localpref: 100
Router ID: 10.0.0.30
Primary Routing Table bgp.mvpn.0
Indirect next hops: 1
  Protocol next hop: 10.0.0.40 Metric: 1
  Indirect next hop: 2 no-forward
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
  10.0.0.40/32 Originating RIB: inet.3
    Metric: 1 Node path count: 1
    Forwarding nexthops: 1
      Nexthop: 10.0.24.4 via lt-0/3/0.24
```

## show route active-path

|                                    |                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3559</a><br><a href="#">Syntax (EX Series Switches) on page 3559</a>                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                      | <pre>show route active-path &lt;brief   detail   extensive   terse&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches)</b> | <pre>show route active-path &lt;brief   detail   extensive   terse&gt;</pre>                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>         | <p>Command introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                               |
| <b>Description</b>                 | Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.                                                                                                                                                                                                                                               |
| <b>Options</b>                     | <p><b>none</b>—Display all active routes.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>       | <a href="#">show route active-path on page 3559</a><br><a href="#">show route active-path brief on page 3560</a><br><a href="#">show route active-path detail on page 3560</a><br><a href="#">show route active-path extensive on page 3561</a><br><a href="#">show route active-path terse on page 3563</a>                                                                               |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                             |

## Sample Output

### show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
                  > via lo0.0
10.255.71.50/32   *[IS-IS/15] 00:18:13, metric 10
                  > to 100.1.2.1 via so-2/1/3.0
100.1.2.0/24      *[Direct/0] 00:18:36
                  > via so-2/1/3.0
```

```
100.1.2.2/32      *[Local/0] 00:18:41
                  Local via so-2/1/3.0
192.168.64.0/21  *[Direct/0] 21:33:52
                  > via fxp0.0
192.168.70.19/32 *[Local/0] 21:33:52
                  Local via fxp0.0
```

### show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see [show route active-path on page 3559](#).

### show route active-path detail

```
user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:37:10
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

10.255.71.50/32 (1 entry, 1 announced)
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Router, Next hop index: 397
    Next-hop reference count: 4
    Next hop: 100.1.2.1 via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:31 Metric: 10
    Task: IS-IS
    Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
    AS path: I

100.1.2.0/24 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:54
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

100.1.2.2/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
```



```

Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:59
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

```

### show route active-path extensive

```

user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via lo0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {100.1.2.1}
IS-IS level 2, LSP fragment 0
*IS-IS Preference: 15
Level: 1
Next hop type: Router, Next hop index: 397

```

```
Next-hop reference count: 4
Next hop: 100.1.2.1 via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:08 Metric: 10
Task: IS-IS
Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
AS path: I

100.1.2.0/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:31
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
AS path: I

100.1.2.2/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 24:36
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
```

AS path: I

### show route active-path terse

```
user@host> show route active-path terse
```

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

| A | Destination      | P | Prf | Metric 1 | Metric 2 | Next hop    | AS path |
|---|------------------|---|-----|----------|----------|-------------|---------|
| * | 10.255.70.19/32  | D | 0   |          |          | >1o0.0      |         |
| * | 10.255.71.50/32  | I | 15  | 10       |          | >100.1.2.1  |         |
| * | 100.1.2.0/24     | D | 0   |          |          | >so-2/1/3.0 |         |
| * | 100.1.2.2/32     | L | 0   |          |          | Local       |         |
| * | 192.168.64.0/21  | D | 0   |          |          | >fxp0.0     |         |
| * | 192.168.70.19/32 | L | 0   |          |          | Local       |         |

## show route all

---

|                                    |                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3564</a><br><a href="#">Syntax (EX Series Switches) on page 3564</a>                                                                                                                                                                                                                                               |
| <b>Syntax</b>                      | show route all<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                        |
| <b>Syntax (EX Series Switches)</b> | show route all                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                         |
| <b>Description</b>                 | Display information about all routes in all routing tables, including private, or internal, tables.                                                                                                                                                                                                                                           |
| <b>Options</b>                     | <b>none</b> —Display information about all routes in all routing tables, including private, or internal, tables.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.                                                                 |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>       | <a href="#">show route all on page 3564</a>                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>               | In Junos OS Release 9.5 and later, only the output fields for the <b>show route all</b> command display all routing tables, including private, or hidden, routing tables. The output field table of the <a href="#">show route</a> command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later. |

## Sample Output

### show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```
user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
1          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
2          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
800017     *[VPLS/7] 1d 14:00:16
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 14:00:26
            > via vt-3/2/0.32772, Pop
```

```
user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
1          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
2          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
800017     *[VPLS/7] 1d 13:54:49
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 13:54:59
            > via vt-3/2/0.32772, Pop
vt-3/2/0.32769 [VPLS/7] 1d 13:54:49
              Unusable
vt-3/2/0.32772 [VPLS/7] 1d 13:54:59
              Unusable
```

## show route aspath-regex

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3566</a><br><a href="#">Syntax (EX Series Switches) on page 3566</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                      | <code>show route aspath-regex <i>regular-expression</i></code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (EX Series Switches)</b> | <code>show route aspath-regex <i>regular-expression</i></code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                 | Display the entries in the routing table that match the specified autonomous system (AS) path regular expression.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                     | <b><i>regular-expression</i></b> —Regular expression that matches an entire AS path.<br><br><b><i>logical-system</i> (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Additional Information</b>      | <p>You can specify a regular expression as:</p> <ul style="list-style-type: none"><li>• An individual AS number</li><li>• A period wildcard used in place of an AS number</li><li>• An AS path regular expression that is enclosed in parentheses</li></ul> <p>You also can include the operators described in the table of AS path regular expression operators in the <i>Junos Policy Framework Configuration Guide</i>. The following list summarizes these operators:</p> <ul style="list-style-type: none"><li>• <b><i>{m,n}</i></b>—At least <i>m</i> and at most <i>n</i> repetitions of the AS path term.</li><li>• <b><i>{m}</i></b>—Exactly <i>m</i> repetitions of the AS path term.</li><li>• <b><i>{m,}</i></b>—<i>m</i> or more repetitions of the AS path term.</li><li>• <b><i>*</i></b>—Zero or more repetitions of an AS path term.</li><li>• <b><i>+</i></b>—One or more repetitions of an AS path term.</li><li>• <b><i>?</i></b>—Zero or one repetition of an AS path term.</li><li>• <b><i>aspath_term</i>   <i>aspath_term</i></b>—Match one of the two AS path terms.</li></ul> <p>When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:</p> <pre>show route aspath-regex ". * 234 ."</pre> |

|                          |                                                                                                                                                                                      |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Required Privilege Level | view                                                                                                                                                                                 |
| Related Documentation    | <ul style="list-style-type: none"> <li>• <i>Example: Using AS Path Regular Expressions</i></li> </ul>                                                                                |
| List of Sample Output    | <a href="#">show route aspath-regex (Matching a Specific AS Number) on page 3567</a><br><a href="#">show route aspath-regex (Matching Any Path with Two AS Numbers) on page 3567</a> |
| Output Fields            | For information about output fields, see the output field table for the <a href="#">show route</a> command.                                                                          |

## Sample Output

### show route aspath-regex (Matching a Specific AS Number)

```

user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25      *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25   *[IS-IS/15] 09:15:37, metric 37, tag 1
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
                   [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
...

```

### show route aspath-regex (Matching Any Path with Two AS Numbers)

```

user@host> show route aspath-regex ?.* 234 3561.*?

inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 2685 2686 Incomplete
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 5696 7369 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 6327 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...

```

## show route best

---

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax              | <a href="#">Syntax on page 3568</a><br><a href="#">Syntax (EX Series Switches) on page 3568</a>                                                                                                                                                                                                                                                                                                                  |
| Syntax                      | <code>show route best <i>destination-prefix</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                      |
| Syntax (EX Series Switches) | <code>show route best <i>destination-prefix</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code>                                                                                                                                                                                                                                                                                                |
| Release Information         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                            |
| Description                 | Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.                                                                                                                                                                                                                                                      |
| Options                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b><i>destination-prefix</i></b> —Address or range of addresses.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level    | view                                                                                                                                                                                                                                                                                                                                                                                                             |
| List of Sample Output       | <a href="#">show route best on page 3568</a><br><a href="#">show route best detail on page 3569</a><br><a href="#">show route best extensive on page 3570</a><br><a href="#">show route best terse on page 3570</a>                                                                                                                                                                                              |
| Output Fields               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                                   |

## Sample Output

### show route best

```
user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[OSPF/10] 1d 13:19:20, metric 2
                  > to 10.31.1.6 via ge-3/1/0.0
                  via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[RSVP/7] 1d 13:20:13, metric 2
```



```

> via so-0/3/0.0, label-switched-path green-r1-r3

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8          *[Direct/0] 2d 01:43:34
                    > via fxp2.0
                    [Direct/0] 2d 01:43:34
                    > via fxp1.0

```

### show route best detail

```

user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
  *OSPF   Preference: 10
          Next-hop reference count: 9
          Next hop: 10.31.1.6 via ge-3/1/0.0, selected
          Next hop: via so-0/3/0.0
          State: <Active Int>
          Local AS: 69
          Age: 1d 13:20:06      Metric: 2
          Area: 0.0.0.0
          Task: OSPF
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 5
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r3
          Label operation: Push 100016
          State: <Active Int>
          Local AS: 69
          Age: 1d 13:20:59      Metric: 2
          Task: RSVP
          Announcement bits (1): 1-Resolve tree 2
          AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via fxp2.0, selected
          State: <Active Int>
          Age: 2d 1:44:20
          Task: IF
          AS path: I
  Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via fxp1.0, selected
          State: <NotBest Int>
          Inactive reason: No difference
          Age: 2d 1:44:20

```

Task: IF  
AS path: I

### show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see [show route best detail on page 3569](#).

### show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.255.70.103/32  0 10      2          >10.31.1.6
                               so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.255.70.103/32  R  7      2          >so-0/3/0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.0.0.0/8        D  0          >fxp2.0
                    D  0          >fxp1.0
```

## show route brief

|                                    |                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3571</a><br><a href="#">Syntax (EX Series Switches) on page 3571</a>                                                                                                                                                                                                                                          |
| <b>Syntax</b>                      | show route brief<br><destination-prefix><br><logical-system (all   logical-system-name)>                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switches)</b> | show route brief<br><destination-prefix>                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                    |
| <b>Description</b>                 | Display brief information about the active entries in the routing tables.                                                                                                                                                                                                                                                                |
| <b>Options</b>                     | <b>none</b> —Display all active entries in the routing table.<br><br><b>destination-prefix</b> —(Optional) Display active entries for the specified address or range of addresses.<br><br><b>logical-system (all   logical-system-name)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>       | <a href="#">show route brief on page 3571</a>                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>               | For information about output fields, see the Output Field table of the <a href="#">show route</a> command.                                                                                                                                                                                                                               |

## Sample Output

### show route brief

```

user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                   Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                   > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.0.0/18      *[Static/5] 1w5d 20:30:29
                   > to 192.168.167.254 via fxp0.0
192.168.40.0/22     *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.64.0/18     *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.164.0/22    *[Direct/0] 2w4d 13:11:14
                   > via fxp0.0

```

```
192.168.164.51/32  *[Local/0] 2w4d 13:11:14
                  Local via fxp0.0
207.17.136.192/32  *[Static/5] 2w4d 13:11:14
                  > to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16     *[Direct/0] 1w5d 20:30:28
                  > via fe-0/0/3.0
100.101.2.3/32    *[Local/0] 1w5d 20:30:28
                  Local via fe-0/0/3.0
224.0.0.5/32      *[OSPF/10] 1w5d 20:30:29, metric 1
                  MultiRecv
```

## show route community

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax              | <a href="#">Syntax on page 3573</a><br><a href="#">Syntax (EX Series Switches) on page 3573</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Syntax                      | show route community <i>as-number:community-value</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Syntax (EX Series Switches) | show route community <i>as-number:community-value</i><br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Release Information         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Description                 | Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Options                     | <p><i>as-number:community-value</i>—One or more community identifiers. <i>as-number</i> is the AS number, and <i>community-value</i> is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks. Community identifiers can include wildcards.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information      | Specifying the community option displays all routes matching the community found within the routing table. The community option does not limit the output to only the routes being advertised to the neighbor after any egress routing policy.                                                                                                                                                                                                                                                                                                                                            |
| Required Privilege Level    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Related Documentation       | <ul style="list-style-type: none"> <li>• <a href="#">show route detail on page 3582</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| List of Sample Output       | <a href="#">show route community on page 3573</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Output Fields               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                                                                                                                                                                                                            |

## Sample Output

### show route community

```

user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

```

```
4.0.0.0/8      *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 1 IGP
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
6.0.0.0/8      *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 568 721 Incomplete
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
9.2.0.0/16     *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 1673 1675 1747 IGP
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
```

## show route community-name

|                                    |                                                                                                                                                                                                                                                                                                                |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3575</a><br><a href="#">Syntax (EX Series Switches) on page 3575</a>                                                                                                                                                                                                                |
| <b>Syntax</b>                      | <b>show route community-name</b> <i>community-name</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                         |
| <b>Syntax (EX Series Switches)</b> | <b>show route community-name</b> <i>community-name</i><br><brief   detail   extensive   terse>                                                                                                                                                                                                                 |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                          |
| <b>Description</b>                 | Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community, specified by a community name.                                                                                                                                                                  |
| <b>Options</b>                     | <i>community-name</i> —Name of the community.<br><br><b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>       | <a href="#">show route community-name on page 3575</a>                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                 |

## Sample Output

### show route community-name

```

user@host> show route community-name red-com
inet.0: 17 destinations, 17 routes (16 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

instance1.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.212/32  *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                   AS path: 300 I
                   > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
20.20.20.20/32    *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                   AS path: I
                   > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
100.1.4.0/24     *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204

```

```
AS path: I
> to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.204:10:10.255.245.212/32
    *[BGP/170] 00:06:40, localpref 100, from 10.255.245.204
    AS path: 300 I
    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:20.20.20.20/32
    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
    AS path: I
    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:100.1.4.0/24
    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
    AS path: I
    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

instance1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```



## show route damping

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3577</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3577</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax</b>                                   | show route damping (decayed   history   suppressed)<br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (EX Series Switch and QFX Series)</b> | show route damping (decayed   history   suppressed)<br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                              | Display the BGP routes for which updates might have been reduced because of route flap damping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                                  | <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p><b>decayed</b>—Display route damping entries that might no longer be valid, but are not suppressed.</p> <p><b>history</b>—Display entries that have already been withdrawn, but have been logged.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>suppressed</b>—Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.</p> |
| <b>Required Privilege Level</b>                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"> <li>• <a href="#">clear bgp damping on page 3038</a></li> <li>• <a href="#">show policy damping on page 3072</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>                    | <a href="#">show route damping decayed detail on page 3580</a><br><a href="#">show route damping history on page 3581</a><br><a href="#">show route damping history detail on page 3581</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>                            | <a href="#">Table 374 on page 3578</a> lists the output fields for the <b>show route damping</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 374: show route damping Output Fields

| Field Name                                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output         |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <i>routing-table-name</i>                    | Name of the routing table—for example, <b>inet.0</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels              |
| <b>destinations</b>                          | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels              |
| <b>number routes</b>                         | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li>• <b>active</b></li> <li>• <b>holddown</b> (routes that are in a pending state before being declared inactive)</li> <li>• <b>hidden</b> (the routes are not used because of a routing policy)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels              |
| <b>destination-prefix (entry, announced)</b> | Destination prefix. The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>[protocol, preference]</b>                | Protocol from which the route was learned and the preference value for the route. <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p> | All levels              |
| <b>Next-hop reference count</b>              | Number of references made to the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> |
| <b>Source</b>                                | IP address of the route source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Next hop</b>                              | Network layer address of the directly reachable neighboring system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> |
| <b>via</b>                                   | Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word <b>Selected</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive</b> |
| <b>Protocol next hop</b>                     | Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b> |
| <b>Indirect next hop</b>                     | Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> |
| <b>State</b>                                 | Flags for this route. For a description of possible values for this field, see the output field table for the <a href="#">show route detail</a> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> |

Table 374: show route damping Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output  |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Local AS          | AS number of the local routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | detail extensive |
| Peer AS           | AS number of the peer routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | detail extensive |
| Age               | How long the route has been known.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | detail extensive |
| Metric            | Metric for the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | detail extensive |
| Task              | Name of the protocol that has added the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | detail extensive |
| Announcement bits | List of protocols that announce this route. <i>n-Resolve inet</i> indicates that the route is used for route resolution for next hops found in the routing table. <i>n</i> is an index used by Juniper Networks customer support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | detail extensive |
| AS path           | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• I—IGP.</li> <li>• E—EGP.</li> <li>• ?—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• [ ]—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured.</li> <li>• { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• ( )—Parentheses enclose a confederation.</li> <li>• ( [ ] )—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> | All levels       |
| to                | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | brief none       |
| via               | Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word <b>Selected</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | brief none       |
| Communities       | Community path attribute for the route. See the output field table for the <a href="#">show route detail</a> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | detail extensive |
| Localpref         | Local preference value included in the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels       |
| Router ID         | BGP router ID as advertised by the neighbor in the open message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | detail extensive |

Table 374: show route damping Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                        | Level of Output         |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Merit (last update/now)</b> | Last updated and current figure-of-merit value.                                                                                                                          | <b>detail extensive</b> |
| <b>damping-parameters</b>      | Name that identifies the damping parameters used, which is defined in the damping statement at the <b>[edit policy-options]</b> hierarchy level.                         | <b>detail extensive</b> |
| <b>Last update</b>             | Time of most recent change in path attributes.                                                                                                                           | <b>detail extensive</b> |
| <b>First update</b>            | Time of first change in path attributes, which started the route damping process.                                                                                        | <b>detail extensive</b> |
| <b>Flaps</b>                   | Number of times the route has gone up or down or its path attributes have changed.                                                                                       | <b>detail extensive</b> |
| <b>Suppressed</b>              | ( <b>suppressed</b> keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it. | All levels              |
| <b>Reusable in</b>             | ( <b>suppressed</b> keyword only) Time when a suppressed route will again be available.                                                                                  | All levels              |
| <b>Preference will be</b>      | ( <b>suppressed</b> keyword only) Preference value that will be applied to the route when it is again active.                                                            | All levels              |

## Sample Output

### show route damping decayed detail

```

user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
  *BGP      Preference: 170/-101
            Next-hop reference count: 151973
            Source: 172.23.2.129
            Next hop: via so-1/2/0.0
            Next hop: via so-5/1/0.0, selected
            Next hop: via so-6/0/0.0
            Protocol next hop: 172.23.2.129
            Indirect next hop: 89a1a00 264185
            State: <Active Ext>
            Local AS: 65000 Peer AS: 65490
            Age: 3:28      Metric2: 0
            Task: BGP_65490.172.23.2.129+179
            Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

        6-Resolve tree 2 7-Resolve tree 3
        AS path: 65490 65520 65525 65525 65525 65525 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:701
        Localpref: 100
        Router ID: 172.23.2.129
        Merit (last update/now): 1934/1790
        damping-parameters: damping-high

```

```

Last update:      00:03:28 First update:      00:06:40
Flaps: 2

```

### show route damping history

```

user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15      [BGP ] 2d 22:47:58, localpref 100
                  AS path: 65220 65501 65502 I
                  > to 192.168.60.85 via so-3/1/0.0

```

### show route damping history detail

```

user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
    BGP                /-101
        Next-hop reference count: 69058
        Source: 192.168.60.85
        Next hop: 192.168.60.85 via so-3/1/0.0, selected
        State: <Hidden Ext>
        Inactive reason: Unusable path
        Local AS: 65000 Peer AS: 65220
        Age: 2d 22:48:10
        Task: BGP_65220.192.168.60.85+179
        AS path: 65220 65501 65502 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:3561
        Localpref: 100
        Router ID: 192.168.80.25
        Merit (last update/now): 1000/932
        damping-parameters: set-normal
        Last update:      00:01:05 First update:      00:01:05
        Flaps: 1

```

## show route detail

|                                    |                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3582</a><br><a href="#">Syntax (EX Series Switches) on page 3582</a>                                                                                                                                                                                                                                                         |
| <b>Syntax</b>                      | show route detail<br><destination-prefix><br><logical-system (all   logical-system-name)>                                                                                                                                                                                                                                                               |
| <b>Syntax (EX Series Switches)</b> | show route detail<br><destination-prefix>                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.                                                                                                                                                         |
| <b>Description</b>                 | Display detailed information about the active entries in the routing tables.                                                                                                                                                                                                                                                                            |
| <b>Options</b>                     | <b>none</b> —Display all active entries in the routing table on all systems.<br><br><b>destination-prefix</b> —(Optional) Display active entries for the specified address or range of addresses.<br><br><b>logical-system (all   logical-system-name)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>       | <a href="#">show route detail on page 3591</a><br><a href="#">show route detail (with BGP Multipath) on page 3597</a><br><a href="#">show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 3597</a><br><a href="#">show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 3598</a>  |
| <b>Output Fields</b>               | Table 375 on page 3582 describes the output fields for the <b>show route detail</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                    |

**Table 375: show route detail Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).                                                                                                                                                                                                                                                                                                                |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                                         |
| <i>number routes</i>       | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li><b>active</b> (routes that are active)</li> <li><b>holddown</b> (routes that are in the pending state before being declared inactive)</li> <li><b>hidden</b> (routes that are not used because of a routing policy)</li> </ul> |

Table 375: show route detail Output Fields (*continued*)

| Field Name                                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>route-destination</i><br>(entry, announced) | <p>Route destination (for example:10.0.0.1/24). The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul> |
| label stacking                                 | <p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| [ <i>protocol, preference</i> ]                | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+—</b>A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>- —</b>A hyphen indicates the last active route.</li> <li>• <b>*—</b>An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>                                                                                                                                                                                                                           |
| Level                                          | <p>(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Route Distinguisher                            | IP subnet augmented with a 64-bit prefix.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| PMSI                                           | Provider multicast service interface (MVPN routing table).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Next-hop type                                  | Type of next hop. For a description of possible values for this field, see <a href="#">Table 376 on page 3587</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 375: show route detail Output Fields (*continued*)

| Field Name                                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Next-hop reference count</b>                      | Number of references made to the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Flood nexthop branches exceed maximum message</b> | Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Source</b>                                        | IP address of the route source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Next hop</b>                                      | Network layer address of the directly reachable neighboring system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>via</b>                                           | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> </ul> |
| <b>Label-switched-path<br/>lsp-path-name</b>         | Name of the LSP used to reach the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Label operation</b>                               | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Interface</b>                                     | (Local only) Local interface name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Protocol next hop</b>                             | Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Indirect next hop</b>                             | Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>State</b>                                         | State of the route (a route can be in more than one state). See <a href="#">Table 377 on page 3588</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Local AS</b>                                      | AS number of the local routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Age</b>                                           | How long the route has been known.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>AIGP</b>                                          | Accumulated interior gateway protocol (AIGP) BGP attribute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Metricn</b>                                       | Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |



Table 375: show route detail Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MED-plus-IGP</b>        | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>TTL-Action</b>          | <p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see <a href="#">show route table</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Task</b>                | Name of the protocol that has added the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Announcement bits</b>   | List of protocols that announce this route. <b>n-Resolve inet</b> indicates that the route is used for route resolution for next hops found in the routing table. <b>n</b> is an index used by Juniper Networks customer support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>AS path</b>             | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893.</li> <li>• <b>[ ]</b>—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| <b>validation-state</b>    | <p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Unverified</b>—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>FECs bound to route</b> | Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 375: show route detail Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                   |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Primary Upstream          | When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path. |
| RPF Nexthops              | When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.                                             |
| Label                     | Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.                         |
| weight                    | Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.                                                                                                                                                  |
| VC Label                  | MPLS label assigned to the Layer 2 circuit virtual connection.                                                                                                                                                                                                                                                      |
| MTU                       | Maximum transmission unit (MTU) of the Layer 2 circuit.                                                                                                                                                                                                                                                             |
| VLAN ID                   | VLAN identifier of the Layer 2 circuit.                                                                                                                                                                                                                                                                             |
| Prefixes bound to route   | Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.                                                                                                                                                                                                                  |
| Communities               | Community path attribute for the route. See <a href="#">Table 378 on page 3590</a> for all possible values for this field.                                                                                                                                                                                          |
| Layer2-info: encaps       | Layer 2 encapsulation (for example, VPLS).                                                                                                                                                                                                                                                                          |
| control flags             | Control flags: <b>none</b> or <b>Site Down</b> .                                                                                                                                                                                                                                                                    |
| mtu                       | Maximum transmission unit (MTU) information.                                                                                                                                                                                                                                                                        |
| Label-Base, range         | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.                                                                                                                                              |
| status vector             | Layer 2 VPN and VPLS network layer reachability information (NLRI).                                                                                                                                                                                                                                                 |
| Accepted Multipath        | Current active path when BGP multipath is configured.                                                                                                                                                                                                                                                               |
| Accepted MultipathContrib | Path currently contributing to BGP multipath.                                                                                                                                                                                                                                                                       |
| Localpref                 | Local preference value included in the route.                                                                                                                                                                                                                                                                       |
| Router ID                 | BGP router ID as advertised by the neighbor in the open message.                                                                                                                                                                                                                                                    |
| Primary Routing Table     | In a routing table group, the name of the primary routing table in which the route resides.                                                                                                                                                                                                                         |
| Secondary Tables          | In a routing table group, the name of one or more secondary tables in which the route resides.                                                                                                                                                                                                                      |

Table 376 on page 3587 describes all possible values for the Next-hop Types output field.

**Table 376: Next-hop Types Output Field Values**

| Next-Hop Type                   | Description                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Broadcast (bcast)</b>        | Broadcast next hop.                                                                                                                                                                                                                                                            |
| <b>Deny</b>                     | Deny next hop.                                                                                                                                                                                                                                                                 |
| <b>Discard</b>                  | Discard next hop.                                                                                                                                                                                                                                                              |
| <b>Flood</b>                    | Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast. |
| <b>Hold</b>                     | Next hop is waiting to be resolved into a unicast or multicast type.                                                                                                                                                                                                           |
| <b>Indexed (idxd)</b>           | Indexed next hop.                                                                                                                                                                                                                                                              |
| <b>Indirect (indr)</b>          | Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.                                                   |
| <b>Interface</b>                | Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.                                                                                                                   |
| <b>Local (locl)</b>             | Local address on an interface. This next-hop type causes packets with this destination address to be received locally.                                                                                                                                                         |
| <b>Multicast (mcst)</b>         | Wire multicast next hop (limited to the LAN).                                                                                                                                                                                                                                  |
| <b>Multicast discard (mdsc)</b> | Multicast discard.                                                                                                                                                                                                                                                             |
| <b>Multicast group (mgrp)</b>   | Multicast group member.                                                                                                                                                                                                                                                        |
| <b>Receive (recv)</b>           | Receive.                                                                                                                                                                                                                                                                       |
| <b>Reject (rjct)</b>            | Discard. An ICMP unreachable message was sent.                                                                                                                                                                                                                                 |
| <b>Resolve (rslv)</b>           | Resolving next hop.                                                                                                                                                                                                                                                            |
| <b>Routed multicast (mcrt)</b>  | Regular multicast next hop.                                                                                                                                                                                                                                                    |

Table 376: Next-hop Types Output Field Values (*continued*)

| Next-Hop Type         | Description                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Router</b>         | <p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> <li>• Must not be a direct or local subnet for the routing device.</li> <li>• Must have a next hop that is directly connected to the routing device.</li> </ul> |
| <b>Table</b>          | Routing table next hop.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Unicast (ucst)</b> | Unicast.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Unilist (ulst)</b> | List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.                                                                                                                                                                                                                                                                                                                |

Table 377 on page 3588 describes all possible values for the State output field. A route can be in more than one state (for example, <Active NoReadvrt Int Ext>).

Table 377: State Output Field Values

| Value                                        | Description                                                                      |
|----------------------------------------------|----------------------------------------------------------------------------------|
| <b>Accounting</b>                            | Route needs accounting.                                                          |
| <b>Active</b>                                | Route is active.                                                                 |
| <b>Always Compare MED</b>                    | Path with a lower multiple exit discriminator (MED) is available.                |
| <b>AS path</b>                               | Shorter AS path is available.                                                    |
| <b>Cisco Non-deterministic MED selection</b> | Cisco nondeterministic MED is enabled, and a path with a lower MED is available. |
| <b>Clone</b>                                 | Route is a clone.                                                                |
| <b>Cluster list length</b>                   | Length of cluster list sent by the route reflector.                              |
| <b>Delete</b>                                | Route has been deleted.                                                          |
| <b>Ex</b>                                    | Exterior route.                                                                  |
| <b>Ext</b>                                   | BGP route received from an external BGP neighbor.                                |

Table 377: State Output Field Values (*continued*)

| Value                                                    | Description                                                                                                                                                                          |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>FlashAll</b>                                          | Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes. |
| <b>Hidden</b>                                            | Route not used because of routing policy.                                                                                                                                            |
| <b>IfCheck</b>                                           | Route needs forwarding RPF check.                                                                                                                                                    |
| <b>IGP metric</b>                                        | Path through next hop with lower IGP metric is available.                                                                                                                            |
| <b>Inactive reason</b>                                   | Flags for this route, which was not selected as best for a particular destination.                                                                                                   |
| <b>Initial</b>                                           | Route being added.                                                                                                                                                                   |
| <b>Int</b>                                               | Interior route.                                                                                                                                                                      |
| <b>Int Ext</b>                                           | BGP route received from an internal BGP peer or a BGP confederation peer.                                                                                                            |
| <b>Interior &gt; Exterior &gt; Exterior via Interior</b> | Direct, static, IGP, or EBGP path is available.                                                                                                                                      |
| <b>Local Preference</b>                                  | Path with a higher local preference value is available.                                                                                                                              |
| <b>Martian</b>                                           | Route is a martian (ignored because it is obviously invalid).                                                                                                                        |
| <b>MartianOK</b>                                         | Route exempt from martian filtering.                                                                                                                                                 |
| <b>Next hop address</b>                                  | Path with lower metric next hop is available.                                                                                                                                        |
| <b>No difference</b>                                     | Path from neighbor with lower IP address is available.                                                                                                                               |
| <b>NoReadvrt</b>                                         | Route not to be advertised.                                                                                                                                                          |
| <b>NotBest</b>                                           | Route not chosen because it does not have the lowest MED.                                                                                                                            |
| <b>Not Best in its group</b>                             | Incoming BGP AS is not the best of a group (only one AS can be the best).                                                                                                            |
| <b>NotInstall</b>                                        | Route not to be installed in the forwarding table.                                                                                                                                   |
| <b>Number of gateways</b>                                | Path with a greater number of next hops is available.                                                                                                                                |
| <b>Origin</b>                                            | Path with a lower origin code is available.                                                                                                                                          |
| <b>Pending</b>                                           | Route pending because of a hold-down configured on another route.                                                                                                                    |

Table 377: State Output Field Values (*continued*)

| Value                                 | Description                                                                                                                                                                                                                       |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release</b>                        | Route scheduled for release.                                                                                                                                                                                                      |
| <b>RIB preference</b>                 | Route from a higher-numbered routing table is available.                                                                                                                                                                          |
| <b>Route Distinguisher</b>            | 64-bit prefix added to IP subnets to make them unique.                                                                                                                                                                            |
| <b>Route Metric or MED comparison</b> | Route with a lower metric or MED is available.                                                                                                                                                                                    |
| <b>Route Preference</b>               | Route with lower preference value is available                                                                                                                                                                                    |
| <b>Router ID</b>                      | Path through a neighbor with lower ID is available.                                                                                                                                                                               |
| <b>Secondary</b>                      | Route not a primary route.                                                                                                                                                                                                        |
| <b>Unusable path</b>                  | Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> <li>• The route is damped.</li> <li>• The route is rejected by an import policy.</li> <li>• The route is unresolved.</li> </ul> |
| <b>Update source</b>                  | Last tiebreaker is the lowest IP address value.                                                                                                                                                                                   |

Table 378 on page 3590 describes the possible values for the Communities output field.

Table 378: Communities Output Field Values

| Value                                                   | Description                                                                                                                                                                                                                                                                           |
|---------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>area-number</i>                                      | 4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.                                        |
| <b>bandwidth: local AS number:link-bandwidth-number</b> | Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute. |
| <b>domain-id</b>                                        | Unique configurable number that identifies the OSPF domain.                                                                                                                                                                                                                           |
| <b>domain-id-vendor</b>                                 | Unique configurable number that further identifies the OSPF domain.                                                                                                                                                                                                                   |
| <i>link-bandwidth-number</i>                            | Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).                                                                                                                                                                                                               |
| <i>local AS number</i>                                  | Local AS number: from 1 through 65,535.                                                                                                                                                                                                                                               |
| <i>options</i>                                          | 1 byte. Currently this is only used if the route type is 5 or 7. Setting the least significant bit in the field indicates that the route carries a type 2 metric.                                                                                                                     |

Table 378: Communities Output Field Values (*continued*)

| Value                                | Description                                                                                                                                                                                                                                           |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>origin</b>                        | (Used with VPNs) Identifies where the route came from.                                                                                                                                                                                                |
| <b>ospf-route-type</b>               | 1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses. |
| <b>route-type-vendor</b>             | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute <b>0x8000</b> . The format is <b>area-number:ospf-route-type:options</b> .                                          |
| <b>rte-type</b>                      | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute <b>0x0306</b> . The format is <b>area-number:ospf-route-type:options</b> .                                          |
| <b>target</b>                        | Defines which VPN the route participates in; <b>target</b> has the format <b>32-bit IP address:16-bit number</b> . For example, 10.19.0.0:100.                                                                                                        |
| <b>unknown IANA</b>                  | Incoming IANA codes with a value between <b>0x1</b> and <b>0x7fff</b> . This code of the BGP extended community attribute is accepted, but it is not recognized.                                                                                      |
| <b>unknown OSPF vendor community</b> | Incoming IANA codes with a value above <b>0x8000</b> . This code of the BGP extended community attribute is accepted, but it is not recognized.                                                                                                       |

## Sample Output

### show route detail

```

user@host> show route detail

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:31:43
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2
    Next hop: via so-0/3/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:30:17
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I
  OSPF Preference: 10

```

```
Next-hop reference count: 1
Next hop: via so-0/3/0.0, selected
State: <Int>
Inactive reason: Route Preference
Local AS: 69
Age: 1:30:17 Metric: 1
Area: 0.0.0.0
Task: OSPF
AS path: I

10.31.1.1/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 7
    Interface: so-0/3/0.0
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:30:20
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
  *OSPF Preference: 10
    Next-hop reference count: 9
    Next hop: via so-0/3/0.0
    Next hop: 10.31.1.6 via ge-3/1/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:29:56 Metric: 2
    Area: 0.0.0.0
    Task: OSPF
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:43
    Task: IGMP
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I
```



```

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
          Label-switched-path green-r1-r3
          Label operation: Push 100096
          State: <Active Int>
          Local AS:    69
          Age: 1:25:49   Metric: 2
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r2
          State: <Active Int>
          Local AS:    69
          Age: 1:25:49   Metric: 1
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

private__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via lo0.0, selected
          State: <Active Int>
          Local AS:    69
          Age: 1:31:44
          Task: IF
          AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0 (1 entry, 1 announced)
  *MPLS   Preference: 0
          Next hop type: Receive
          Next-hop reference count: 6
          State: <Active Int>
          Local AS:    69
          Age: 1:31:45   Metric: 1
          Task: MPLS
          Announcement bits (1): 0-KRT
          AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

299840 (1 entry, 1 announced)

```

```
TSI:
KRT in-kerne 299840 /52 -> {indirect(1048575)}
    *RSVP   Preference: 7/2
            Next hop type: Flood
            Address: 0x9174a30
            Next-hop reference count: 4
            Next hop type: Router, Next hop index: 798
            Address: 0x9174c28
            Next-hop reference count: 2
            Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
            Label-switched-path R2-to-R4-2p2mp
            Label operation: Pop
            Next hop type: Router, Next hop index: 1048574
            Address: 0x92544f0
            Next-hop reference count: 2
            Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
            Label-switched-path R2-to-R200-p2mp
            Label operation: Pop
            Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
            Label operation: Pop
            State: <Active Int>
            Age: 1:29      Metric: 1
            Task: RSVP
            Announcement bits (1): 0-KRT
            AS path: I...

800010 (1 entry, 1 announced)
    *VPLS   Preference: 7
            Next-hop reference count: 2
            Next hop: via vt-3/2/0.32769, selected
            Label operation: Pop
            State: <Active Int>
            Age: 1:29:30
            Task: Common L2 VC
            Announcement bits (1): 0-KRT
            AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
    *VPLS   Preference: 7
            Next-hop reference count: 2
            Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
            Label-switched-path green-r1-r3
            Label operation: Push 800012, Push 100096(top)
            Protocol next hop: 10.255.70.103
            Push 800012
            Indirect next hop: 87272e4 1048574
            State: <Active Int>
            Age: 1:29:30   Metric2: 2
            Task: Common L2 VC
            Announcement bits (2): 0-KRT 1-Common L2 VC
            AS path: I
            Communities: target:11111:1 Layer2-info: encaps:VPLS,
            control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
    *Direct Preference: 0
            Next hop type: Interface
            Next-hop reference count: 1
            Next hop: via lo0.0, selected
```

```

        State: <Active Int>
        Local AS: 69
        Age: 1:31:44
        Task: IF
        AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::16/128 (1 entry, 1 announced)
  *MLD Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:43
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.16385, selected
    State: <Active NoReadvrt Int>
    Age: 1:31:44
    Task: IF
    AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

```

```
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 1:25:49 Metric2: 1
    AIGP 210
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-1
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-101
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:31:40 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
    Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
```

```

Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512

```

### show route detail (with BGP Multipath)

```
user@host> show route detail
```

```

10.1.1.8/30 (2 entries, 1 announced)
  *BGP   Preference: 170/-101
        Next hop type: Router, Next hop index: 262142
        Address: 0x901a010
        Next-hop reference count: 2
        Source: 10.1.1.2
        Next hop: 10.1.1.2 via ge-0/3/0.1, selected
        Next hop: 10.1.1.6 via ge-0/3/0.5
        State: <Active Ext>
        Local AS:      1 Peer AS:      2
        Age: 5:04:43
        Validation State: unverified
        Task: BGP_2.10.1.1.2+59955
        Announcement bits (1): 0-KRT
        AS path: 2 I
        Accepted Multipath
        Localpref: 100
        Router ID: 1.1.1.2
  BGP   Preference: 170/-101
        Next hop type: Router, Next hop index: 678
        Address: 0x8f97520
        Next-hop reference count: 9
        Source: 10.1.1.6
        Next hop: 10.1.1.6 via ge-0/3/0.5, selected
        State: <NotBest Ext>
        Inactive reason: Not Best in its group - Active preferred
        Local AS:      1 Peer AS:      2
        Age: 5:04:43
        Validation State: unverified
        Task: BGP_2.10.1.1.6+58198
        AS path: 2 I
        Accepted MultipathContrib
        Localpref: 100
        Router ID: 1.1.1.3

```

### show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
  *LDP   Preference: 9
        Next hop type: Flood
        Next-hop reference count: 3
        Address: 0x9097d90
        Next hop: via vt-0/1/0.1
        Next-hop index: 661
        Label operation: Pop
        Address: 0x9172130
        Next hop: via so-0/0/3.0
        Next-hop index: 654
        Label operation: Swap 299872
        State: **Active Int>
        Local AS: 1001

```

```
Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2
```

#### show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```
mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
    Next hop type: Flood
    Address: 0x2735208
    Next-hop reference count: 3
    Next hop type: Router, Next hop index: 1397
    Address: 0x2735d2c
    Next-hop reference count: 3
    Next hop: 1.3.8.2 via ge-1/2/22.0
    Label operation: Pop
    Load balance label: None;
    Next hop type: Router, Next hop index: 1395
    Address: 0x2736290
    Next-hop reference count: 3
    Next hop: 1.3.4.2 via ge-1/2/18.0
    Label operation: Pop
    Load balance label: None;
    State: <Active Int AckRequest MulticastRPF>
    Local AS: 10
    Age: 54:05      Metric: 1
    Validation State: unverified
    Task: LDP
    Announcement bits (1): 0-KRT
    AS path: I
    FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
      Primary Upstream : 1.1.1.3:0--1.1.1.2:0
        RPF Nexthops :
          ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
          ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
      Backup Upstream : 1.1.1.3:0--1.1.1.6:0
        RPF Nexthops :
          ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
          ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe
```

## show route exact

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3599</a><br><a href="#">Syntax (EX Series Switches) on page 3599</a>                                                                                                                                                                                                                                                                                                           |
| <b>Syntax</b>                      | show route exact <i>destination-prefix</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                |
| <b>Syntax (EX Series Switches)</b> | show route exact <i>destination-prefix</i><br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                 | Display only the routes that exactly match the specified address or range of addresses.                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b>destination-prefix</b> —Address or range of addresses.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>       | <a href="#">show route exact on page 3599</a><br><a href="#">show route exact detail on page 3599</a><br><a href="#">show route exact extensive on page 3600</a><br><a href="#">show route exact terse on page 3600</a>                                                                                                                                                                                   |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                            |

## Sample Output

### show route exact

```

user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                  > to 192.168.71.254 via fxp0.0

```

### show route exact detail

```

user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)

```

```
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2d 3:30:26
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I
```

#### show route exact extensive

```
user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:25:18
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I
```

#### show route exact terse

```
user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 207.17.136.0/24  S  5                >192.168.71.254
```



## show route export

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3601</a><br><a href="#">Syntax (EX Series Switches) on page 3601</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                      | <pre>show route export &lt;brief   detail&gt; &lt;instance &lt;instance-name&gt;   routing-table-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switches)</b> | <pre>show route export &lt;brief   detail&gt; &lt;instance &lt;instance-name&gt;   routing-table-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                 | Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                     | <p><b>none</b>—(Same as <b>brief</b>.) Display standard information about policy-based export for all instances and routing tables on all systems.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance &lt;instance-name&gt;</b>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>routing-table-name</b>—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the <b>show route export inet</b> command).</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>       | <a href="#">show route export on page 3602</a><br><a href="#">show route export detail on page 3602</a><br><a href="#">show route export instance detail on page 3602</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>               | <a href="#">Table 379 on page 3601</a> lists the output fields for the <b>show route export</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

**Table 379: show route export Output Fields**

| Field Name                 | Field Description                                                                                                                                           | Level of Output   |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| Table or <i>table-name</i> | Name of the routing tables that either import or export routes.                                                                                             | All levels        |
| Routes                     | Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one. | <b>brief</b> none |

Table 379: show route export Output Fields (*continued*)

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                           | Level of Output   |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| Export        | Whether the table is currently exporting routes to other tables: <b>Y</b> or <b>N</b> (Yes or No).                                                                                                                                                                                                                                                                          | <b>brief</b> none |
| Import        | Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)                                                                                                                                                                                                                                                  | <b>detail</b>     |
| Flags         | ( <b>instance</b> keyword only) Flags for this feature on this instance: <ul style="list-style-type: none"> <li><b>config auto-policy</b>—The policy was deduced from the configured IGP export policies.</li> <li><b>cleanup</b>—Configuration information for this instance is no longer valid.</li> <li><b>config</b>—The instance was explicitly configured.</li> </ul> | <b>detail</b>     |
| Options       | ( <b>instance</b> keyword only) Configured option displays the type of routing tables the feature handles: <ul style="list-style-type: none"> <li><b>unicast</b>—Indicates <i>instance.inet.0</i>.</li> <li><b>multicast</b>—Indicates <i>instance.inet.2</i>.</li> <li><b>unicast multicast</b>—Indicates <i>instance.inet.0</i> and <i>instance.inet.2</i>.</li> </ul>    | <b>detail</b>     |
| Import policy | ( <b>instance</b> keyword only) Policy that <b>route export</b> uses to construct the import-export matrix. Not displayed if the instance type is <b>vrf</b> .                                                                                                                                                                                                              | <b>detail</b>     |
| Instance      | ( <b>instance</b> keyword only) Name of the routing instance.                                                                                                                                                                                                                                                                                                               | <b>detail</b>     |
| Type          | ( <b>instance</b> keyword only) Type of routing instance: <b>forwarding</b> , <b>non-forwarding</b> , or <b>vrf</b> .                                                                                                                                                                                                                                                       | <b>detail</b>     |

## Sample Output

### show route export

```

user@host> show route export
Table      Export      Routes
inet.0     N            0
black.inet.0 Y           3
red.inet.0 Y            4

```

### show route export detail

```

user@host> show route export detail
inet.0                                Routes:      0
black.inet.0                          Routes:      3
  Import: [ inet.0 ]
red.inet.0                            Routes:      4
  Import: [ inet.0 ]

```

### show route export instance detail

```

user@host> show route export instance detail
Instance: master                      Type: forwarding
Flags: <config auto-policy> Options: <unicast multicast>
Import policy: [ (ospf-master-from-red || isis-master-from-black) ]

```

Instance: black  
Instance: red

Type: non-forwarding  
Type: non-forwarding

## show route extensive

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3604</a><br><a href="#">Syntax (EX Series Switches) on page 3604</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                      | show route extensive<br><destination-prefix><br><logical-system (all   logical-system-name)>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switches)</b> | show route extensive<br><destination-prefix>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                 | Display extensive information about the active entries in the routing tables.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                     | <b>none</b> —Display all active entries in the routing table.<br><br><b>destination-prefix</b> —(Optional) Display active entries for the specified address or range of addresses.<br><br><b>logical-system (all   logical-system-name)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>       | <a href="#">show route extensive on page 3611</a><br><a href="#">show route extensive (Access Route) on page 3617</a><br><a href="#">show route extensive (BGP PIC Edge) on page 3618</a><br><a href="#">show route extensive (FRR and LFA) on page 3618</a><br><a href="#">show route extensive (Route Reflector) on page 3619</a><br><a href="#">show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 3619</a><br><a href="#">show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 3620</a> |
| <b>Output Fields</b>               | <a href="#">Table 158 on page 1760</a> describes the output fields for the <b>show route extensive</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                              |

**Table 380: show route extensive Output Fields**

| Field Name                 | Field Description                                                       |
|----------------------------|-------------------------------------------------------------------------|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).                        |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table. |

Table 380: show route extensive Output Fields (*continued*)

| Field Name                                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>number routes</i>                           | <p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active).</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive).</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <i>route-destination</i><br>(entry, announced) | <p>Route destination (for example: 10.0.0.1/24). The <b>entry</b> value is the number of route for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul> |
| <b>TSI</b>                                     | Protocol header information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>label stacking</b>                          | <p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>[protocol, preference]</b>                  | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>                                                                                                                                                                                                                            |

Table 380: show route extensive Output Fields (*continued*)

| Field Name                                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Level</b>                                         | (IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Route Distinguisher</b>                           | IP subnet augmented with a 64-bit prefix.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>PMSI</b>                                          | Provider multicast service interface (MVPN routing table).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Next-hop type</b>                                 | Type of next hop. For a description of possible values for this field, see the Output Field table in the <a href="#">show route detail</a> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Next-hop reference count</b>                      | Number of references made to the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Flood nexthop branches exceed maximum message</b> | Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Source</b>                                        | IP address of the route source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Next hop</b>                                      | Network layer address of the directly reachable neighboring system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>via</b>                                           | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> </ul> |
| <b>Label-switched-path<br/>lsp-path-name</b>         | Name of the LSP used to reach the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Label operation</b>                               | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Offset</b>                                        | Whether the metric has been increased or decreased by an offset value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Interface</b>                                     | (Local only) Local interface name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Protocol next hop</b>                             | Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

Table 380: show route extensive Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b><i>label-operation</i></b> | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).                                                                                                                                                                                                                                                 |
| <b>Indirect next hops</b>     | <p>When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.</p> <p>When BGP PIC Edge is enabled, the output lines that contain <b>Indirect next hop: weight</b> follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values:</p> <ul style="list-style-type: none"> <li>• 0x1 indicates active next hops.</li> <li>• 0x4000 indicates passive next hops.</li> </ul> |
| <b>State</b>                  | State of the route (a route can be in more than one state). See the Output Field table in the <a href="#">show route detail</a> command.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Session ID</b>             | The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Weight</b>                 | <p>Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown.</p> <p>For sample output, see <a href="#">show route table</a>.</p>                                                                                                                                                                                                                                                                                                                                     |

Table 380: show route extensive Output Fields (*continued*)

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Inactive reason | <p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> <li>• <b>Active preferred</b>—Currently active route was selected over this route.</li> <li>• <b>Always compare MED</b>—Path with a lower multiple exit discriminator (MED) is available.</li> <li>• <b>AS path</b>—Shorter AS path is available.</li> <li>• <b>Cisco Non-deterministic MED selection</b>—Cisco nondeterministic MED is enabled and a path with a lower MED is available.</li> <li>• <b>Cluster list length</b>—Path with a shorter cluster list length is available.</li> <li>• <b>Forwarding use only</b>—Path is only available for forwarding purposes.</li> <li>• <b>IGP metric</b>—Path through the next hop with a lower IGP metric is available.</li> <li>• <b>IGP metric type</b>—Path with a lower OSPF link-state advertisement type is available.</li> <li>• <b>Interior &gt; Exterior &gt; Exterior via Interior</b>—Direct, static, IGP, or EBGP path is available.</li> <li>• <b>Local preference</b>—Path with a higher local preference value is available.</li> <li>• <b>Next hop address</b>—Path with a lower metric next hop is available.</li> <li>• <b>No difference</b>—Path from a neighbor with a lower IP address is available.</li> <li>• <b>Not Best in its group</b>—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed).</li> <li>• <b>Number of gateways</b>—Path with a higher number of next hops is available.</li> <li>• <b>Origin</b>—Path with a lower origin code is available.</li> <li>• <b>OSPF version</b>—Path does not support the indicated OSPF version.</li> <li>• <b>RIB preference</b>—Route from a higher-numbered routing table is available.</li> <li>• <b>Route distinguisher</b>—64-bit prefix added to IP subnets to make them unique.</li> <li>• <b>Route metric or MED comparison</b>—Route with a lower metric or MED is available.</li> <li>• <b>Route preference</b>—Route with a lower preference value is available.</li> <li>• <b>Router ID</b>—Path through a neighbor with a lower ID is available.</li> <li>• <b>Unusable path</b>—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved.</li> <li>• <b>Update source</b>—Last tiebreaker is the lowest IP address value.</li> </ul> |
| Local AS        | Autonomous system (AS) number of the local routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Age             | How long the route has been known.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| AIGP            | Accumulated interior gateway protocol (AIGP) BGP attribute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Metric          | Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| MED-plus-IGP    | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| TTL-Action      | <p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see <a href="#">show route table</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |



Table 380: show route extensive Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Task                    | Name of the protocol that has added the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Announcement bits       | List of protocols that announce this route. <b>n-Resolve inet</b> indicates that the route is used for route resolution for next hops found in the routing table. <b>n</b> is an index used by Juniper Networks customer support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| AS path                 | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| validation-state        | <p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Unverified</b>—Indicates that origin validation is not enabled for the BGP peers.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| FECs bound to route     | Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| AS path: I <Originator> | (For route reflected output only) Originator ID attribute set by the route reflector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Primary Upstream        | When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| RPF Nexthops            | When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 380: show route extensive Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Label                   | Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.                      |
| weight                  | Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.                                                                                                                                               |
| VC Label                | MPLS label assigned to the Layer 2 circuit virtual connection.                                                                                                                                                                                                                                                   |
| MTU                     | Maximum transmission unit (MTU) of the Layer 2 circuit.                                                                                                                                                                                                                                                          |
| VLAN ID                 | VLAN identifier of the Layer 2 circuit.                                                                                                                                                                                                                                                                          |
| Cluster list            | (For route reflected output only) Cluster ID sent by the route reflector.                                                                                                                                                                                                                                        |
| Originator ID           | (For route reflected output only) Address of router that originally sent the route to the route reflector.                                                                                                                                                                                                       |
| Prefixes bound to route | Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.                                                                                                                                                                                                               |
| Communities             | Community path attribute for the route. See the Output Field table in the <a href="#">show route detail</a> command for all possible values for this field.                                                                                                                                                      |
| Layer2-info: encaps     | Layer 2 encapsulation (for example, VPLS).                                                                                                                                                                                                                                                                       |
| control flags           | Control flags: <b>none</b> or Site Down.                                                                                                                                                                                                                                                                         |
| mtu                     | Maximum transmission unit (MTU) information.                                                                                                                                                                                                                                                                     |
| Label-Base, range       | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.                                                                                                                                           |
| status vector           | Layer 2 VPN and VPLS network layer reachability information (NLRI).                                                                                                                                                                                                                                              |
| Localpref               | Local preference value included in the route.                                                                                                                                                                                                                                                                    |
| Router ID               | BGP router ID as advertised by the neighbor in the open message.                                                                                                                                                                                                                                                 |
| Primary Routing Table   | In a routing table group, the name of the primary routing table in which the route resides.                                                                                                                                                                                                                      |
| Secondary Tables        | In a routing table group, the name of one or more secondary tables in which the route resides.                                                                                                                                                                                                                   |
| Originating RIB         | Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix. |
| Node path count         | Number of nodes in the path.                                                                                                                                                                                                                                                                                     |

Table 380: show route extensive Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                     |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Forwarding nexthops</b> | Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it. |

## Sample Output

### show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 29
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 69
        Age: 1:34:06
        Task: RT
        Announcement bits (2): 0-KRT 3-Resolve tree 2
        AS path: I

10.31.1.0/30 (2 entries, 1 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 2
        Next hop: via so-0/3/0.0, selected
        State: <Active Int>
        Local AS: 69
        Age: 1:32:40
        Task: IF
        Announcement bits (1): 3-Resolve tree 2
        AS path: I
    OSPF Preference: 10
        Next-hop reference count: 1
        Next hop: via so-0/3/0.0, selected
        State: <Int>
        Inactive reason: Route Preference
        Local AS: 69
        Age: 1:32:40 Metric: 1
        Area: 0.0.0.0
        Task: OSPF
        AS path: I

10.31.1.1/32 (1 entry, 1 announced)
    *Local Preference: 0
        Next hop type: Local
        Next-hop reference count: 7
        Interface: so-0/3/0.0
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:32:43
        Task: IF
        Announcement bits (1): 3-Resolve tree 2
        AS path: I

```

```
...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.31.2.0/30 -> {10.31.1.6}
    *OSPF   Preference: 10
            Next-hop reference count: 9
            Next hop: via so-0/3/0.0
            Next hop: 10.31.1.6 via ge-3/1/0.0, selected
            State: <Active Int>
            Local AS:    69
            Age: 1:32:19   Metric: 2
            Area: 0.0.0.0
            Task: OSPF
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kerne1 224.0.0.2/32 -> {}
    *PIM    Preference: 0
            Next-hop reference count: 18
            State: <Active NoReadvrt Int>
            Local AS:    69
            Age: 1:34:08
            Task: PIM Recv
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kerne1 224.0.0.22/32 -> {}
    *IGMP   Preference: 0
            Next-hop reference count: 18
            State: <Active NoReadvrt Int>
            Local AS:    69
            Age: 1:34:06
            Task: IGMP
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

...

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
    State: <FlashAll>
    *RSVP   Preference: 7
            Next-hop reference count: 6
            Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
            Label-switched-path green-r1-r3
            Label operation: Push 100096
            State: <Active Int>
            Local AS:    69
            Age: 1:28:12   Metric: 2
            Task: RSVP
            Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
            AS path: I
```

```

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r2
          State: <Active Int>
          Local AS: 69
          Age: 1:28:12    Metric: 1
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via lo0.0, selected
          State: <Active Int>
          Local AS: 69
          Age: 1:34:07
          Task: IF
          AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
  *MPLS   Preference: 0
          Next hop type: Receive
          Next-hop reference count: 6
          State: <Active Int>
          Local AS: 69
          Age: 1:34:08    Metric: 1
          Task: MPLS
          Announcement bits (1): 0-KRT
          AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299840 (1 entry, 1 announced)
TSI:
KRT in-kernel 299840 /52 -> {indirect(1048575)}
  *RSVP   Preference: 7/2
          Next hop type: Flood
          Address: 0x9174a30
          Next-hop reference count: 4
          Next hop type: Router, Next hop index: 798
          Address: 0x9174c28
          Next-hop reference count: 2
          Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
          Label-switched-path R2-to-R4-2p2mp

```

```

Label operation: Pop
Next hop type: Router, Next hop index: 1048574
Address: 0x92544f0
Next-hop reference count: 2
Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
Label-switched-path R2-to-R200-p2mp
Label operation: Pop
Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
Label operation: Pop
State: <Active Int>
Age: 1:29      Metric: 1
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I...

```

800010 (1 entry, 1 announced)

TSI:

```

KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: via vt-3/2/0.32769, selected
    Label operation: Pop
    State: <Active Int>
    Age: 1:31:53
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

```

vt-3/2/0.32769 (1 entry, 1 announced)

TSI:

```

KRT in-kernel vt-3/2/0.32769.0 /16 -> {indirect(1048574)}
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
    Label-switched-path green-r1-r3
    Label operation: Push 800012, Push 100096(top)
    Protocol next hop: 10.255.70.103
    Push 800012
    Indirect next hop: 87272e4 1048574
    State: <Active Int>
    Age: 1:31:53      Metric2: 2
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Indirect next hops: 1
      Protocol next hop: 10.255.70.103 Metric: 2
      Push 800012
      Indirect next hop: 87272e4 1048574
      Indirect path forwarding next hops: 1
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
        10.255.70.103/32 Originating RIB: inet.3
        Metric: 2      Node path count: 1
        Forwarding nexthops: 1
        Nexthop: 10.31.1.6 via ge-3/1/0.0

```

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)

```

*Direct Preference: 0
  Next hop type: Interface
  Next-hop reference count: 1
  Next hop: via lo0.0, selected
  State: <Active Int>
  Local AS: 69
  Age: 1:34:07
  Task: IF
  AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0
  Next hop type: Interface
  Next-hop reference count: 1
  Next hop: via lo0.0, selected
  State: <Active NoReadvrt Int>
  Local AS: 69
  Age: 1:34:07
  Task: IF
  AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
  *MLD Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:06
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0

```

```
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.16385, selected
State: <Active NoReadvrt Int>
Age: 1:34:07
Task: IF
AS path: I
```

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)

```
*BGP Preference: 170/-101
Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 1:28:12 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-green-l2vpn
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0
```

10.255.71.52:1:1:1/96 (1 entry, 1 announced)

TSI:

Page 0 idx 0 Type 1 val 8699540

```
*L2VPN Preference: 170/-1
Next-hop reference count: 5
Protocol next hop: 10.255.71.52
Indirect next hop: 0 -
State: <Active Int Ext>
Age: 1:34:03 Metric2: 1
Task: green-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8, status-vector: 0x9F
```

10.255.71.52:1:5:1/96 (1 entry, 1 announced)

TSI:

Page 0 idx 0 Type 1 val 8699528

```
*L2VPN Preference: 170/-101
Next-hop reference count: 5
Protocol next hop: 10.255.71.52
Indirect next hop: 0 -
State: <Active Int Ext>
Age: 1:34:03 Metric2: 1
Task: green-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
Label-base: 800008, range: 8, status-vector: 0x9F
```



...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)

```
*L2CKT Preference: 7
  Next hop: via so-1/1/2.0 weight 1, selected
  Label-switched-path my-lsp
  Label operation: Push 100000[0]
  Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
  State: <Active Int>
  Local AS: 99
  Age: 10:21
  Task: l2 circuit
  Announcement bits (1): 0-LDP
  AS path: I
  VC Label 100000, MTU 1500, VLAN ID 512
```

55.0.0.0/24 (1 entry, 1 announced)

TSI:

KRT queued (pending) add

55.0.0.0/24 -> {Push 300112}

```
*BGP Preference: 170/-101
  Next hop type: Router
  Address: 0x925c208
  Next-hop reference count: 2
  Source: 10.0.0.9
  Next hop: 10.0.0.9 via ge-1/2/0.15, selected
  Label operation: Push 300112
  Label TTL action: prop-ttl
  State: <Active Ext>
  Local AS: 7019 Peer AS: 13979
  Age: 1w0d 23:06:56
  AIGP: 25
  Task: BGP_13979.10.0.0.9+56732
  Announcement bits (1): 0-KRT
  AS path: 13979 7018 I
  Accepted
  Route Label: 300112
  Localpref: 100
  Router ID: 10.9.9.1
```

### show route extensive (Access Route)

user@host> show route 13.160.0.102 extensive

inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)

13.160.0.102/32 (1 entry, 1 announced)

TSI:

KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}

OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern

```
*Access Preference: 13
  Next-hop reference count: 78472
  Next hop: 13.160.0.2 via fe-0/0/0.0, selected
  State: <Active Int>
```

Age: 12

```
Task: RPD Unix Domain Server./var/run/rpd_serv.local
Announcement bits (2): 0-KRT 1-OSPFv2
AS path: I
```

### show route extensive (BGP PIC Edge)

```
user@host> show route 1.1.1.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  1.1.1.6/32 (3 entries, 2 announced)
    State: <CalcForwarding>
    TSI:
    KRT in-kerne 1.1.1.6/32 -> {indirect(1048574), indirect(1048577)}
    Page 0 idx 0 Type 1 val 9219e30
    Nexthop: Self
    AS path: [2] 3 I
    Communities: target:2:1
    Path 1.1.1.6 from 1.1.1.4 Vector len 4. Val: 0
  ..
    #Multipath Preference: 255
    Next hop type: Indirect
    Address: 0x93f4010
    Next-hop reference count: 2
  ..
    Protocol next hop: 1.1.1.4
    Push 299824
    Indirect next hop: 944c000 1048574 INH Session ID: 0x3
    Indirect next hop: weight 0x1
    Protocol next hop: 1.1.1.5
    Push 299824
    Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
    Indirect next hop: weight 0x4000
    State: <ForwardingOnly Int Ext>
    Inactive reason: Forwarding use only
    Age: 25 Metric2: 15
    Validation State: unverified
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: 3 I
    Communities: target:2:1
```

### show route extensive (FRR and LFA)

```
user@host> show route 20.31.2.0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
  20.31.2.0/24 (2 entries, 1 announced)
    State: FlashAll
    TSI:
    KRT in-kerne 20.31.2.0/24 -> {Push 299776, Push 299792}
    *RSVP Preference: 7/1
    Next hop type: Router, Next hop index: 1048574
    Address: 0xbbbc010
    Next-hop reference count: 5
    Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
    Label-switched-path europa-d-to-europa-e
    Label operation: Push 299776
    Label TTL action: prop-ttl
    Session Id: 0x201
    Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
    Label-switched-path europa-d-to-europa-e
    Label operation: Push 299792
    Label TTL action: prop-ttl
    Session Id: 0x202
    State: Active Int
    Local AS: 100
    Age: 5:31 Metric: 2
```

```

Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I

```

### show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
*BGP Preference: 170/-101
Source: 192.168.4.214
Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
State: <Active Int Ext>
Local AS: 10458 Peer AS: 10458
Age: 3:09 Metric: 0 Metric2: 0
Task: BGP_10458.192.168.4.214+1033
Announcement bits (2): 0-KRT 4-Resolve inet.0
AS path: 3944 7777 I <Originator>
Cluster list: 1.1.1.1
Originator ID: 10.255.245.88
Communities: 7777:7777
Localpref: 100
Router ID: 4.4.4.4
Indirect next hops: 1
    Protocol next hop: 207.17.136.192 Metric: 0
    Indirect next hop: 84ac908 40
    Indirect path forwarding next hops: 0
    Next hop type: Discard

```

### show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
*LDP Preference: 9
Next hop type: Flood
Next-hop reference count: 3
Address: 0x9097d90
Next hop: via vt-0/1/0.1
Next-hop index: 661
Label operation: Pop
Address: 0x9172130
Next hop: via so-0/0/3.0
Next-hop index: 654
Label operation: Swap 299872
State: **Active Int>

```

```
Local AS: 1001
Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2
```

#### show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```
mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
    Next hop type: Flood
    Address: 0x2735208
    Next-hop reference count: 3
    Next hop type: Router, Next hop index: 1397
    Address: 0x2735d2c
    Next-hop reference count: 3
    Next hop: 1.3.8.2 via ge-1/2/22.0
    Label operation: Pop
    Load balance label: None;
    Next hop type: Router, Next hop index: 1395
    Address: 0x2736290
    Next-hop reference count: 3
    Next hop: 1.3.4.2 via ge-1/2/18.0
    Label operation: Pop
    Load balance label: None;
    State: <Active Int AckRequest MulticastRPF>
    Local AS: 10
    Age: 54:05      Metric: 1
    Validation State: unverified
    Task: LDP
    Announcement bits (1): 0-KRT
    AS path: I
    FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
      Primary Upstream : 1.1.1.3:0--1.1.1.2:0
        RPF Nexthops :
          ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
          ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
      Backup Upstream : 1.1.1.3:0--1.1.1.6:0
        RPF Nexthops :
          ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
          ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe
```

## show route flow validation

**List of Syntax** [Syntax on page 3621](#)  
[Syntax \(EX Series Switches\) on page 3621](#)

**Syntax** show route flow validation  
 <brief | detail>  
 <ip-prefix>  
 <table *table-name*>  
 <logical-system (all | *logical-system-name*)>

**Syntax (EX Series Switches)** show route flow validation  
 <brief | detail>  
 <ip-prefix>  
 <table *table-name*>

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Display flow route information.

**Options** none—Display flow route information.

**brief | detail**—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

**ip-prefix**—(Optional) IP address for the flow route.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**table *table-name***—(Optional) Display flow route information for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route flow validation inet** command).

**Required Privilege Level** view

**List of Sample Output** [show route flow validation on page 3622](#)

**Output Fields** [Table 381 on page 3621](#) lists the output fields for the **show route flow validation** command. Output fields are listed in the approximate order in which they appear.

**Table 381: show route flow validation Output Fields**

| Field Name                | Field Description                                | Level of Output |
|---------------------------|--------------------------------------------------|-----------------|
| <i>routing-table-name</i> | Name of the routing table (for example, inet.0). | All levels      |
| <i>prefix</i>             | Route address.                                   | All levels      |
| Active unicast route      | Active route in the routing table.               | All levels      |

Table 381: show route flow validation Output Fields (*continued*)

| Field Name                  | Field Description                                                       | Level of Output |
|-----------------------------|-------------------------------------------------------------------------|-----------------|
| Dependent flow destinations | Number of flows for which there are routes in the routing table.        | All levels      |
| Origin                      | Source of the route flow.                                               | All levels      |
| Neighbor AS                 | Autonomous system identifier of the neighbor.                           | All levels      |
| Flow destination            | Number of entries and number of destinations that match the route flow. | All levels      |
| Unicast best match          | Destination that is the best match for the route flow.                  | All levels      |
| Flags                       | Information about the route flow.                                       | All levels      |

## Sample Output

### show route flow validation

```
user@host> show route flow validation
inet.0:
10.0.5.0/24Active unicast route
Dependent flow destinations: 1
Origin: 192.168.224.218, Neighbor AS: 65001
Flow destination (3 entries, 1 match origin)
Unicast best match: 10.0.5.0/24
Flags: SubtreeApex Consistent
```

## show route inactive-path

|                                    |                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3623</a><br><a href="#">Syntax (EX Series Switches) on page 3623</a>                                                                                                                                                                                                                                                                                              |
| <b>Syntax</b>                      | show route inactive-path<br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches)</b> | show route inactive-path<br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                        |
| <b>Description</b>                 | Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.                                                                                                                                                                                                                                                              |
| <b>Options</b>                     | <p><b>none</b>—Display all inactive routes.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>       | <a href="#">show route inactive-path on page 3623</a><br><a href="#">show route inactive-path detail on page 3624</a><br><a href="#">show route inactive-path extensive on page 3625</a><br><a href="#">show route inactive-path terse on page 3625</a>                                                                                                                                      |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                               |

## Sample Output

### show route inactive-path

```

user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30      [OSPF/10] 03:57:28, metric 1
> via so-0/3/0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```
10.0.0.0/8          [Direct/0] 04:39:56
                   > via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
                   AS path: 100 I
                   > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

#### show route inactive-path detail

```
user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
  OSPF   Preference: 10
         Next-hop reference count: 1
         Next hop: via so-0/3/0.0, selected
         State: <Int>
         Inactive reason: Route Preference
         Local AS: 1
         Age: 3:58:24   Metric: 1
         Area: 0.0.0.0
         Task: OSPF
         AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
  Direct Preference: 0
         Next hop type: Interface
         Next-hop reference count: 1
         Next hop: via fxp1.0, selected
         State: <NotBest Int>
         Inactive reason: No difference
         Age: 4:40:52
         Task: IF
         AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
```



```

BGP      Preference: 170/-101
        Next-hop reference count: 6
        Source: 10.12.80.1
        Next hop: 10.12.80.1 via ge-6/3/2.0, selected
        State: <Ext>
        Inactive reason: Route Preference
        Peer AS: 100
        Age: 4:39:13
        Task: BGP_100.10.12.80.1+179
        AS path: 100 I
        Localpref: 100
        Router ID: 10.0.0.0

```

### show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-path detail on page 3624](#).

### show route inactive-path terse

```

user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.12.100.12/30   0 10           1           >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.0.0.0/8        D  0           0           >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
  10.12.80.0/30     B 170          100          >10.12.80.1    100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

## show route inactive-prefix

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3626</a><br><a href="#">Syntax (EX Series Switches) on page 3626</a>                                                                                                                                                                                                                                                                                              |
| <b>Syntax</b>                      | <code>show route inactive-prefix</code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switches)</b> | <code>show route inactive-prefix</code><br><code>&lt;brief   detail   extensive   terse&gt;</code>                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                        |
| <b>Description</b>                 | Display inactive route destinations in each routing table.                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                     | <b>none</b> —Display all inactive route destination.<br><br><b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>       | <a href="#">show route inactive-prefix on page 3626</a><br><a href="#">show route inactive-prefix detail on page 3626</a><br><a href="#">show route inactive-prefix extensive on page 3627</a><br><a href="#">show route inactive-prefix terse on page 3627</a>                                                                                                                              |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                               |

## Sample Output

### show route inactive-prefix

```
user@host> show route inactive-prefix

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32          [Direct/0] 00:04:54
> via lo0.0
```

### show route inactive-prefix detail

```
user@host> show route inactive-prefix detail

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
```

```

127.0.0.1/32 (1 entry, 0 announced)
  Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Hidden Martian Int>
    Age: 4:51
    Task: IF
    AS path: I00:04:54
      > via lo0.0

```

### show route inactive-prefix extensive

The output for the **show route inactive-prefix extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-prefix detail on page 3626](#).

### show route inactive-prefix terse

```
user@host> show route inactive-prefix terse
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A Destination | P Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---------------|-------|----------|----------|----------|---------|
| 127.0.0.1/32  | D 0   |          |          | >lo0.0   |         |

## show route instance

---

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3628</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3628</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax</b>                                     | <pre>show route instance &lt;brief   detail   summary&gt; &lt;instance-name&gt; &lt;logical-system (all   logical-system-name)&gt; &lt;operational&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show route instance &lt;brief   detail   summary&gt; &lt;instance-name&gt; &lt;operational&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                | Display routing instance information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                    | <p><b>none</b>—(Same as <b>brief</b>) Display standard information about all routing instances.</p> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>. (These options are not available with the <b>operational</b> keyword.)</p> <p><b>instance-name</b>—(Optional) Display information for all routing instances whose name begins with this string (for example, <b>cust1</b>, <b>cust11</b>, and <b>cust111</b> are all displayed when you run the <b>show route instance cust1</b> command).</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>operational</b>—(Optional) Display operational routing instances.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <i>Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling</i></li><li>• <i>Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>                      | <a href="#">show route instance on page 3630</a><br><a href="#">show route instance detail (Graceful Restart Complete) on page 3630</a><br><a href="#">show route instance detail (Graceful Restart Incomplete) on page 3632</a><br><a href="#">show route instance detail (VPLS Routing Instance) on page 3633</a><br><a href="#">show route instance operational on page 3634</a><br><a href="#">show route instance summary on page 3634</a>                                                                                                                                                                                                                                                                                                                                                                                                        |

**Output Fields** Table 382 on page 3629 lists the output fields for the **show route instance** command. Output fields are listed in the approximate order in which they appear.

**Table 382: show route instance Output Fields**

| Field Name                       | Field Description                                                                                                                                                                                                                                                  | Level of Output                  |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|
| Instance or <i>instance-name</i> | Name of the routing instance.                                                                                                                                                                                                                                      | All levels                       |
| Operational Routing Instances    | ( <b>operational</b> keyword only) Names of all operational routing instances.                                                                                                                                                                                     | —                                |
| Type                             | Type of routing instance: <b>forwarding</b> , <b>l2vpn</b> , <b>no-forwarding</b> , <b>vpls</b> , <b>virtual-router</b> , or <b>vrf</b> .                                                                                                                          | All levels                       |
| State                            | State of the routing instance: <b>active</b> or <b>inactive</b> .                                                                                                                                                                                                  | <b>brief detail</b> none         |
| Interfaces                       | Name of interfaces belonging to this routing instance.                                                                                                                                                                                                             | <b>brief detail</b> none         |
| Restart State                    | Status of graceful restart for this instance: <b>Pending</b> or <b>Complete</b> .                                                                                                                                                                                  | <b>detail</b>                    |
| Path selection timeout           | Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is <b>300</b> .                                                                                                                                             | <b>detail</b>                    |
| Tables                           | Tables (and number of routes) associated with this routing instance.                                                                                                                                                                                               | <b>brief detail</b> none         |
| Route-distinguisher              | Unique route distinguisher associated with this routing instance.                                                                                                                                                                                                  | <b>detail</b>                    |
| Vrf-import                       | VPN routing and forwarding instance import policy name.                                                                                                                                                                                                            | <b>detail</b>                    |
| Vrf-export                       | VPN routing and forwarding instance export policy name.                                                                                                                                                                                                            | <b>detail</b>                    |
| Vrf-import-target                | VPN routing and forwarding instance import target community name.                                                                                                                                                                                                  | <b>detail</b>                    |
| Vrf-export-target                | VPN routing and forwarding instance export target community name.                                                                                                                                                                                                  | <b>detail</b>                    |
| Fast-reroute-priority            | Fast reroute priority setting for a VPLS routing instance: <b>high</b> , <b>medium</b> , or <b>low</b> . The default is <b>low</b> .                                                                                                                               | <b>detail</b>                    |
| Restart State                    | Restart state: <ul style="list-style-type: none"> <li><b>Pending;protocol-name</b>—List of protocols that have not yet completed graceful restart for this routing table.</li> <li><b>Complete</b>—All protocols have restarted for this routing table.</li> </ul> | <b>detail</b>                    |
| Primary rib                      | Primary table for this routing instance.                                                                                                                                                                                                                           | <b>brief</b> none <b>summary</b> |
| Active/holddown/hidden           | Number of active, hold-down, and hidden routes.                                                                                                                                                                                                                    | All levels                       |

## Sample Output

### show route instance

```

user@host> show route instance
Instance           Type
Primary RIB
master             forwarding
inet.0             16/0/1
iso.0               1/0/0
mpls.0             0/0/0
inet6.0            2/0/0
l2circuit.0        0/0/0
__juniper_private1__ forwarding
__juniper_private1__.inet.0 12/0/0
__juniper_private1__.inet6.0 1/0/0

```

### show route instance detail (Graceful Restart Complete)

```

user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding      State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0               : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3               : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0               : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0          : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0              : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0          : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
  BGP-INET:
    Router ID: 10.69.103.1
    Type: vrf            State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.103
    Route-distinguisher: 10.255.14.176:103
    Vrf-import: [ BGP-INET-import ]
    Vrf-export: [ BGP-INET-export ]
    Tables:
      BGP-INET.inet.0     : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
  BGP-L:
    Router ID: 10.69.104.1
    Type: vrf            State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.104
    Route-distinguisher: 10.255.14.176:104
    Vrf-import: [ BGP-L-import ]
    Vrf-export: [ BGP-L-export ]
    Tables:

```

```

BGP-L.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
BGP-L.mpls.0          : 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
L2VPN:
Router ID: 0.0.0.0
Type: l2vpn           State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.512
Route-distinguisher: 10.255.14.176:512
Vrf-import: [ L2VPN-import ]
Vrf-export: [ L2VPN-export ]
Tables:
  L2VPN.l2vpn.0        : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
LDP:
Router ID: 10.69.105.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.105
Route-distinguisher: 10.255.14.176:105
Vrf-import: [ LDP-import ]
Vrf-export: [ LDP-export ]
Tables:
  LDP.inet.0           : 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
OSPF:
Router ID: 10.69.101.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.101
Route-distinguisher: 10.255.14.176:101
Vrf-import: [ OSPF-import ]
Vrf-export: [ OSPF-export ]
Vrf-import-target: [ target:11111
Tables:
  OSPF.inet.0          : 8 routes (7 active, 0 holddown, 0 hidden)
Restart Complete
RIP:
Router ID: 10.69.102.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.102
Route-distinguisher: 10.255.14.176:102
Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
  RIP.inet.0           : 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
STATIC:
Router ID: 10.69.100.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.100
Route-distinguisher: 10.255.14.176:100
Vrf-import: [ STATIC-import ]

```

```
Vrf-export: [ STATIC-export ]
Tables:
  STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
  Restart Complete
```

#### show route instance detail (Graceful Restart Incomplete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding          State: Active
  Restart State: Pending    Path selection timeout: 300
  Tables:
    inet.0                  : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                  : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0             : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
  BGP-INET:
    Router ID: 10.69.103.1
    Type: vrf                State: Active
    Restart State: Pending    Path selection timeout: 300
    Interfaces:
      t3-0/0/0.103
    Route-distinguisher: 10.255.14.176:103
    Vrf-import: [ BGP-INET-import ]
    Vrf-export: [ BGP-INET-export ]
    Tables:
      BGP-INET.inet.0       : 6 routes (5 active, 0 holddown, 0 hidden)
      Restart Pending: VPN
  BGP-L:
    Router ID: 10.69.104.1
    Type: vrf                State: Active
    Restart State: Pending    Path selection timeout: 300
    Interfaces:
      t3-0/0/0.104
    Route-distinguisher: 10.255.14.176:104
    Vrf-import: [ BGP-L-import ]
    Vrf-export: [ BGP-L-export ]
    Tables:
      BGP-L.inet.0          : 6 routes (5 active, 0 holddown, 0 hidden)
      Restart Pending: VPN
      BGP-L.mpls.0          : 2 routes (2 active, 0 holddown, 0 hidden)
      Restart Pending: VPN
  L2VPN:
    Router ID: 0.0.0.0
    Type: l2vpn              State: Active
    Restart State: Pending    Path selection timeout: 300
    Interfaces:
      t3-0/0/0.512
    Route-distinguisher: 10.255.14.176:512
    Vrf-import: [ L2VPN-import ]
```



```

Vrf-export: [ L2VPN-export ]
Tables:
  L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
  Restart Pending: VPN L2VPN
LDP:
  Router ID: 10.69.105.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0           : 5 routes (4 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF LDP VPN
OSPF:
  Router ID: 10.69.101.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0          : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN
RIP:
  Router ID: 10.69.102.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0           : 8 routes (6 active, 2 holddown, 0 hidden)
    Restart Pending: RIP VPN
STATIC:
  Router ID: 10.69.100.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0        : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Pending: VPN

```

### show route instance detail (VPLS Routing Instance)

```

user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls              State: Active
  Interfaces:
    lsi.1048833

```

```

    lsi.1048832
    fe-0/1/0.513
    Route-distinguisher: 10.255.37.65:1
    Vrf-import: [ __vrf-import-test-vpls-internal__ ]
    Vrf-export: [ __vrf-export-test-vpls-internal__ ]
    Vrf-import-target: [ target:300:1 ]
    Vrf-export-target: [ target:300:1 ]
    Fast-reroute-priority: high
    Tables:
        test-vpls.l2vpn.0          : 3 routes (3 active, 0 holddown, 0 hidden)

```

### show route instance operational

```

user@host> show route instance operational
Operational Routing Instances:

```

```

master
default

```

### show route instance summary

```

user@host> show route instance summary

```

| Instance | Type       | Primary rib      | Active/holddown/hidden |
|----------|------------|------------------|------------------------|
| master   | forwarding |                  |                        |
|          |            | inet.0           | 15/0/1                 |
|          |            | iso.0            | 1/0/0                  |
|          |            | mpls.0           | 35/0/0                 |
|          |            | l3vpn.0          | 0/0/0                  |
|          |            | inet6.0          | 2/0/0                  |
|          |            | l2vpn.0          | 0/0/0                  |
|          |            | l2circuit.0      | 0/0/0                  |
| BGP-INET | vrf        |                  |                        |
|          |            | BGP-INET.inet.0  | 5/0/0                  |
|          |            | BGP-INET.iso.0   | 0/0/0                  |
|          |            | BGP-INET.inet6.0 | 0/0/0                  |
| BGP-L    | vrf        |                  |                        |
|          |            | BGP-L.inet.0     | 5/0/0                  |
|          |            | BGP-L.iso.0      | 0/0/0                  |
|          |            | BGP-L.mpls.0     | 4/0/0                  |
|          |            | BGP-L.inet6.0    | 0/0/0                  |
| L2VPN    | l2vpn      |                  |                        |
|          |            | L2VPN.inet.0     | 0/0/0                  |
|          |            | L2VPN.iso.0      | 0/0/0                  |
|          |            | L2VPN.inet6.0    | 0/0/0                  |
|          |            | L2VPN.l2vpn.0    | 2/0/0                  |
| LDP      | vrf        |                  |                        |
|          |            | LDP.inet.0       | 4/0/0                  |
|          |            | LDP.iso.0        | 0/0/0                  |
|          |            | LDP.mpls.0       | 0/0/0                  |
|          |            | LDP.inet6.0      | 0/0/0                  |
|          |            | LDP.l2circuit.0  | 0/0/0                  |
| OSPF     | vrf        |                  |                        |
|          |            | OSPF.inet.0      | 7/0/0                  |
|          |            | OSPF.iso.0       | 0/0/0                  |
|          |            | OSPF.inet6.0     | 0/0/0                  |
| RIP      | vrf        |                  |                        |
|          |            | RIP.inet.0       | 6/0/0                  |
|          |            | RIP.iso.0        | 0/0/0                  |
|          |            | RIP.inet6.0      | 0/0/0                  |
| STATIC   | vrf        |                  |                        |
|          |            | STATIC.inet.0    | 4/0/0                  |

|                |       |
|----------------|-------|
| STATIC.iso.0   | 0/0/0 |
| STATIC.inet6.0 | 0/0/0 |

## show route label

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3636</a><br><a href="#">Syntax (EX Series Switches) on page 3636</a>                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax</b>                      | <code>show route label <i>label</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                                                          |
| <b>Syntax (EX Series Switches)</b> | <code>show route label <i>label</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code>                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>                 | Display the routes based on a specified Multiprotocol Label Switching (MPLS) label value.                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                     | <b><i>label</i></b> —Value of the MPLS label.<br><br><b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.                                                                    |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs</a></li></ul>                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>       | <a href="#">show route label terse on page 3636</a><br><a href="#">show route label on page 3637</a><br><a href="#">show route label detail on page 3637</a><br><a href="#">show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 3637</a><br><a href="#">show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 3638</a><br><a href="#">show route label extensive on page 3638</a> |
| <b>Output Fields</b>               | For information about output fields, see the output field table for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                                                                            |

## Sample Output

### show route label terse

```
user@host> show route label 100016 terse

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
```

| A Destination | P Prf | Metric 1 | Metric 2 | Next hop    | AS path |
|---------------|-------|----------|----------|-------------|---------|
| * 100016      | V 170 |          |          | >10.12.80.1 |         |

### show route label

```
user@host> show route label 100016
```

```
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
100016          *[VPN/170] 03:25:41
                > to 10.12.80.1 via ge-6/3/2.0, Pop
```

### show route label detail

```
user@host> show route label 100016 detail
```

```
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
100016 (1 entry, 1 announced)
    *VPN      Preference: 170
              Next-hop reference count: 2
              Source: 10.12.80.1
              Next hop: 10.12.80.1 via ge-6/3/2.0, selected
              Label operation: Pop
              State: <Active Int Ext>
              Local AS: 1
              Age: 3:23:31
              Task: BGP.0.0.0.0+179
              Announcement bits (1): 0-KRT
              AS path: 100 I
              Ref Cnt: 2
```

### show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
```

```
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
    *LDP      Preference: 9
              Next hop type: Flood
              Next-hop reference count: 3
              Address: 0x9097d90
              Next hop: via vt-0/1/0.1
              Next-hop index: 661
              Label operation: Pop
              Address: 0x9172130
              Next hop: via so-0/0/3.0
              Next-hop index: 654
              Label operation: Swap 299872
              State: **Active Int>
              Local AS: 1001
              Age: 8:20      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2
```

**show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)**

```
user@host> show route label 301568 detail
```

```
mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
          Next hop type: Flood
          Address: 0x2735208
          Next-hop reference count: 3
          Next hop type: Router, Next hop index: 1397
          Address: 0x2735d2c
          Next-hop reference count: 3
          Next hop: 1.3.8.2 via ge-1/2/22.0
          Label operation: Pop
          Load balance label: None;
          Next hop type: Router, Next hop index: 1395
          Address: 0x2736290
          Next-hop reference count: 3
          Next hop: 1.3.4.2 via ge-1/2/18.0
          Label operation: Pop
          Load balance label: None;
          State: <Active Int AckRequest MulticastRPF>
          Local AS: 10
          Age: 54:05      Metric: 1
          Validation State: unverified
          Task: LDP
          Announcement bits (1): 0-KRT
          AS path: I
          FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
          Primary Upstream : 1.1.1.3:0--1.1.1.2:0
          RPF Nexthops :
            ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
            ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
          Backup Upstream : 1.1.1.3:0--1.1.1.6:0
          RPF Nexthops :
            ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
            ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe
```

**show route label extensive**

The output for the **show route label extensive** command is identical to that of the **show route label detail** command. For sample output, see [show route label detail on page 3637](#).

## show route label-switched-path

|                                    |                                                                                                                                                                                                                                                                                                                              |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3639</a><br><a href="#">Syntax (EX Series Switches) on page 3639</a>                                                                                                                                                                                                                              |
| <b>Syntax</b>                      | <code>show route label-switched-path <i>path-name</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                            |
| <b>Syntax (EX Series Switches)</b> | <code>show route label-switched-path <i>path-name</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code>                                                                                                                                                                                                      |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.                                                                                                                                                                                                        |
| <b>Description</b>                 | Display the routes used in an MPLS label-switched path (LSP).                                                                                                                                                                                                                                                                |
| <b>Options</b>                     | <code>brief   detail   extensive   terse</code> —(Optional) Display the specified level of output.<br><br><code><i>path-name</i></code> —LSP tunnel name.<br><br><code>logical-system (all   <i>logical-system-name</i>)</code> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>       | <a href="#">show route label-switched-path on page 3639</a>                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                               |

## Sample Output

### show route label-switched-path

```

user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32          *[MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32          *[MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
4.4.4.4/32          *[MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path abc
> to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

```

```
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0  
                   > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, * = Both
```

```
mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, * = Both
```



## show route martians

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax              | <a href="#">Syntax on page 3641</a><br><a href="#">Syntax (EX Series Switches) on page 3641</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Syntax                      | <pre>show route martians &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;table <i>routing-table-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Syntax (EX Series Switches) | <pre>show route martians &lt;table <i>routing-table-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Release Information         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Description                 | Display the martian (invalid and ignored) entries associated with each routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Options                     | <p><b>none</b>—Display standard information about route martians for all routing tables.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>table <i>routing-table-name</i></b>—(Optional) Display information about route martians for all routing tables whose name begins with this string (for example, <b>inet.0</b> and <b>inet6.0</b> are both displayed when you run the <b>show route martians table inet</b> command).</p> |
| Required Privilege Level    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Related Documentation       | <ul style="list-style-type: none"> <li><a href="#">Example: Configuring Martian Addresses</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| List of Sample Output       | <a href="#">show route martians on page 3642</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Output Fields               | <p><a href="#">Table 383 on page 3641</a> lists the output fields for the <b>show route martians</b> command. Output fields are listed in the approximate order in which they appear</p>                                                                                                                                                                                                                                                                                                                                                                |

**Table 383: show route martians Output Fields**

| Field Name                | Field Description                                           |
|---------------------------|-------------------------------------------------------------|
| <i>table-name</i>         | Name of the route table in which the route martians reside. |
| <i>destination-prefix</i> | Route destination.                                          |
| <i>match value</i>        | Route match parameter.                                      |
| <i>status</i>             | Status of the route: <b>allowed</b> or <b>disallowed</b> .  |

## Sample Output

### show route martians

```
user@host> show route martians

inet.0:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.1:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed

inet.2:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.3:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

...

inet6.0:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.1:
    ::1/128 exact -- disallowed

inet6.2:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.3:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

...
```

## show route next-hop

|                                    |                                                                                                                                                                                                                                                                                                            |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3643</a><br><a href="#">Syntax (EX Series Switches) on page 3643</a>                                                                                                                                                                                                            |
| <b>Syntax</b>                      | <b>show route next-hop</b> <i>next-hop</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                 |
| <b>Syntax (EX Series Switches)</b> | <b>show route next-hop</b> <i>next-hop</i><br><brief   detail   extensive   terse>                                                                                                                                                                                                                         |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                      |
| <b>Description</b>                 | Display the entries in the routing table that are being sent to the specified next-hop address.                                                                                                                                                                                                            |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>next-hop</i></b> —Next-hop address. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>       | <a href="#">show route next-hop on page 3643</a><br><a href="#">show route next-hop detail on page 3644</a><br><a href="#">show route next-hop extensive on page 3646</a><br><a href="#">show route next-hop terse on page 3647</a>                                                                        |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                             |

## Sample Output

### show route next-hop

```

user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
172.16.0.0/12    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
192.168.0.0/16   *[Static/5] 06:26:25

```

```

> to 192.168.71.254 via fxp0.0
192.168.102.0/23  *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0
207.17.136.0/24  *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0
207.17.136.192/32 *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

#### show route next-hop detail

```

user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
```

```

AS path: I

192.168.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

192.168.102.0/23 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.192/32 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

## show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
```

```
10.10.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
10.209.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
172.16.0.0/12 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.102.0/23 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```

Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:02:28
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:02:28
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

### show route next-hop terse

```

user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.10.0.0/16     S  5         0         0  >192.168.71.254
* 10.209.0.0/16    S  5         0         0  >192.168.71.254
* 172.16.0.0/12    S  5         0         0  >192.168.71.254

```

```
* 192.168.0.0/16      S   5                >192.168.71.254
* 192.168.102.0/23   S   5                >192.168.71.254
* 207.17.136.0/24    S   5                >192.168.71.254
* 207.17.136.192/32 S   5                >192.168.71.254

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```



## show route no-community

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3649</a><br><a href="#">Syntax (EX Series Switches) on page 3649</a>                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                      | show route no-community<br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switches)</b> | show route no-community<br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                |
| <b>Description</b>                 | Display the route entries in each routing table that are not associated with any community.                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                     | <p><b>none</b>—(Same as <b>brief</b>) Display the route entries in each routing table that are not associated with any community.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>       | <a href="#">show route no-community on page 3649</a><br><a href="#">show route no-community detail on page 3650</a><br><a href="#">show route no-community extensive on page 3650</a><br><a href="#">show route no-community terse on page 3651</a>                                                                                                                                                  |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                       |

## Sample Output

### show route no-community

```

user@host> show route no-community
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:36:27
> to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:36:27
> to 192.168.71.254 via fxp0.0
10.255.71.52/32  *[Direct/0] 00:36:27
> via lo0.0
10.255.71.63/32  *[OSPF/10] 00:04:39, metric 1
> to 35.1.1.2 via ge-3/1/0.0
10.255.71.64/32  *[OSPF/10] 00:00:08, metric 2

```

```

> to 35.1.1.2 via ge-3/1/0.0
10.255.71.240/32  * [OSPF/10] 00:05:04, metric 2
                  via so-0/1/2.0
> via so-0/3/2.0
10.255.71.241/32  * [OSPF/10] 00:05:14, metric 1
> via so-0/1/2.0
10.255.71.242/32  * [OSPF/10] 00:05:19, metric 1
> via so-0/3/2.0
12.1.1.0/24       * [OSPF/10] 00:05:14, metric 2
> via so-0/3/2.0
14.1.1.0/24       * [OSPF/10] 00:00:08, metric 3
> to 35.1.1.2 via ge-3/1/0.0
                  via so-0/1/2.0
                  via so-0/3/2.0
16.1.1.0/24       * [OSPF/10] 00:05:14, metric 2
> via so-0/1/2.0
.....

```

### show route no-community detail

```
user@host> show route no-community detail
```

```

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

....

```

### show route no-community extensive

```
user@host> show route no-community extensive
```

```

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

```

```

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

```

### show route no-community terse

```
user@host> show route no-community terse
```

```

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

```

| A   | Destination      | P | Prf | Metric 1 | Metric 2 | Next hop        | AS path |
|-----|------------------|---|-----|----------|----------|-----------------|---------|
| *   | 10.10.0.0/16     | S | 5   |          |          | >192.168.71.254 |         |
| *   | 10.209.0.0/16    | S | 5   |          |          | >192.168.71.254 |         |
| *   | 10.255.71.52/32  | D | 0   |          |          | >lo0.0          |         |
| *   | 10.255.71.63/32  | 0 | 10  | 1        |          | >35.1.1.2       |         |
| *   | 10.255.71.64/32  | 0 | 10  | 2        |          | >35.1.1.2       |         |
| *   | 10.255.71.240/32 | 0 | 10  | 2        |          | so-0/1/2.0      |         |
|     |                  |   |     |          |          | >so-0/3/2.0     |         |
| *   | 10.255.71.241/32 | 0 | 10  | 1        |          | >so-0/1/2.0     |         |
| *   | 10.255.71.242/32 | 0 | 10  | 1        |          | >so-0/3/2.0     |         |
| *   | 12.1.1.0/24      | 0 | 10  | 2        |          | >so-0/3/2.0     |         |
| *   | 14.1.1.0/24      | 0 | 10  | 3        |          | >35.1.1.2       |         |
|     |                  |   |     |          |          | so-0/1/2.0      |         |
|     |                  |   |     |          |          | so-0/3/2.0      |         |
| *   | 16.1.1.0/24      | 0 | 10  | 2        |          | >so-0/1/2.0     |         |
| ... |                  |   |     |          |          |                 |         |

## show route protocol

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3652</a><br><a href="#">Syntax (EX Series Switches) on page 3652</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                      | <code>show route protocol <i>protocol</i></code><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switches)</b> | <code>show route protocol <i>protocol</i></code><br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>ospf2</b> and <b>ospf3</b> options introduced in Junos OS Release 9.2.<br><b>ospf2</b> and <b>ospf3</b> options introduced in Junos OS Release 9.2 for EX Series switches.<br><b>flow</b> option introduced in Junos OS Release 10.0.<br><b>flow</b> option introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                 | Display the route entries in the routing table that were learned from a particular protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>protocol</i></b> —Protocol from which the route was learned: <ul style="list-style-type: none"><li>• <b>access</b>—Access route for use by DHCP application</li><li>• <b>access-internal</b>—Access-internal route for use by DHCP application</li><li>• <b>aggregate</b>—Locally generated aggregate route</li><li>• <b>arp</b>—Route learned through the Address Resolution Protocol</li><li>• <b>atmvpn</b>—Asynchronous Transfer Mode virtual private network</li><li>• <b>bgp</b>—Border Gateway Protocol</li><li>• <b>ccc</b>—Circuit cross-connect</li><li>• <b>direct</b>—Directly connected route</li><li>• <b>dvmrp</b>—Distance Vector Multicast Routing Protocol</li><li>• <b>esis</b>—End System-to-Intermediate System</li><li>• <b>flow</b>—Locally defined flow-specification route</li><li>• <b>frr</b>—Precomputed protection route or backup route used when a link goes down</li><li>• <b>isis</b>—Intermediate System-to-Intermediate System</li><li>• <b>ldp</b>—Label Distribution Protocol</li><li>• <b>l2circuit</b>—Layer 2 circuit</li></ul> |

- **l2vpn**—Layer 2 virtual private network
- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



**NOTE:** EX Series switches run a subset of these protocols. See the switch CLI for details.

**Required Privilege Level** view

**Related Documentation** • *MPLS Feature Support on Juniper Switches*

**List of Sample Output**

- [show route protocol access on page 3654](#)
- [show route protocol access-internal extensive on page 3654](#)
- [show route protocol arp on page 3654](#)
- [show route protocol bgp on page 3655](#)
- [show route protocol bgp detail on page 3655](#)
- [show route protocol bgp extensive on page 3655](#)
- [show route protocol bgp terse on page 3656](#)
- [show route protocol direct on page 3656](#)
- [show route protocol frr on page 3657](#)
- [show route protocol l2circuit detail on page 3657](#)
- [show route protocol l2vpn extensive on page 3658](#)
- [show route protocol ldp on page 3659](#)
- [show route protocol ldp extensive on page 3659](#)
- [show route protocol ospf \(Layer 3 VPN\) on page 3660](#)
- [show route protocol ospf detail on page 3661](#)

[show route protocol rip on page 3661](#)  
[show route protocol rip detail on page 3661](#)  
[show route protocol ripng table inet6 on page 3662](#)  
[show route protocol static detail on page 3662](#)

**Output Fields** For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

## Sample Output

### [show route protocol access](#)

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
```

### [show route protocol access-internal extensive](#)

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
    *Access-internal Preference: 12
        Next-hop reference count: 200000
        Next hop: 13.160.0.2 via fe-0/0/0.0, selected
        State: <Active Int>
    Age: 36
        Task: RPD Unix Domain Server./var/run/rpd_serv.local
        Announcement bits (1): 0-KRT
        AS path: I
```

### [show route protocol arp](#)

```
user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.4/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.5/32      [ARP/4294967293] 00:04:32, from 20.20.1.1
                  Unusable
20.20.1.6/32      [ARP/4294967293] 00:04:34, from 20.20.1.1
                  Unusable
20.20.1.7/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
```

```

20.20.1.8/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.9/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.10/32     [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.11/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
20.20.1.12/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
20.20.1.13/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
...

```

### show route protocol bgp

```

user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21    *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                  AS path: 10458 14203 2914 4788 4788 I
                  > to 192.168.167.254 via fxp0.0

```

### show route protocol bgp detail

```

user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24    (1 entry, 1 announced)
   *BGP           Preference: 170/-101
                   Next hop type: Indirect
                   Next-hop reference count: 1006436
                   Source: 192.168.69.71
                   Next hop type: Router, Next hop index: 324
                   Next hop: 192.168.167.254 via fxp0.0, selected
                   Protocol next hop: 192.168.69.71
                   Indirect next hop: 8e166c0 342
                   State: <Active Ext>
                   Local AS: 69 Peer AS: 10458
                   Age: 6d 10:42:42 Metric2: 0
                   Task: BGP_10458.192.168.69.71+179
                   Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1

   AS path: 10458 14203 2914 4788 4788 I
   Communities: 2914:410 2914:2403 2914:3400
   Accepted
   Localpref: 100
   Router ID: 207.17.136.192

```

### show route protocol bgp extensive

```

user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
  Nexthop: Self
    AS path: [69] 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400
  Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1

```

```

*BGP      Preference: 170/-101
          Next hop type: Indirect
          Next-hop reference count: 1006502
          Source: 192.168.69.71
          Next hop type: Router, Next hop index: 324
          Next hop: 192.168.167.254 via fxp0.0, selected
          Protocol next hop: 192.168.69.71
          Indirect next hop: 8e166c0 342
          State: <Active Ext>
          Local AS: 69 Peer AS: 10458
          Age: 6d 10:44:45 Metric2: 0
          Task: BGP_10458.192.168.69.71+179
          Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1
  AS path: 10458 14203 2914 4788 4788 I
  Communities: 2914:410 2914:2403 2914:3400
  Accepted
  Localpref: 100
  Router ID: 207.17.136.192
  Indirect next hops: 1
    Protocol next hop: 192.168.69.71
    Indirect next hop: 8e166c0 342
    Indirect path forwarding next hops: 1
      Next hop type: Router
      Next hop: 192.168.167.254 via fxp0.0
    192.168.0.0/16 Originating RIB: inet.0
    Node path count: 1
    Forwarding nexthops: 1
      Nexthop: 192.168.167.254 via fxp0.0

```

### show route protocol bgp terse

```

user@host> show route protocol bgp 192.168.64.0/21 terse

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
192.168.64.0/21   B 170      100          >100.1.3.2    10023 21 I

```

### show route protocol direct

```

user@host> show route protocol direct

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

8.8.8.0/24         *[Direct/0] 17w0d 10:31:49
> via fe-1/3/1.0
10.255.165.1/32    *[Direct/0] 25w4d 04:13:18
> via lo0.0
30.30.30.0/24      *[Direct/0] 17w0d 23:06:26
> via fe-1/3/2.0
192.168.164.0/22   *[Direct/0] 25w4d 04:13:20
> via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
*[Direct/0] 25w4d 04:13:21

```



```

> via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

abcd::10:255:165:1/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0

```

### show route protocol frr

```

user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.3 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.4 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32      *[FRR/200] 00:05:35, from 20.20.1.1
                  > to 20.20.1.5 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32      *[FRR/200] 00:05:37, from 20.20.1.1
                  > to 20.20.1.6 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.7 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.8 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.9 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32     *[FRR/200] 00:05:38, from 20.20.1.1
...

```

### show route protocol l2circuit detail

```

user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop: via ge-2/0/0.0, selected
        Label operation: Pop      Offset: 4
        State: <Active Int>
        Local AS: 99
        Age: 9:52
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

```

```
ge-2/0/0.0 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000, Push 100000(top)[0] Offset: -4
    Protocol next hop: 10.245.255.63
    Push 100000 Offset: -4
    Indirect next hop: 86af0c0 298
    State: <Active Int>
    Local AS: 99
    Age: 9:52
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512
```

### show route protocol l2vpn extensive

```
user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
  *L2VPN Preference: 7
    Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
    Label operation: Pop Offset: 4
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
  *L2VPN Preference: 7
    Next hop: via so-0/0/1.0, selected
```

```

Label operation: Push 800000 Offset: -4
Protocol next hop: 10.255.14.220
Push 800000 Offset: -4
  Indirect next hop: 85142a0 288
State: <Active Int>
Local AS: 69
Age: 7:48
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I
Communities: target:69:1 Layer2-info: encaps:PPP,
control flags:2, mtu: 0

```

### show route protocol ldp

```

user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100080            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Swap 100000

```

### show route protocol ldp extensive

```

user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          Label operation: Push 100000
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

192.168.17.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>
          Local AS: 65500

```

```

Age: 1d 23:03:58      Metric: 1
Task: LDP
Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
AS path: I

private1__inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              Label operation: Pop
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              Label operation: Swap 100000
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 192.168.16.1/32

```

### show route protocol ospf (Layer 3 VPN)

```

user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30      * [OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.39.1.8/30      [OSPF/10] 00:05:18, metric 2

```

```

> via t3-3/2/0.0
10.255.14.171/32  *[OSPF/10] 00:05:18, metric 4
> via t3-3/2/0.0
10.255.14.179/32  *[OSPF/10] 00:05:18, metric 2
> via t3-3/2/0.0
224.0.0.5/32     *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30    [OSPF/10] 00:05:43, metric 1
> via so-0/2/2.0
10.255.14.173/32 *[OSPF/10] 00:05:43, metric 1
> via so-0/2/2.0
224.0.0.5/32    *[OSPF/10] 20:26:20, metric 1

```

### show route protocol ospf detail

```

user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
  OSPF   Preference: 10
         Nexthop: via so-0/2/2.0, selected
         State: <Int>
         Inactive reason: Route Preference
         Age: 6:25      Metric: 1
         Area: 0.0.0.0
         Task: VPN-AB-OSPF
         AS path: I
         Communities: Route-Type:0.0.0.0:1:0

...

```

### show route protocol rip

```

user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32  *[RIP/100] 20:24:34, metric 2
> to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32     *[RIP/100] 00:03:59, metric 1

```

### show route protocol rip detail

```

user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
  *RIP   Preference: 100
         Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
         State: <Active Int>
         Age: 20:25:02  Metric: 2
         Task: VPN-AB-RIPv2
         Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179

```

AS path: I  
Route learned from 10.39.1.22 expires in 96 seconds

### show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

### show route protocol static detail

```
user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.13.10.0/23 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
```

State: <Active NoReadvrt Int Ext>  
Age: 7w3d 21:24:25  
Validation State: unverified  
Task: RT  
Announcement bits (1): 0-KRT  
AS path: I

## show route range

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3664</a><br><a href="#">Syntax (EX Series Switches) on page 3664</a>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax</b>                      | <code>show route range</code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;destination-prefix&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (EX Series Switches)</b> | <code>show route range</code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;destination-prefix&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                 | Display routing table entries using a prefix range.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                     | <b>none</b> —Display standard information about all routing table entries using a prefix range.<br><br><b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b>destination-prefix</b> —Destination and prefix mask for the range.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>       | <a href="#">show route range on page 3664</a><br><a href="#">show route range destination-prefix on page 3665</a><br><a href="#">show route range detail on page 3665</a><br><a href="#">show route range extensive on page 3666</a><br><a href="#">show route range terse on page 3667</a>                                                                                                                                                                                                                                  |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                                                                                                                                               |

## Sample Output

### show route range

```
user@host> show route range

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:30:01
```



```

10.255.71.14/32      > to 192.168.71.254 via fxp0.0
                    *[Direct/0] 00:30:01
                    > via lo0.0
172.16.0.0/12       *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21     *[Direct/0] 00:30:01
                    > via fxp0.0
192.168.71.14/32    *[Local/0] 00:30:01
                    Local via fxp0.0
192.168.102.0/23    *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
...

```

### show route range destination-prefix

```

user@host> show route range 192.168.0.0/16

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16      *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21     *[Direct/0] 00:31:14
                    > via fxp0.0
192.168.71.14/32    *[Local/0] 00:31:14
                    Local via fxp0.0
192.168.102.0/23    *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0

```

### show route range detail

```

user@host> show route range detail

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected

```

```
State: <Active Int>
Age: 30:05
Task: IF
AS path: I

172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

...
```

### show route range extensive

```
user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Age: 30:17
    Task: IF
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
```

```

Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Age: 30:17
Task: RT
Announcement bits (1): 0-KRT
AS path: I

```

```
...
```

### show route range terse

```
user@host> show route range terse
```

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A | Destination       | P | Prf | Metric 1 | Metric 2 | Next hop        | AS path |
|---|-------------------|---|-----|----------|----------|-----------------|---------|
| * | 10.10.0.0/16      | S | 5   |          |          | >192.168.71.254 |         |
| * | 10.209.0.0/16     | S | 5   |          |          | >192.168.71.254 |         |
| * | 10.255.71.14/32   | D | 0   |          |          | >lo0.0          |         |
| * | 172.16.0.0/12     | S | 5   |          |          | >192.168.71.254 |         |
| * | 192.168.0.0/16    | S | 5   |          |          | >192.168.71.254 |         |
| * | 192.168.64.0/21   | D | 0   |          |          | >fxp0.0         |         |
| * | 192.168.71.14/32  | L | 0   |          |          | Local           |         |
| * | 192.168.102.0/23  | S | 5   |          |          | >192.168.71.254 |         |
| * | 207.17.136.0/24   | S | 5   |          |          | >192.168.71.254 |         |
| * | 207.17.136.192/32 | S | 5   |          |          | >192.168.71.254 |         |

```

__juniper_private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

```

| A | Destination | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|-------------|---|-----|----------|----------|----------|---------|
| * | 10.0.0.0/8  | D | 0   |          |          | >fxp2.0  |         |
|   |             | D | 0   |          |          | >fxp1.0  |         |
| * | 10.0.0.4/32 | L | 0   |          |          | Local    |         |

```

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

| A | Destination                                         | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|-----------------------------------------------------|---|-----|----------|----------|----------|---------|
|   | 47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152 |   |     |          |          |          |         |
| * |                                                     | D | 0   |          |          | >lo0.0   |         |

```

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

| A | Destination                  | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|------------------------------|---|-----|----------|----------|----------|---------|
|   | abcd::10:255:71:14/128       |   |     |          |          |          |         |
| * |                              | D | 0   |          |          | >lo0.0   |         |
|   | fe80::280:42ff:fe11:226f/128 |   |     |          |          |          |         |
| * |                              | D | 0   |          |          | >lo0.0   |         |

```

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

```

| A | Destination                  | P | Prf | Metric 1 | Metric 2 | Next hop   | AS path |
|---|------------------------------|---|-----|----------|----------|------------|---------|
|   | fe80::280:42ff:fe11:226f/128 |   |     |          |          |            |         |
| * |                              | D | 0   |          |          | >lo0.16385 |         |

## show route receive-protocol

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3668</a><br><a href="#">Syntax (EX Series Switches) on page 3668</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax</b>                      | show route receive-protocol <i>protocol neighbor-address</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switches)</b> | show route receive-protocol <i>protocol neighbor-address</i><br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>                 | Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>protocol neighbor-address</i></b> —Protocol transmitting the route ( <b>bgp</b> , <b>dvmrp</b> , <b>msdp</b> , <b>pim</b> , <b>rip</b> , or <b>ripng</b> ) and address of the neighboring router from which the route entry was received.                                                                                                                                                                                                                                        |
| <b>Additional Information</b>      | The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>       | <a href="#">show route receive-protocol bgp on page 3671</a><br><a href="#">show route receive-protocol bgp extensive on page 3671</a><br><a href="#">show route receive-protocol bgp table extensive on page 3671</a><br><a href="#">show route receive-protocol bgp logical-system extensive on page 3672</a><br><a href="#">show route receive-protocol bgp detail (Layer 2 VPN) on page 3673</a><br><a href="#">show route receive-protocol bgp extensive (Layer 2 VPN) on page 3673</a><br><a href="#">show route receive-protocol bgp (Layer 3 VPN) on page 3674</a><br><a href="#">show route receive-protocol bgp detail (Layer 3 VPN) on page 3674</a><br><a href="#">show route receive-protocol bgp extensive (Layer 3 VPN) on page 3675</a> |
| <b>Output Fields</b>               | <a href="#">Table 384 on page 3668</a> describes the output fields for the <b>show route receive-protocol</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

**Table 384: show route receive-protocol Output Fields**

| Field Name                | Field Description                              | Level of Output |
|---------------------------|------------------------------------------------|-----------------|
| <i>routing-table-name</i> | Name of the routing table—for example, inet.0. | All levels      |

Table 384: show route receive-protocol Output Fields (*continued*)

| Field Name                                   | Field Description                                                                                                                                                                                                                                                                                                                               | Level of Output         |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <i>number destinations</i>                   | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                         | All levels              |
| <i>number routes</i>                         | <p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> <li>• <b>active</b></li> <li>• <b>holddown</b> (routes that are in pending state before being declared inactive)</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy)</li> </ul> | All levels              |
| Prefix                                       | Destination prefix.                                                                                                                                                                                                                                                                                                                             | none <b>brief</b>       |
| MED                                          | Multiple exit discriminator value included in the route.                                                                                                                                                                                                                                                                                        | none <b>brief</b>       |
| <i>destination-prefix (entry, announced)</i> | Destination prefix. The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination.                                                                                                                                                           | <b>detail extensive</b> |
| Route Distinguisher                          | 64-bit prefix added to IP subnets to make them unique.                                                                                                                                                                                                                                                                                          | <b>detail extensive</b> |
| Label-Base, range                            | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.                                                                                                                                                                          | <b>detail extensive</b> |
| VPN Label                                    | Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.                                                                                                                                           | <b>detail extensive</b> |
| Next hop                                     | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.                                                                                                                                                                                                                                               | All levels              |
| Localpref or Lclpref                         | Local preference value included in the route.                                                                                                                                                                                                                                                                                                   | All levels              |

Table 384: show route receive-protocol Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output         |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| AS path             | <p>Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893.</li> <li>• <b>[ ]</b>—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> | All levels              |
| Cluster list        | (For route reflected output only) Cluster ID sent by the route reflector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> |
| Originator ID       | (For route reflected output only) Address of routing device that originally sent the route to the route reflector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> |
| Communities         | Community path attribute for the route. See the Output Field table in the <a href="#">show route detail</a> command for all possible values for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| AIGP                | Accumulated interior gateway protocol (AIGP) BGP attribute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| Attrset AS          | Number, local preference, and path of the AS that originated the route. These values are stored in the <b>Attrset</b> attribute at the originating routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> |
| Layer2-info: encaps | Layer 2 encapsulation (for example, VPLS).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b> |
| control flags       | Control flags: <b>none</b> or <b>Site Down</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |
| mtu                 | Maximum transmission unit (MTU) of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> |

## Sample Output

### show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
10.22.1.0/24     10.255.245.215    0         100      I
10.22.2.0/24     10.255.245.215    0         100      I
```

### show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
1.1.1.0/24 (1 entry, 1 announced)
  Next hop: 10.0.50.3
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
165.3.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
165.4.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
195.1.2.0/24 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

### show route receive-protocol bgp table extensive

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
  Nexthop: 207.17.136.29
  Localpref: 100
  AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
  AS path: AS4 PA[2]: 33437 393219
  AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
  Communities: 2914:420
```

**show route receive-protocol bgp logical-system extensive**

```
user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

10.0.0.8/30 (2 entries, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 44.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300096
  Nexthop: 10.0.0.9
  AS path: 13979 I
  AIGP: 203

* 55.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300112
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I
  AIGP: 25

* 66.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300144
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I

* 99.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300160
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I
```



### show route receive-protocol bgp detail (Layer 2 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED      Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

### show route receive-protocol bgp extensive (Layer 2 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED      Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100

```

```

AS path: I
Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

### show route receive-protocol bgp (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171
inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lclpref AS path
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.179/32 10.255.14.171          2    100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.177/32 10.255.14.171          100 I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.171:300:10.255.14.177/32
                  10.255.14.171          100 I
10.255.14.171:100:10.255.14.179/32
                  10.255.14.171          2    100 I
10.255.14.171:200:10.255.14.175/32
                  10.255.14.171          100 2 I

```

### show route receive-protocol bgp detail (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

```

* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

#### show route receive-protocol bgp extensive (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
  1.1.1.0/24 (1 entry, 1 announced)
    Nexthop: 10.0.50.3
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  165.3.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
  165.4.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  195.1.2.0/24 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)

```

## show route resolution

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3676</a><br><a href="#">Syntax (EX Series Switches) on page 3676</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                      | <pre>show route resolution &lt;brief   detail   extensive   summary&gt; &lt;index <i>index</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;prefix&gt; &lt;table <i>routing-table-name</i>&gt; &lt;unresolved&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax (EX Series Switches)</b> | <pre>show route resolution &lt;brief   detail   extensive   summary&gt; &lt;index <i>index</i>&gt; &lt;prefix&gt; &lt;table <i>routing-table-name</i>&gt; &lt;unresolved&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                 | Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                     | <p><b>none</b>—Display standard information about all entries in the next-hop resolution database.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>index <i>index</i></b>—(Optional) Show the index of the resolution tree.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>prefix <i>network/destination-prefix</i></b>—(Optional) Display database entries for the specified address.</p> <p><b>table <i>routing-table-name</i></b>—(Optional) Display information about a particular routing table (for example, <b>inet.0</b>) where policy-based export is currently enabled.</p> <p><b>unresolved</b>—(Optional) Display routes that could not be resolved.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Route Resolution on PE Routers</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>       | <a href="#">show route resolution detail on page 3677</a><br><a href="#">show route resolution summary on page 3678</a><br><a href="#">show route resolution unresolved on page 3678</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

**Output Fields** Table 385 on page 3677 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

**Table 385: show route resolution Output Fields**

| Field Name                         | Field Description                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>routing-table-name</i>          | Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.                                                                                     |
| <b>Tree index</b>                  | Tree index identifier.                                                                                                                                                                                                                                                                                                                                |
| <b>Nodes</b>                       | Number of nodes in the tree.                                                                                                                                                                                                                                                                                                                          |
| <b>Reference count</b>             | Number of references made to the next hop.                                                                                                                                                                                                                                                                                                            |
| <b>Contributing routing tables</b> | Routing tables used for next-hop resolution.                                                                                                                                                                                                                                                                                                          |
| <b>Originating RIB</b>             | Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of <b>inet.0</b> resolving through <b>inet.0</b> and <b>inet.3</b> , this field indicates which routing table, <b>inet.0</b> or <b>inet.3</b> , provided the best path for a particular prefix. |
| <b>Metric</b>                      | Metric associated with the forwarding next hop.                                                                                                                                                                                                                                                                                                       |
| <b>Node path count</b>             | Number of nodes in the path.                                                                                                                                                                                                                                                                                                                          |
| <b>Forwarding next hops</b>        | Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.                                                                                                                                                                 |

## Sample Output

### show route resolution detail

```

user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0

```

```
Node path count: 1
Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
Metric: 2 Node path count: 1
Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1
```

#### show route resolution summary

```
user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
Contributing routing tables: :voice.inet.0 :voice.inet.3
Tree Index: 2, Nodes 2, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 3, Nodes 43, Reference Count 1
Contributing routing tables: inet.0 inet.3
```

#### show route resolution unresolved

```
user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
  Protocol Nexthop: 10.255.71.238 Push 800000
  Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
  Protocol Nexthop: 10.255.70.103 Push 800008
  Indirect nexthop: 0 -
Tree Index 2
```

## show route snooping

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show route snooping &lt;brief   detail   extensive   terse&gt; &lt;all&gt; &lt;best address/prefix&gt; &lt;exact address&gt; &lt;range prefix-range&gt; &lt;summary&gt; &lt;table table-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Display the entries in the routing table that were learned from snooping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>none</b>—Display the entries in the routing table that were learned from snooping.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>.</p> <p><b>all</b>—(Optional) Display all entries, including hidden entries.</p> <p><b>best address/prefix</b>—(Optional) Display the longest match for the provided address and optional prefix.</p> <p><b>exact address/prefix</b>—(Optional) Display exact matches for the provided address and optional prefix.</p> <p><b>range prefix-range</b>—(Optional) Display information for the provided address range.</p> <p><b>summary</b>—(Optional) Display route snooping summary statistics.</p> <p><b>table table-name</b>—(Optional) Display information for the named table.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show route snooping detail on page 3679</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### show route snooping detail

```
user@host> show route snooping detail
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
224.0.0.2/32 (1 entry, 1 announced)
  *IGMP    Preference: 0
           Next hop type: MultiRecv
           Next-hop reference count: 4
           State: <Active NoReadvrt Int>
```

```
Age: 2:24
Task: IGMP
Announcement bits (1): 0-KRT
AS path: I

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
    Next hop type: MultiRecv
    Next-hop reference count: 4
    State: <Active NoReadvrt Int>
    Age: 2:24
    Task: IGMP
    Announcement bits (1): 0-KRT
    AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4), Next hop index: 1048584
    Next-hop reference count: 4
    State: <Active Int>
    Age: 2:24
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.2.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.3.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.4.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.5.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
```



```
State: <Active Int>
Age: 1:58
Task: MC
Announcement bits (1): 0-KRT
AS path: I

225.0.0.6.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.7.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.9.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.10.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.1.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.2.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
```

```
Age: 8
Task: MC
Announcement bits (1): 0-KRT
AS path: I

226.0.0.4.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.8.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.10.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:56
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.1.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.2.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.3.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:16
```

```
Task: MC
Announcement bits (1): 0-KRT
AS path: I

227.0.0.4.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.5.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.7.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.8.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.10.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.1.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
```

```
Announcement bits (1): 0-KRT
AS path: I

228.0.0.2.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:18
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.7.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:11
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.8.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.9.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 8
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.10.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.3.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
```

```
AS path: I

229.0.0.4.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.5.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 9
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.6.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.7.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.8.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.9.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

```
229.0.0.10.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

## show route source-gateway

|                                    |                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3687</a><br><a href="#">Syntax (EX Series Switches) on page 3687</a>                                                                                                                                                                                                                                                                                           |
| <b>Syntax</b>                      | <b>show route source-gateway</b> <i>address</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switches)</b> | <b>show route source-gateway</b> <i>address</i><br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                     |
| <b>Description</b>                 | Display the entries in the routing table that were learned from a particular address. The <b>Source</b> field in the <b>show route detail</b> command output lists the source for each route, if known.                                                                                                                                                                                   |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b>address</b> —IP address of the system.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>       | <a href="#">show route source-gateway on page 3687</a><br><a href="#">show route source-gateway detail on page 3688</a><br><a href="#">show route source-gateway extensive on page 3690</a>                                                                                                                                                                                               |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                            |

## Sample Output

### show route source-gateway

```

user@host> show route source-gateway 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

```
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3
```

#### show route source-gateway detail

```
user@host> show route source-gateway 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```



Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 1 announced)

```
*BGP    Preference: 170/-101
        Route Distinguisher: 10.255.70.103:1
        Next-hop reference count: 7
        Source: 10.255.70.103
        Protocol next hop: 10.255.70.103
        Indirect next hop: 2 no-forward
        State: <Secondary Active Int Ext>
        Local AS:    69 Peer AS:    69
        Age: 12:14:00  Metric2: 1
        Task: BGP_69.10.255.70.103+179
        Announcement bits (1): 0-green-12vpn
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0
        Label-base: 800008, range: 8
        Localpref: 100
        Router ID: 10.255.70.103
        Primary Routing Table bgp.12vpn.0
```

red.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)

```
*BGP    Preference: 170/-1
        Route Distinguisher: 10.255.70.103:2
        Next-hop reference count: 7
        Source: 10.255.70.103
        Protocol next hop: 10.255.70.103
        Indirect next hop: 2 no-forward
        State: <Secondary Active Int Ext>
        Local AS:    69 Peer AS:    69
        Age: 12:14:00  Metric2: 1
        Task: BGP_69.10.255.70.103+179
        Announcement bits (1): 0-red-12vpn
        AS path: I
        Communities: target:11111:2 Layer2-info: encaps:VPLS,
        control flags:Site-Down, mtu: 0
        Label-base: 800016, range: 8
        Localpref: 0
        Router ID: 10.255.70.103
        Primary Routing Table bgp.12vpn.0
```

bgp.12vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)

```
*BGP    Preference: 170/-101
        Route Distinguisher: 10.255.70.103:1
        Next-hop reference count: 7
        Source: 10.255.70.103
        Protocol next hop: 10.255.70.103
        Indirect next hop: 2 no-forward
        State: <Active Int Ext>
        Local AS:    69 Peer AS:    69
        Age: 12:14:00  Metric2: 1
        Task: BGP_69.10.255.70.103+179
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS, control
flags:, mtu: 0
```

```

Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Secondary Tables: green.l2vpn.0
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
  *BGP Preference: 170/-1
    Route Distinguisher: 10.255.70.103:2
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:14:00 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    AS path: I
    Communities: target:11111:2 Layer2-info: encaps:VPLS,
    control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8
    Localpref: 0
    Router ID: 10.255.70.103
    Secondary Tables: red.l2vpn.0

```

#### show route source-gateway extensive

```

user@host> show route source-gateway 10.255.70.103 extensive
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:15:24 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0

```

```

Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-1
    Route Distinguisher: 10.255.70.103:2
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:15:24 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-red-l2vpn
    AS path: I
    Communities: target:11111:2 Layer2-info: encaps:VPLS,
    control flags:Site-Down, mtu: 0
    Label-base: 800016, range: 8
    Localpref: 0
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:15:24 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Secondary Tables: green.l2vpn.0
    Indirect next hops: 1
      Protocol next hop: 10.255.70.103 Metric: 2
      Indirect next hop: 2 no-forward
      Indirect path forwarding next hops: 1
    Next hop: via so-0/3/0.0 weight 0x1
      10.255.70.103/32 Originating RIB: inet.3
      Metric: 2 Node path count: 1
      Forwarding nexthops: 1
      Nexthop: via so-0/3/0.0

10.255.70.103:2:3:1/96 (1 entry, 0 announced)
  *BGP Preference: 170/-1

```

```
Route Distinguisher: 10.255.70.103:2
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Secondary Tables: red.12vpn.0
Indirect next hops: 1
    Protocol next hop: 10.255.70.103 Metric: 2
    Indirect next hop: 2 no-forward
    Indirect path forwarding next hops: 1
Next hop: via so-0/3/0.0 weight 0x1
    10.255.70.103/32 Originating RIB: inet.3
    Metric: 2 Node path count: 1
    Forwarding nexthops: 1
    Nexthop: via so-0/3/0.0
```

## show route summary

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3693</a><br><a href="#">Syntax (EX Series Switches) on page 3693</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                      | <pre>show route summary &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;table <i>routing-table-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switches)</b> | show route summary                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                 | <p>Display summary statistics about the entries in the routing table.</p> <p>CPU utilization might increase while the device learns routes. We recommend that you use the <b>show route summary</b> command after the device learns and enters the routes into the routing table. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the <b>show route summary</b> command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the command-line interface (CLI).</p> |
| <b>Options</b>                     | <p><b>none</b>—Display summary statistics about the entries in the routing table.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>table <i>routing-table-name</i></b>—(Optional) Display summary statistics for all routing tables whose name begins with this string (for example, <b>inet.0</b> and <b>inet6.0</b> are both displayed when you run the <b>show route summary table inet</b> command). If you only want to display statistics for a specific routing table, make sure to enter the exact name of that routing table.</p>      |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>       | <a href="#">show route summary on page 3694</a><br><a href="#">show route summary table on page 3695</a><br><a href="#">show route summary table (with Route Limits Configured for the Routing Table) on page 3695</a>                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>               | <p><a href="#">Table 386 on page 3693</a> lists the output fields for the <b>show route summary</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Table 386: show route summary Output Fields**

| Field Name | Field Description                    |
|------------|--------------------------------------|
| Router ID  | Address of the local routing device. |

Table 386: show route summary Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>routing-table-name</i> | Name of the routing table (for example, <b>inet.0</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>destinations</b>       | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>routes</b>             | Number of routes in the routing table: <ul style="list-style-type: none"> <li><b>active</b>—Number of routes that are active.</li> <li><b>holddown</b>—Number of routes that are in the hold-down state before being declared inactive.</li> <li><b>hidden</b>—Number of routes that are not used because of routing policy.</li> </ul>                                                                                                                                                                                                                                                                                                                                                     |
| <b>Limit/Threshold</b>    | Displays the configured route limits for the routing table set with the <b>maximum-prefixes</b> and the <b>maximum-paths</b> statements. If you do not configure route limits for the routing table, the show output does not display this information. <ul style="list-style-type: none"> <li><b>destinations</b>—The first number represents the maximum number of route prefixes installed in the routing table. The second number represents the number of route prefixes that trigger a warning message.</li> <li><b>routes</b>—The first number represents the maximum number of routes. The second number represents the number of routes that trigger a warning message.</li> </ul> |
| <b>Direct</b>             | Routes on the directly connected network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Local</b>              | Local routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <i>protocol-name</i>      | Name of the protocol from which the route was learned. For example, <b>OSPF</b> , <b>RSVP</b> , and <b>Static</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## Sample Output

### show route summary

```

user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
      Direct:    6 routes,      5 active
      Local:    4 routes,      4 active
      OSPF:     5 routes,      4 active
      Static:   7 routes,      7 active
      IGMP:     1 routes,      1 active
      PIM:      2 routes,      2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
      RSVP:      2 routes,      2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

```

Restart Complete
  Direct:      1 routes,      1 active

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
  MPLS:       3 routes,      3 active
  VPLS:       4 routes,      2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
  Direct:      2 routes,      2 active
  PIM:         2 routes,      2 active
  MLD:         1 routes,      1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
  BGP:         2 routes,      2 active
  L2VPN:       2 routes,      2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
  BGP:         2 routes,      2 active
  L2VPN:       1 routes,      1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
  BGP:         4 routes,      4 active

```

#### show route summary table

```

user@host> show route summary table inet
Router ID: 192.168.0.1

inet.0: 32 destinations, 34 routes (31 active, 0 holddown, 1 hidden)
  Direct:      6 routes,      5 active
  Local:       9 routes,      9 active
  OSPF:        3 routes,      1 active
  Static:     13 routes,     13 active
  ICMP:        1 routes,      1 active
  PIM:         2 routes,      2 active

inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Multicast:    1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
  Local:        1 routes,      1 active
  PIM:          2 routes,      2 active

inet6.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Multicast:    1 routes,      1 active

```

#### show route summary table (with Route Limits Configured for the Routing Table)

```

user@host> show route summary table VPN-A.inet.0
Autonomous system number: 100
Router ID: 10.255.182.142

VPN-A.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
Limit/Threshold: 2000/200 destinations 20/12 routes
  Direct:      2 routes,      2 active
  Local:       1 routes,      1 active

```

|       |           |          |
|-------|-----------|----------|
| OSPF: | 4 routes, | 3 active |
| BGP:  | 4 routes, | 4 active |
| IGMP: | 1 routes, | 1 active |
| PIM:  | 2 routes, | 2 active |



## show route table

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3697</a><br><a href="#">Syntax (EX Series Switches) on page 3697</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax</b>                      | show route table <i>routing-table-name</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (EX Series Switches)</b> | show route table <i>routing-table-name</i><br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                 | Display the route entries in a particular routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                     | <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>routing-table-name</i></b>—Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the <b>show route table inet</b> command).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li>• <a href="#">show route summary on page 3693</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>       | <a href="#">show route table bgp.l2.vpn on page 3698</a><br><a href="#">show route table bgp.l3vpn.0 on page 3698</a><br><a href="#">show route table bgp.l3vpn.0 detail on page 3698</a><br><a href="#">show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 3700</a><br><a href="#">show route table bgp.evpn.0 on page 3700</a><br><a href="#">show route table inet.0 on page 3700</a><br><a href="#">show route table inet6.0 on page 3701</a><br><a href="#">show route table inet6.3 on page 3701</a><br><a href="#">show route table inetflow detail on page 3701</a><br><a href="#">show route table l2circuit.0 on page 3702</a><br><a href="#">show route table mpls on page 3702</a><br><a href="#">show route table mpls extensive on page 3702</a><br><a href="#">show route table mpls.0 on page 3703</a><br><a href="#">show route table mpls.0 detail (PTX Series) on page 3704</a><br><a href="#">show route table mpls.0 extensive (PTX Series) on page 3704</a><br><a href="#">show route table mpls.0 (RSVP Route—Transit LSP) on page 3705</a><br><a href="#">show route table vpls_1 detail on page 3705</a> |

[show route table vpn-a on page 3705](#)  
[show route table vpn-a.mdt.0 on page 3706](#)  
[show route table VPN-A detail on page 3706](#)  
[show route table VPN-AB.inet.0 on page 3707](#)  
[show route table VPN\\_blue.mvpn-inet6.0 on page 3707](#)  
[show route table vrf1.mvpn.0 extensive on page 3707](#)  
[show route table inetflow detail on page 3708](#)

**Output Fields** For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

## Sample Output

### [show route table bgp.l2vpn](#)

```
user@host> show route table bgp.l2vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
    *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

### [show route table bgp.l3vpn.0](#)

```
user@host> show route table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.15:100:10.255.71.17/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.15:200:10.255.71.18/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100021, Push 100011(top)
```

### [show route table bgp.l3vpn.0 detail](#)

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
    Label operation: Push 182449
    Protocol next hop: 10.255.245.12
    Push 182449
    Indirect next hop: 863a630 297
    State: <Active Int Ext>
    Local AS: 35 Peer AS: 35
    Age: 12:19 Metric2: 1
```

```

Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

Communities: 2914:420 target:11111:1 origin:56:78
VPN Label: 182449
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 863a8f0 305
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>

```

```

Local AS:    35 Peer AS:    35
Age: 12:19    Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100

```

### show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```

user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
                *[RTarget/5] 00:03:14
                  Type Proxy
                    for 10.255.165.103
                    for 10.255.166.124
                  Local

```

### show route table bgp.evpn.0

```

user@host> show route table bgp.evpn.0
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:100.100.100.2:100::0:00:26:88:5f:67:b0/304
                *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                  AS path: I, validation-state: unverified
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.2:100::0:00:51:51:51:51:51/304
                *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                  AS path: I, validation-state: unverified
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.3:100::0:00:52:52:52:52:52/304
                *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                  AS path: I, validation-state: unverified
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
2:100.100.100.3:100::0:a8:d0:e5:5b:01:c8/304
                *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                  AS path: I, validation-state: unverified
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
3:100.100.100.2:100::1000::100.100.100.2/304
                *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                  AS path: I, validation-state: unverified
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
3:100.100.100.2:100::2000::100.100.100.2/304
                *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                  AS path: I, validation-state: unverified
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1

```

### show route table inet.0

```

user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0      *[Static/5] 00:51:57

```

```

> to 111.222.5.254 via fxp0.0
1.0.0.1/32      *[Direct/0] 00:51:58
> via at-5/3/0.0
1.0.0.2/32      *[Local/0] 00:51:58
                Local
12.12.12.21/32  *[Local/0] 00:51:57
                Reject
13.13.13.13/32  *[Direct/0] 00:51:58
> via t3-5/2/1.0
13.13.13.14/32  *[Local/0] 00:51:58
                Local
13.13.13.21/32  *[Local/0] 00:51:58
                Local
13.13.13.22/32  *[Direct/0] 00:33:59
> via t3-5/2/0.0
127.0.0.1/32    [Direct/0] 00:51:58
> via lo0.0
111.222.5.0/24  *[Direct/0] 00:51:58
> via fxp0.0
111.222.5.81/32 *[Local/0] 00:51:58
                Local

```

#### show route table inet6.0

```

user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0

```

#### show route table inet6.3

```

user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                *[LDP/9] 00:00:22, metric 1
                > via so-1/0/0.0
::10.255.245.196/128
                *[LDP/9] 00:00:08, metric 1
                > via so-1/0/0.0, Push 100008

```

#### show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
              Next-hop reference count: 2
              State: <Active Ext>
              Local AS: 65002 Peer AS: 65000
              Age: 4
              Task: BGP_65000.10.12.99.5+3792
              Announcement bits (1): 0-Flow

```

```

AS path: 65000 I
Communities: traffic-rate:0:0
Validation state: Accept, Originator: 10.12.99.5
Via: 10.12.44.0/24, Active
Localpref: 100
Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
*Flow Preference: 5
Next-hop reference count: 2
State: <Active>
Local AS: 65002
Age: 6:30
Task: RT Flow
Announcement bits (2): 0-Flow 1-BGP.0.0.0+179
AS path: I
Communities: 1:1

```

### show route table l2circuit.0

```

user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
    *[LDP/9] 00:50:14
    Discard
10.1.1.195:CtrlWord:1:2:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
    *[LDP/9] 00:50:14
    Discard

```

### show route table mpls

```

user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 00:13:55, metric 1
            Receive
1          *[MPLS/0] 00:13:55, metric 1
            Receive
2          *[MPLS/0] 00:13:55, metric 1
            Receive
1024       *[VPN/0] 00:04:18
            to table red.inet.0, Pop

```

### show route table mpls extensive

```

user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
    *LDP Preference: 9
    Next hop: via so-1/0/0.0, selected

```

```

Pop
State: <Active Int>
Age: 29:50      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
Prefixes bound to route: 10.0.0.194/32

```

### show route table mpls.0

```

user@host> show route table mpls.0
mpls.0: 18 destinations, 19 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 11:39:56, metric 1
                  to table inet.0
0(S=0)           *[MPLS/0] 11:39:56, metric 1
                  to table mpls.0
1                *[MPLS/0] 11:39:56, metric 1
                  Receive
2                *[MPLS/0] 11:39:56, metric 1
                  to table inet6.0
2(S=0)           *[MPLS/0] 11:39:56, metric 1
                  to table mpls.0
13              *[MPLS/0] 11:39:56, metric 1
                  Receive
303168           *[EVPN/7] 11:00:49, routing-instance pbbn10, route-type
Ingress-MAC, ISID 0
                  to table pbbn10.evpn-mac.0
303184           *[EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 1000
                  to table pbbn10.evpn-mac.0
                  [EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 2000
                  to table pbbn10.evpn-mac.0
303264           *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-IM, ISID 1000
                  > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303280           *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-IM, ISID 2000
                  > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303328           *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
                  > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303344           *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
                  > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303360           *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:26:88:5f:67:b0
                  > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303376           *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:51:51:51:51:51
                  > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303392           *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
                  > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303408           *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
                  > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303424           *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC a8:d0:e5:5b:01:c8

```

```

> to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303440      *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:52:52:52:52:52
> to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2

```

### show route table mpls.0 detail (PTX Series)

```

user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
  *L2VPN Preference: 7
    Next hop type: Indirect
    Address: 0x9438f34
    Next-hop reference count: 2
    Next hop type: Router, Next hop index: 567
    Next hop: 3.0.0.1 via ge-0/0/1.0, selected
    Label operation: Push 299808
    Label TTL action: prop-ttl
    Load balance label: Label 299808:None;
    Session Id: 0x1
    Protocol next hop: 10.255.255.1
    Label operation: Push 299872 Offset: 252
    Label TTL action: no-prop-ttl
    Load balance label: Label 299872:Flow label PUSH;
    Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
    Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
    State: <Active Int>
    Age: 21 Metric2: 1
    Validation State: unverified
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 2-Common L2 VC
    AS path: I

```

### show route table mpls.0 extensive (PTX Series)

```

user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0 /32 -> {composite(570)}
  *L2VPN Preference: 7
    Next hop type: Indirect
    Address: 0x9438f34
    Next-hop reference count: 2
    Next hop type: Router, Next hop index: 567
    Next hop: 3.0.0.1 via ge-0/0/1.0, selected
    Label operation: Push 299808
    Label TTL action: prop-ttl
    Load balance label: Label 299808:None;
    Session Id: 0x1
    Protocol next hop: 10.255.255.1
    Label operation: Push 299872 Offset: 252
    Label TTL action: no-prop-ttl
    Load balance label: Label 299872:Flow label PUSH;
    Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
    Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
    State: <Active Int>
    Age: 47 Metric2: 1
    Validation State: unverified
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 2-Common L2 VC
    AS path: I
    Composite next hops: 1

```



```

Protocol next hop: 10.255.255.1 Metric: 1
Label operation: Push 299872 Offset: 252
Label TTL action: no-prop-ttl
Load balance label: Label 299872:Flow label PUSH;
Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 3.0.0.1 via ge-0/0/1.0
    Session Id: 0x1
10.255.255.1/32 Originating RIB: inet.3
    Metric: 1                      Node path count: 1
    Forwarding nexthops: 1
    Nexthop: 3.0.0.1 via ge-0/0/1.0

```

### show route table mpls.0 (RSVP Route—Transit LSP)

```

user@host> show route table mpls.0

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 00:37:31, metric 1
                 Receive
1                *[MPLS/0] 00:37:31, metric 1
                 Receive
2                *[MPLS/0] 00:37:31, metric 1
                 Receive
13               *[MPLS/0] 00:37:31, metric 1
                 Receive
300352           *[RSVP/7/1] 00:08:00, metric 1
                 > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0)      *[RSVP/7/1] 00:08:00, metric 1
                 > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384           *[RSVP/7/2] 00:05:20, metric 1
                 > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0)      *[RSVP/7/2] 00:05:20, metric 1
                 > to 8.64.1.106 via ge-1/0/0.0, Pop

```

### show route table vpls\_1 detail

```

user@host> show route table vpls_1 detail
vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF

```

### show route table vpn-a

```

user@host> show route table vpn-a
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

```

```

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
    *[VPN/7] 05:48:27
    Discard
192.168.24.1:1:2:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

### show route table vpn-a.mdt.0

```

user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
    *[MVPN/70] 01:23:05, metric2 1
    Indirect
1:1:1:10.255.14.218:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
    AS path: I
    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
    AS path: I
    > via so-0/0/1.0, label-switched-path r0-to-r2

```

### show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
    *BGP    Preference: 170/-101
    Route Distinguisher: 10.255.179.13:200
    Next hop type: Indirect
    Next-hop reference count: 5
    Source: 10.255.179.13
    Next hop type: Router, Next hop index: 732
    Next hop: 10.39.1.14 via fe-0/3/0.0, selected
    Label operation: Push 299824, Push 299824(top)
    Protocol next hop: 10.255.179.13
    Push 299824
    Indirect next hop: 8f275a0 1048574
    State: (Secondary Active Int Ext)
    Local AS: 1 Peer AS: 1
    Age: 3:41:06 Metric: 1 Metric2: 1
    Task: BGP_1.10.255.179.13+64309
    Announcement bits (2): 0-KRT 1-BGP RT Background
    AS path: I
    Communities: target:1:200 rte-type:0.0.0.0:1:0
    Import Accepted
    VPN Label: 299824 TTL Action: vrf-ttl-propagate
    Localpref: 100
    Router ID: 10.255.179.13
    Primary Routing Table bgp.13vpn.0

```

### show route table VPN-AB.inet.0

```

user@host> show route table VPN-AB.inet.0
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.0/30      *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0
10.39.1.4/30      *[Direct/0] 00:08:42
                  > via so-5/1/0.0
10.39.1.6/32      *[Local/0] 00:08:46
                  Local
10.255.71.16/32   *[Static/5] 00:07:24
                  > via so-2/0/0.0
10.255.71.17/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32 *[BGP/170] 00:08:35, localpref 100
                  AS path: 2 I
                  > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32 *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0

```

### show route table VPN\_blue.mvpn-inet6.0

```

user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                  AS path: I
                  > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                  AS path: I
                  > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
                  *[MVPN/70] 00:57:23, metric2 1
                  Indirect
5:10.255.2.202:65535:128::192.168.90.2:128:ffff::1/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                  AS path: I
                  > via so-0/1/3.0
6:10.255.2.203:65535:65000:128::10.12.53.12:128:ffff::1/432
                  *[PIM/105] 00:02:37
                  Multicast (IPv6)
7:10.255.2.202:65535:65000:128::192.168.90.2:128:ffff::1/432
                  *[MVPN/70] 00:02:37, metric2 1
                  Indirect

```

### show route table vrf1.mvpn.0 extensive

```

user@host> show route table vrf1.mvpn.0 extensive
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
    *MVPN    Preference: 70

```

```

PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
  Next hop type: Indirect
  Address: 0xbb2c944
  Next-hop reference count: 360
  Protocol next hop: 10.255.50.77
  Indirect next hop: 0x0 - INH Session ID: 0x0
  State: <Active Int Ext>
  Age: 53:03      Metric2: 1
  Validation State: unverified
  Task: mvpn global task
  Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export

AS path: I

```

### show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
    Next-hop reference count: 2
    State: <Active Ext>
    Local AS: 65002 Peer AS: 65000
    Age: 4
    Task: BGP_65000.10.12.99.5+3792
    Announcement bits (1): 0-Flow
    AS path: 65000 I
    Communities: traffic-rate:0:0
    Validation state: Accept, Originator: 10.12.99.5
    Via: 10.12.44.0/24, Active
    Localpref: 100
    Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
  *Flow      Preference: 5
    Next-hop reference count: 2
    State: <Active>
    Local AS: 65002
    Age: 6:30
    Task: RT Flow
    Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
    AS path: I
    Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1.1.1.4/96 AD
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/1.5
1.1.1.2:100:1:0/96 MH
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1:0/96 MH
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4

```

```

        AS path: I, validation-state: unverified
        > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
        *[VPLS/7] 1d 03:11:02, metric2 1
        > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
        *[LDP/9] 1d 03:11:02
        Discard

user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
hidden)
22.0.0.0/32 (3 entries, 1 announced)
        State: <OnList CalcForwarding>
TSI:
KRT in-kernel 22.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

        Nexthop: Self
        AS path: [2] I
        Communities: target:2:1
Path 22.0.0.0 from 2.3.0.0 Vector len 4. Val: 1
    @BGP      Preference: 170/-1
                Route Distinguisher: 2:1
                Next hop type: Indirect
                Address: 0x258059e4
                Next-hop reference count: 2
                Source: 2.2.0.0
                Next hop type: Router
                Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                Label operation: Push 707633
                Label TTL action: prop-ttl
                Session Id: 0x17d8
                Protocol next hop: 2.2.0.0
                Push 16
                Composite next hop: 0x25805988 - INH Session ID: 0x193c
                Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                State: <Secondary Active Int Ext ProtectionPath ProtectionCand>
                Local AS:      2 Peer AS:      2
                Age: 23      Metric2: 35
                Validation State: unverified
                Task: BGP_2.2.2.0.0+34549
                AS path: I
                Communities: target:2:1
                Import Accepted
                VPN Label: 16
                Localpref: 0
                Router ID: 2.2.0.0
                Primary Routing Table bgp.13vpn.0
                Composite next hops: 1
                    Protocol next hop: 2.2.0.0 Metric: 35
                    Push 16
                    Composite next hop: 0x25805988 - INH Session ID: 0x193c
                    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                    Indirect path forwarding next hops: 1
                        Next hop type: Router
                        Next hop: 10.1.1.1 via ge-1/1/9.0
                        Session Id: 0x17d8
                2.2.0.0/32 Originating RIB: inet.3
                    Metric: 35      Node path count: 1
                    Forwarding nexthops: 1
                        Nexthop: 10.1.1.1 via ge-1/1/9.0

```

```

BGP      Preference: 170/-1
         Route Distinguisher: 2:1
         Next hop type: Indirect
         Address: 0x9347028
         Next-hop reference count: 3
         Source: 2.3.0.0
         Next hop type: Router, Next hop index: 702
         Next hop: 10.1.4.2 via ge-1/0/0.0, selected
         Label operation: Push 634278
         Label TTL action: prop-ttl
         Session Id: 0x17d9
         Protocol next hop: 2.3.0.0
         Push 16
         Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
         Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
         State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

         Inactive reason: Not Best in its group - IGP metric
         Local AS:      2 Peer AS:      2
         Age: 3:34      Metric2: 70
         Validation State: unverified
         Task: BGP_2.2.3.0.0+32805
         Announcement bits (2): 0-KRT 1-BGP_RT_Background
         AS path: I
         Communities: target:2:1
         Import Accepted
         VPN Label: 16
         Localpref: 0
         Router ID: 2.3.0.0
         Primary Routing Table bgp.l3vpn.0
         Composite next hops: 1
             Protocol next hop: 2.3.0.0 Metric: 70
             Push 16
             Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
             Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da
             Indirect path forwarding next hops: 1
                 Next hop type: Router
                 Next hop: 10.1.4.2 via ge-1/0/0.0
                 Session Id: 0x17d9
                 2.3.0.0/32 Originating RIB: inet.3
                 Metric: 70                      Node path count: 1
                 Forwarding nexthops: 1
                 Nexthop: 10.1.4.2 via ge-1/0/0.0
#Multipath Preference: 255
         Next hop type: Indirect
         Address: 0x24afca30
         Next-hop reference count: 1
         Next hop type: Router
         Next hop: 10.1.1.1 via ge-1/1/9.0, selected
         Label operation: Push 707633
         Label TTL action: prop-ttl
         Session Id: 0x17d8
         Next hop type: Router, Next hop index: 702
         Next hop: 10.1.4.2 via ge-1/0/0.0
         Label operation: Push 634278
         Label TTL action: prop-ttl
         Session Id: 0x17d9
         Protocol next hop: 2.2.0.0
         Push 16

```


Composite next hop: 0x25805988 - INH Session ID: 0x193c  
Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

Protocol next hop: 2.3.0.0  
Push 16  
Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da  
Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight

0x4000

State: <ForwardingOnly Int Ext>  
Inactive reason: Forwarding use only  
Age: 23 Metric2: 35  
Validation State: unverified  
Task: RT  
AS path: I  
Communities: target:2:1

## show route terse

|                                                                                                                                                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                   | <a href="#">Syntax on page 3712</a><br><a href="#">Syntax (EX Series Switches) on page 3712</a>                                                                                                                                                 |
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                           | <pre>show route terse &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                           |
| <b>Syntax (EX Series Switches)</b>                                                                                                                                                                                                                                                                                                                                                                                                      | show route terse                                                                                                                                                                                                                                |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                              | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                           |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                      | Display a high-level summary of the routes in the routing table.                                                                                                                                                                                |
| <div>  <p><b>NOTE:</b> For BGP routes, the <code>show route terse</code> command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.</p> <p>To display the metric1 and metric2 value of a BGP route, use the <a href="#">show route extensive</a> command.</p> </div> |                                                                                                                                                                                                                                                 |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                          | <p><b>none</b>—Display a high-level summary of the routes in the routing table.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                         | view                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>                                                                                                                                                                                                                                                                                                                                                                                                            | <a href="#">show route terse on page 3714</a>                                                                                                                                                                                                   |
| <b>Output Fields</b>                                                                                                                                                                                                                                                                                                                                                                                                                    | <a href="#">Table 387 on page 3712</a> describes the output fields for the <code>show route terse</code> command. Output fields are listed in the approximate order in which they appear.                                                       |

**Table 387: show route terse Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).                                                                                                                                                                                                                                                                                                                |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                                         |
| <i>number routes</i>       | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li><b>active</b> (routes that are active)</li> <li><b>holddown</b> (routes that are in the pending state before being declared inactive)</li> <li><b>hidden</b> (routes that are not used because of a routing policy)</li> </ul> |



Table 387: show route terse Output Fields (*continued*)

| Field Name       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>route key</i> | Key for the state of the route: <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul>                                                                                                                                                                                                                |
| <b>A</b>         | Active route. An asterisk (*) indicates this is the active route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>V</b>         | Validation status of the route: <ul style="list-style-type: none"> <li>• <b>?</b>—Not evaluated. Indicates that the route was not learned through BGP.</li> <li>• <b>I</b>—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>N</b>—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>V</b>—Valid. Indicates that the prefix and autonomous system pair are found in the database.</li> </ul> |
| Destination      | Destination of the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>P</b>         | Protocol through which the route was learned: <ul style="list-style-type: none"> <li>• <b>A</b>—Aggregate</li> <li>• <b>B</b>—BGP</li> <li>• <b>C</b>—CCC</li> <li>• <b>D</b>—Direct</li> <li>• <b>G</b>—GMPLS</li> <li>• <b>I</b>—IS-IS</li> <li>• <b>L</b>—L2CKT, L2VPN, LDP, Local</li> <li>• <b>K</b>—Kernel</li> <li>• <b>M</b>—MPLS, MSDP</li> <li>• <b>O</b>—OSPF</li> <li>• <b>P</b>—PIM</li> <li>• <b>R</b>—RIP, RIPng</li> <li>• <b>S</b>—Static</li> <li>• <b>T</b>—Tunnel</li> </ul>                                                                                                                                                                                     |
| <b>Prf</b>       | Preference value of the route. In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.                                                                                              |
| Metric 1         | First metric value in the route. For routes learned from BGP, this is the MED metric.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Metric 2         | Second metric value in the route. For routes learned from BGP, this is the IGP metric.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 387: show route terse Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                          |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Next hop   | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.                                                                                                                                                                                                                                                          |
| AS path    | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>I—IGP.</li> <li>E—EGP.</li> <li>?—Incomplete; typically, the AS path was aggregated.</li> </ul> |

## Sample Output

### show route terse

```

user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* ? 1.0.1.1/32       0 10      1           >10.0.0.2      I
?                               B 170      100           >10.0.0.2      I
  unverified
* ? 1.1.1.1/32       D 0           >10.0.0.2
* V 2.2.0.2/32       B 170     110      >10.0.0.2      200 I
  valid
* ? 10.0.0.0/30      D 0           >1t-1/2/0.1
?                               B 170     100           >10.0.0.2      I
  unverified
* ? 10.0.0.1/32      L 0           Local
* ? 10.0.0.4/30      B 170     100           >10.0.0.2      I
  unverified
* ? 10.0.0.8/30      B 170     100           >10.0.0.2      I
  unverified
* I 172.16.1.1/32     B 170      90      >10.0.0.2      200 I
  invalid
* N 192.168.2.3/32   B 170     100      >10.0.0.2      200 I
  unknown
* ? 224.0.0.5/32     O 10      1           MultiRecv

```

## PART 20

# Multicast

- [Overview on page 3717](#)
- [Configuration on page 3727](#)
- [Administration on page 3823](#)



## CHAPTER 63

# Overview

- [IGMP Snooping Overview on page 3717](#)

## IGMP Snooping Overview

---

- [IGMP Snooping on EX Series Switches Overview on page 3717](#)

## IGMP Snooping on EX Series Switches Overview

Internet Group Management Protocol (IGMP) snooping constrains the flooding of IPv4 multicast traffic on VLANs on a switch. When IGMP snooping is enabled on a VLAN, a Juniper Networks EX Series Ethernet Switch examines IGMP messages between hosts and multicast routers and learns which hosts are interested in receiving traffic for a multicast group. Based on what it learns, the switch then forwards multicast traffic only to those interfaces in the VLAN that are connected to interested receivers instead of flooding the traffic to all interfaces.

IGMP snooping on EX Series switches supports IGMP version 1 (IGMPv1), IGMPv2, and IGMPv3. For details on IGMP, see the following standards:

- IGMPv1—See RFC 1112, *Host extensions for IP multicasting*.
- IGMPv2—See RFC 2236, *Internet Group Management Protocol, Version 2*.
- For IGMPv3—See RFC 3376, *Internet Group Management Protocol, Version 3*.

This topic covers:

- [How IGMP Snooping Works on page 3718](#)
- [IGMP Message Types on page 3718](#)
- [How Hosts Join and Leave Multicast Groups on page 3719](#)
- [Support for IGMPv3 Multicast Sources on page 3719](#)
- [IGMP Snooping and Forwarding Interfaces on page 3720](#)
- [General Forwarding Rules on page 3720](#)
- [Examples of IGMP Snooping Multicast Forwarding on page 3721](#)


## How IGMP Snooping Works

A Layer 2 switch usually learns *unicast* media access control (MAC) addresses by checking the source address field of the frames it receives. However, a *multicast* MAC address can never be the source address for a packet. As a result, the switch floods multicast traffic on the VLAN, consuming significant amounts of bandwidth.

You can enable IGMP snooping on a switch to avoid this flooding. When IGMP snooping is enabled, the switch monitors IGMP messages between receivers and multicast routers and uses the content of the messages to build an IPv4 multicast forwarding table—a database of multicast groups and the interfaces that are connected to members of the groups. When the switch receives multicast traffic for a multicast group, it uses the forwarding table to forward the traffic only to interfaces that are connected to receivers that belong to the multicast group.

Figure 45 on page 3718 shows an example of multicast traffic flow with IGMP snooping enabled.

**Figure 45: Multicast Traffic Flow with IGMP Snooping Enabled**



## IGMP Message Types

Multicast routers use IGMP to learn, for each of their attached physical networks, which groups have interested listeners. In any given subnet, one multicast router acts as an IGMP querier. The IGMP querier sends out the following types of queries to hosts:

- General query—Asks whether any host is listening to any group.
- Group-specific query—(IGMPv2 and IGMPv3 only) Asks whether any host is listening to a specific multicast group. This query is sent in response to a host leaving the multicast group and allows the router to quickly determine if any remaining hosts are interested in the group.

- Group-and-source-specific query—(IGMPv3 only) Asks whether any host is listening to group multicast traffic from a specific multicast source. This query is sent in response to a host indicating that it is not longer interested in receiving group multicast traffic from the multicast source and allows the router to quickly determine any remaining hosts are interested in receiving group multicast traffic from that source.

Hosts that are multicast listeners send the following kinds of messages:

- Membership report—Indicates that the host wants to join a particular multicast group.
- Leave report—(IGMPv2 and IGMPv3 only) Indicates that the host wants to leave a particular multicast group.

### How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in either of two ways:

- By sending an unsolicited IGMP join message to a multicast router that specifies the IP multicast group that the host is attempting to join.
- By sending an IGMP join message in response to a general query from a multicast router.

A multicast router continues to forward multicast traffic to a VLAN provided that at least one host on that VLAN responds to the periodic general IGMP queries. For a host to remain a member of a multicast group, therefore, it must continue to respond to the periodic general IGMP queries.

Hosts can leave a multicast group in either of two ways:

- By not responding to periodic queries within a set interval of time. This results in what is known as a “silent leave.” This is the only method available to IGMPv1 hosts.
- By sending a leave report. This method can be used by IGMPv2 and IGMPv3 hosts.



**NOTE:** If a host is connected to the switch through a hub, the host does not automatically leave the multicast group if it disconnects from the hub. The host remains a member of the group until group membership times out and a silent leave occurs. If another host connects to the hub port before the silent leave occurs, the new host might receive the group multicast traffic until the silent leave, even though it never sent an membership report.

### Support for IGMPv3 Multicast Sources

In IGMPv3, a host can send a membership report that includes a list of source addresses. When the host sends a membership report in INCLUDE mode, the host is interested in group multicast traffic only from those sources in the source address list. If host sends a membership report in EXCLUDE mode, the host is interested in group multicast traffic from any source *except* the sources in the source address list. A host can also send an EXCLUDE report in which the source-list parameter is empty, which is known as an

EXCLUDE NULL report. An EXCLUDE NULL report indicates that the host wants to join the multicast group and receive packets from all sources.

EX Series switches support IGMPv3 membership reports that are in INCLUDE and EXCLUDE mode. However, EX Series switches do not support forwarding on a per-source basis. Instead, a switch consolidates all INCLUDE and EXCLUDE mode reports it receives on a VLAN for a specified group into a single route that includes all multicast sources for that group, with the next hop being all interfaces that have interested receivers for the group. As a result, interested receivers on the VLAN can receive traffic from a source that they did not include in their INCLUDE report or from a source they excluded in their EXCLUDE report. For example, if Host 1 wants traffic for G from Source A and Host 2 wants traffic for G from Source B, they both receive traffic for G regardless of whether A or B sends the traffic.

### IGMP Snooping and Forwarding Interfaces

---

To determine how to forward multicast traffic, a switch with IGMP snooping enabled maintains information about the following interfaces in its multicast forwarding table:

- Multicast-router interfaces—These interfaces lead toward multicast routers or IGMP queriers.
- Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups.

The switch learns about these interfaces by monitoring IGMP traffic. If an interface receives IGMP queries or Protocol Independent Multicast (PIM) updates, the switch adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the switch adds the interface to its multicast forwarding table as a group-member interface.

Table entries for interfaces that the switch learns about are subject to aging. For example, if a learned multicast-router interface does not receive IGMP queries or PIM hellos within a certain interval, the switch removes the entry for that interface from its multicast forwarding table.



**NOTE:** For a switch to learn multicast-router interfaces and group-member interfaces, an IGMP querier must exist in the network. For the switch itself to function as an IGMP querier, IGMP must be enabled on the switch.

---

You can statically configure an interface to be a multicast-router interface or a group-member interface. The switch adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject to aging. You can have a mix of statically configured and dynamically learned interfaces on a switch.

### General Forwarding Rules

---

Multicast traffic received on a switch interface in a VLAN on which IGMP snooping is enabled is forwarded according to the following rules.



IGMP traffic is forwarded as follows:

- IGMP general queries received on a multicast-router interface are forwarded to all other interfaces in the VLAN.
- IGMP group-specific queries received on a multicast-router interface are forwarded to only those interfaces in the VLAN that are members of the group.
- IGMP reports received on a host interface are forwarded to multicast-router interfaces in the same VLAN, but not to the other host interfaces in the VLAN.

Multicast traffic that is not IGMP traffic is forwarded as follows:

- A multicast packet with a destination address of 224.0.0.0/24 is flooded to all other interfaces on the VLAN.
- An unregistered multicast packet—that is, a packet for a group that has no current members—is forwarded to all multicast-router interfaces in the VLAN.
- A registered multicast packet is forwarded only to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN.

### Examples of IGMP Snooping Multicast Forwarding

The following examples are provided to illustrate how IGMP snooping forwards multicast traffic in different topologies:

- [Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts on page 3721](#)
- [Scenario 2: Switch Forwarding Multicast Traffic to Another Switch on page 3722](#)
- [Scenario 3: Switch Connected to Hosts Only \(No IGMP Querier\) on page 3723](#)
- [Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs on page 3724](#)

#### ***Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts***

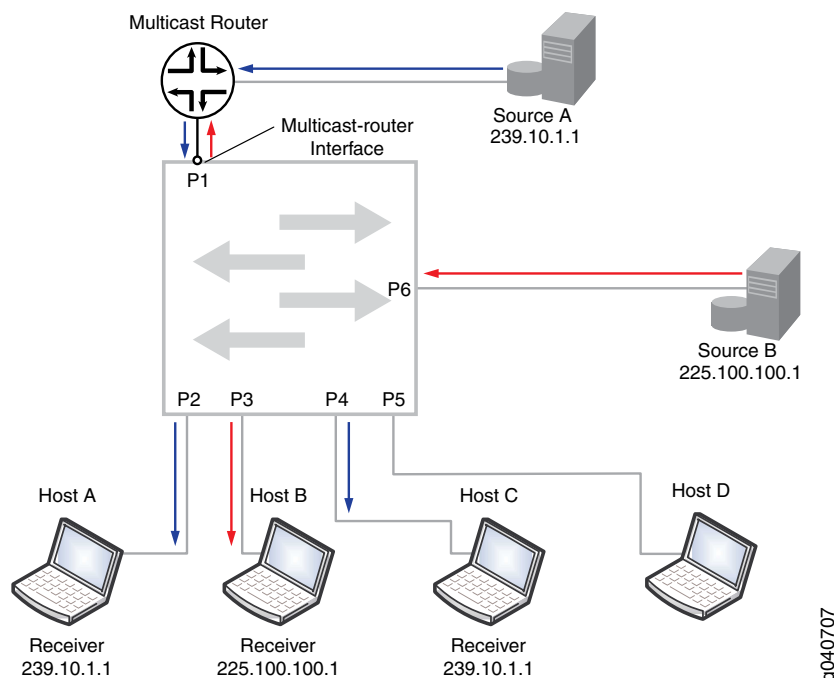
In the topology shown in [Figure 46 on page 3722](#), a switch acting as a Layer 2 device receives multicast traffic belonging to multicast group **239.10.1.1** from Source A, which is connected to the multicast router. It also receives multicast traffic belonging to multicast group **225.100.100.1** from Source B, which is connected directly to the switch. All interfaces on the switch belong to the same VLAN.

Because the switch receives IGMP queries from the multicast router on interface P1, IGMP snooping learns that interface P1 is a multicast-router interface and adds the interface to its multicast cache table. It forwards any IGMP general queries it receives on this interface to all host interfaces on the switch, and, in turn, forwards membership reports it receives from hosts to the multicast-router interface.

In the example, Hosts A and C have responded to the membership queries with membership reports for group **239.10.1.1**. IGMP snooping adds interfaces P2 and P4 to its multicast cache table as member interfaces for group **239.10.1.1**. It forwards the group multicast traffic received from Source A to Hosts A and C, but not to Hosts B and D.

Host B has responded to the membership queries with a membership report for group **225.100.100.1**. The switch adds interface P3 to its multicast cache table as a member interface for group **225.100.100.1** and forwards multicast traffic it receives from Source B to Host B. The switch also forwards the multicast traffic it receives from Source B to the multicast-router interface P1.

**Figure 46: Scenario 1: Switch Forwarding Multicast Traffic to a Multicast Router and Hosts**

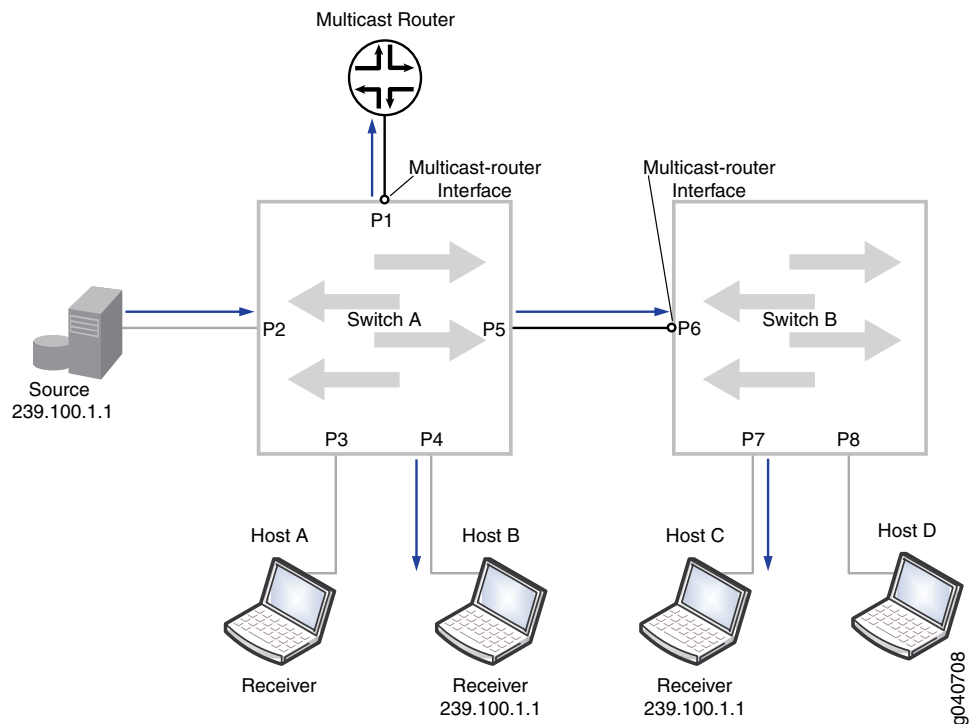


#### **Scenario 2: Switch Forwarding Multicast Traffic to Another Switch**

In the topology shown in [Figure 47 on page 3723](#), a multicast source is connected to Switch A. Switch A in turn is connected to another switch, Switch B. Hosts on both Switch A and B are potential members of the multicast group. Both switches are acting as Layer 2 devices and all interfaces on the switches are members of the same VLAN.

Switch A receives IGMP queries from the multicast router on interface P1, making interface P1 a multicast-router interface for Switch A. Switch A forwards all general IGMP queries it receives on this interface to the other interfaces on the switch, including the interface connecting Switch B. Because Switch B receives the forwarded IGMP queries on interface P6, P6 is the multicast-router interface for Switch B. Switch B forwards the group membership report it receives from Host C to Switch A through its multicast-router interface. Switch A forwards the membership report to its multicast-router interface, includes interface P5 in its multicast cache table as a group-member interface, and forwards multicast traffic from the source to Switch B.

**Figure 47: Scenario 2: Switch Forwarding Multicast Traffic to Another Switch**



In certain implementations, you might have to configure P6 on Switch B as a static multicast-router interface to avoid a delay in a host receiving multicast traffic. For example, if Switch B receives unsolicited membership reports from its hosts before it learns which interface is its multicast-router interface, it does not forward those reports to Switch A. If Switch A then receives multicast traffic, it does not forward the traffic to Switch B, because it has not received any membership reports on interface P5. This issue will resolve when the multicast router sends out its next general query; however, it can cause a delay in the host receiving multicast traffic. You can statically configure interface P6 as a multicast-router interface to solve this issue.


### **Scenario 3: Switch Connected to Hosts Only (No IGMP Querier)**

In the topology shown in [Figure 48 on page 3724](#), a switch is connected to a multicast source and to hosts. There is no multicast router in this topology—hence there is no IGMP querier. Without an IGMP querier to respond to, a host does not send periodic membership reports. As a result, even if the host sends an unsolicited join to join a multicast group, its membership in the multicast group times out.

For IGMP snooping to work correctly in this network so that the switch forwards multicast traffic to Hosts A and C only, you can either:

- Configure interfaces P2 and P4 as static group-member interfaces.
- Configure a routed VLAN interface (RVI) on the VLAN and enable IGMP on it. In this case, the switch itself acts as an IGMP querier, and the hosts can dynamically join the multicast group and refresh their group membership by responding to the queries.


Figure 48: Scenario 3: Switch Connected to Hosts Only (No IGMP Querier)

**Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs**

In the topology shown in [Figure 49 on page 3725](#), a multicast source, Multicast Router A, and Hosts A and B are connected to the switch and are in VLAN 10. Multicast Router B and Hosts C and D are also connected to the switch and are in VLAN 20.

In a pure Layer 2 environment, traffic is not forwarded between VLANs. For Host C to receive the multicast traffic from the source on VLAN 10, RVIs must be created on VLAN 10 and VLAN 20 to permit routing of the multicast traffic between the VLANs. In addition, PIM must be enabled on the switch to perform the multicast routing.

Figure 49: Scenario 4: Layer 2/Layer 3 Switch Forwarding Multicast Traffic Between VLANs



**Related Documentation**

- *Understanding Multicast VLAN Registration*
- *Example: Configuring IGMP Snooping on EX Series Switches*
- *Configuring IGMP Snooping (CLI Procedure)*
- *Configuring Routed VLAN Interfaces (CLI Procedure)*



# Configuration

- [Configuration Examples on page 3727](#)
- [Configuration Tasks on page 3730](#)
- [Configuration Statements on page 3739](#)

## Configuration Examples

---

- [Example: Configuring IGMP Snooping on page 3727](#)

### Example: Configuring IGMP Snooping



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see, *Example: Configuring IGMP Snooping on EX Series Switches*. For ELS details, see, “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

You can enable IGMP snooping on a VLAN to constrain the flooding of IPv4 multicast traffic on a VLAN. When IGMP snooping is enabled, a switch examines IGMP messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards multicast traffic only to those interfaces that are connected to relevant receivers instead of flooding the traffic to all interfaces.

This example describes how to configure IGMP snooping:

- [Requirements on page 3727](#)
- [Overview and Topology on page 3728](#)
- [Configuration on page 3729](#)
- [Verifying IGMP Snooping Operation on page 3730](#)

#### Requirements

---

This example uses the following hardware and software components:

- One EX4300 Series switch

- Junos OS Release 13.2X50-D10 or later for EX Series switches

Before you configure IGMP snooping, be sure you have:

- Configured a VLAN, vlan100, on the switch
- Assigned interfaces ge-0/0/0, ge-0/0/1, ge-0/0/2, and ge-0/0/12 to vlan100
- Configured ge-0/0/12 as a trunk interface


See “Configuring VLANs for EX Series Switches (CLI Procedure)” on page 2337.

### Overview and Topology

In this example, interfaces ge-0/0/0, ge-0/0/1, and ge-0/0/2 on the switch are in vlan100 and are connected to hosts that are potential multicast receivers. Interface ge-0/0/12, a trunk interface also in vlan100, is connected to a multicast router. The router acts as the IGMP querier and forwards multicast traffic for group 225.100.100.100 to the switch from a multicast source.

The sample topology is illustrated in [Figure 50 on page 3728](#).

**Figure 50: IGMP Snooping Topology Sample Topology**



In this sample topology, the multicast router forwards multicast traffic to the switch from the source when it receives a membership report for group 255.100.100.100 from one of the hosts—for example, Host B. If IGMP snooping is not enabled on vlan100, the switch floods the multicast traffic on all interfaces in vlan100 (except for interface ge-0/0/12). If IGMP snooping is enabled on vlan100, the switch monitors the IGMP messages between the hosts and router, allowing it to determine that only Host B is interested in receiving the multicast traffic. The switch then forwards the multicast traffic only to interface ge-0/0/1.



This example shows how to perform the following optional configurations, which can reduce group join and leave latency:

- Configure immediate leave on the VLAN. When immediate leave is configured, the switch stops forwarding multicast traffic on an interface when it detects that the last member of the multicast group has left the group. If immediate leave is not configured, the switch waits until the group-specific queries time out before it stops forwarding traffic.

Immediate leave is supported by IGMP version 2 (IGMPv2) and IGMPv3. With IGMPv2, we recommend that you configure immediate leave only when there is only one IGMP host on an interface. In IGMPv2, only one host on a interface sends a membership report in response to a group-specific query—any other interested hosts suppress their reports to avoid a flood of reports for the same group. This report-suppression feature implies that the switch knows about only one interested host at any given time.

- Configure ge-0/0/12 as a static multicast-router interface. In this topology, ge-0/0/12 always leads to the multicast router. By statically configuring ge-0/0/12 as a multicast-router interface, you avoid any delay imposed by the switch having to learn that ge-0/0/12 is a multicast-router interface.

### Configuration

To configure IGMP snooping on a switch:

#### CLI Quick Configuration

To quickly configure IGMP snooping, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols igmp-snooping vlan vlan100 immediate-leave
set protocols igmp-snooping vlan vlan100 interface ge-0/0/12 multicast-router-interface
```

#### Step-by-Step Procedure

To configure IGMP snooping on vlan100:

1. Configure the switch to immediately remove a group membership from an interface when it receives a leave report from the last member of the group on the interface:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan100 immediate-leave
```

2. Statically configure interface ge-0/0/12 as a multicast-router interface:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan100 interface ge-0/0/12
multicast-router-interface
```

**Results** Check the results of the configuration:

```
[edit protocols]
user@switch# show igmp-snooping
vlan vlan100 {
    immediate-leave;
    interface ge-0/0/12.0 {
        multicast-router-interface;
    }
}
```

### Verifying IGMP Snooping Operation

---

To verify that IGMP snooping is operating as configured, perform the following task:

- [Displaying IGMP Snooping Information for VLAN vlan100 on page 3730](#)

#### *Displaying IGMP Snooping Information for VLAN vlan100*

**Purpose** Verify that IGMP snooping is enabled on vlan100 and that ge-0/0/12 is recognized as a multicast-router interface.

**Action** Enter the following command:

```
user@switch>show igmp snooping membership
VLAN: vlan100
    Interfaces: ge-0/0/12.0,
```

**Meaning** By showing information for vlan100, the command output confirms that IGMP snooping is configured on the VLAN. Interface ge-0/0/12.0 is listed as a multicast-router interface, as configured. Because none of the host interfaces are listed, none of the hosts are currently receivers for the multicast group.

**Related Documentation**

- [Configuring IGMP Snooping \(CLI Procedure\)](#)
- [Verifying IGMP Snooping \(CLI Procedure\) on page 3826](#)
- [IGMP Snooping on EX Series Switches Overview on page 3717](#)

## Configuration Tasks

---

- [Configuring IGMP Snooping \(CLI Procedure\) on page 3730](#)
- [Configuring IGMP Snooping \(J-Web Procedure\) on page 3735](#)
- [Configuring IGMP Snooping Tracing Operations \(CLI Procedure\) on page 3737](#)

### Configuring IGMP Snooping (CLI Procedure)



**NOTE:** This topic uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see: [Configuring IGMP Snooping](#). For ELS details see, [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

Internet Group Management Protocol (IGMP) snooping constrains the flooding of IPv4 multicast traffic on a VLAN. When IGMP snooping is enabled, a switch examines IGMP messages between hosts and multicast routers and learns which hosts are interested in receiving multicast traffic for a multicast group. Based on what it learns, the switch then forwards multicast traffic only to those interfaces connected to interested receivers instead of flooding the traffic to all interfaces. In IGMPv2, only one host on a interface sends a membership report in response to a group-specific query—any other interested

hosts suppress their reports to avoid a flood of reports for the same group. This report-suppression feature means that the switch only knows about only one interested host at any given time.

By default, IGMP snooping is enabled on the default VLAN. For many networks, IGMP snooping requires no further configuration.

You can perform the following optional configurations for each VLAN:

- Selectively enable IGMP snooping on specific VLANs.
- Enable immediate leave on a VLAN or all VLANs. Enabling immediate leave ensures that the switch stops forwarding multicast traffic immediately after the last member host on the interface leaves the group.
- Configure an interface as a static multicast-router interface for a VLAN so that the switch does not need to dynamically learn that the interface is a multicast-router interface.
- Configure an interface as a static member of a multicast group so that the switch does not need to dynamically learn the interface's membership.
- Change the value for certain timers and counters to match the values configured on the multicast router serving as the IGMP querier.

This topic covers:

- [Enabling IGMP Snooping on VLANs on page 3731](#)
- [Enabling Immediate Leave on page 3731](#)
- [Configuring an Interface as a Multicast-Router Interface on page 3732](#)
- [Configuring Static Group Membership on an Interface on page 3733](#)
- [Changing the Timer and Counter Values on page 3734](#)

### Enabling IGMP Snooping on VLANs

This topic describes how you can selectively enable IGMP snooping on VLANs. It assumes that you are beginning with the factory default configuration.

- To enable IGMP snooping on a VLAN:

```
[edit protocols igmp-snooping]
user@switch# set vlan vlan-name
```

You can also deactivate the IGMP snooping protocol on the switch without changing the IGMP snooping VLAN configurations:

```
[edit]
user@switch# deactivate protocols igmp-snooping
```

### Enabling Immediate Leave

By default, when a switch with IGMP snooping enabled receives an IGMP leave report on a member interface, it waits for hosts on the interface to respond to IGMP group-specific queries to determine whether there still are hosts on the interface interested

in receiving the group multicast traffic. If the switch does not see any membership reports for the group within a set interval of time, it removes the interface's group membership from the multicast forwarding table and stops forwarding multicast traffic for the group to the interface.

You can decrease the leave latency created by this default behavior by enabling immediate leave on a VLAN.

When you enable immediate leave on a VLAN, host tracking is also enabled, which allows the switch to keep track of the hosts on an interface that have joined a multicast group. When the switch receives a leave report from the last member of the group, it immediately stops forwarding traffic to the interface and does not wait for the interface group membership to time out.

Immediate leave is supported for both IGMP version 2 (IGMPv2) and IGMPv3. However, with IGMPv2, we recommend that you configure immediate leave only when there is only one IGMP host on an interface.

To enable immediate leave on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name immediate-leave
```

---

### Configuring an Interface as a Multicast-Router Interface

---

When IGMP snooping is enabled on a switch, the switch determines which interfaces face a multicast router by monitoring interfaces for IGMP queries. If the switch receives these messages on an interface, it adds the interface to its multicast forwarding table as a multicast-router interface.

In addition to dynamically learned interfaces, the multicast forwarding table can include interfaces that you explicitly configure to be multicast-router interfaces. Unlike the table entries for dynamically learned interfaces, table entries for statically configured interfaces are not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure a static multicast-router interface include:

- You have an unusual network configuration that prevents IGMP snooping from reliably learning about a multicast-router interface through monitoring IGMP queries.
- Your implementation does not require an IGMP querier.
- You have a stable topology and want to avoid the delay the dynamic learning process entails.



**NOTE:** All unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded only to the multicast-router interface.

---

To configure an interface as a static multicast-router interface:

```
[edit protocols]
```

```
user@switch# set igmp-snooping vlan vlan-name interface interface-name
multicast-router-interface
```

For example, to configure ge-0/0/5.0 as a multicast-router interface on a VLAN vlan100 on the switch:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan100 interface ge-0/0/5.0
multicast-router-interface
```

### Configuring Static Group Membership on an Interface

To determine how to forward multicast packets, a switch with IGMP snooping enabled maintains a multicast forwarding table containing a list of host interfaces that have interested listeners for a specific multicast group. The switch learns which host interfaces to add or delete from this table by examining IGMP membership reports as they arrive on interfaces on which IGMP snooping is enabled.

In addition to such dynamically learned interfaces, the multicast forwarding table can include interfaces that you statically configure to be members of multicast groups. When you configure a static group interface, the switch adds the interface to the forwarding table as a host interface for the group. Unlike an entry for a dynamically learned interface, a static interface entry is not subject to aging and deletion from the forwarding table.

Examples of when you might want to configure static group membership on an interface include:

- You want to simulate an attached multicast receiver for testing purposes.
- The interface has receivers that cannot send IGMP membership reports.
- You want the multicast traffic for a specific group to be immediately available to a receiver without any delay imposed by the dynamic join process.



**NOTE:** The switch does not simulate IGMP membership reports on behalf of a statically configured interface. Thus a multicast router might be unaware that the switch has an interface that is a member of the multicast group. You can configure a static group interface on the router to ensure that the switch receives the group multicast traffic.

To configure a host interface as a static member of a multicast group:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name interface interface-name static group
ip-address
```

For example, to configure interface ge-0/0/11.0 in vlan100 as a static member of multicast group 225.0.0.1:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan100 interface ge-0/0/11.0 static group
225.0.0.1
```

## Changing the Timer and Counter Values

---

IGMP uses various timers and counters to determine how often an IGMP querier sends out membership queries and when group memberships time out. On EX Series switches, the default values of the IGMP and IGMP snooping timers and counters are set to the values recommended in RFC 2236, *Internet Group Management Protocol, Version 2*. These values work well for most multicast implementations.

There might be cases, however, where you might want to modify the timer and counter values—for example, to reduce burstiness, to reduce leave latency, or to adjust for expected packet loss on a subnet. If you change a timer or counter value for the IGMP querier on a VLAN, we recommend that you change the value for all multicast routers and switches on the VLAN so that all devices time out group memberships at approximately the same time.

You can configure the following timers and counters on a switch:

- **query-interval**—The length of time the IGMP querier waits between sending general queries (the default is 125 seconds). You can change this interval to tune the number of IGMP messages on the subnet; larger values cause general queries to be sent less often.

You cannot configure this value directly for IGMP snooping. IGMP snooping inherits the value from the IGMP value configured on the switch, which is applied to all VLANs on the switch.

To configure the IGMP **query-interval**:

```
[edit protocols]
user@switch# set igmp-snooping vlan
vlan-name query-interval seconds
```

- **query-response-interval**—The maximum length of time the host can wait until it responds (the default is 10 seconds). You can change this interval to adjust the burst peaks of IGMP messages on the subnet. Set a larger interval to make the traffic less bursty.

You cannot configure this value directly for IGMP snooping. IGMP snooping inherits the value from the IGMP value configured on the switch, which is applied to all VLANs on the switch.

To configure the IGMP **query-response-interval**:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name

query-response-interval seconds
```

- **query-last-member-interval**—The length of time the IGMP querier waits between sending group-specific membership queries (the default is 1 second). The IGMP querier sends a group-specific query after receiving a leave report from a host. You can decrease this interval to reduce the amount of time it takes for multicast traffic to stop forwarding traffic after the last member leaves a group.

You cannot configure this value directly for IGMP snooping. IGMP snooping inherits the value from the IGMP value configured on the switch, which is applied to all VLANs on the switch.

To configure the IGMP **query-last-member-interval**:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name query-last-member-interval seconds
```

- **robust-count**—The number of times the querier resends a general membership query or a group-specific membership query (the default is 2 times). You can increase this count to tune for higher expected packet loss.

For IGMP snooping, you can configure **robust-count** for a specific VLAN. If a VLAN does not have **robust-count** configured, the robust-count value is the value configured for IGMP.

To configure robust-count for IGMP snooping on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name robust-count number
```

The values configured for **query-interval**, **query-response-interval**, and **robust-count** determine the multicast listener interval—the length of time the switch waits for a group membership report after a general query before removing a multicast group from its multicast forwarding table. The switch calculates the multicast listener interval by multiplying **query-interval** by the **robust-count** and then adding **query-response-interval** to the product:

$(\text{query-interval} * \text{robust-count}) + \text{query-response-interval} = \text{multicast listener interval}$ .

For example, the multicast listener interval is 260 seconds when the default settings for **query-interval**, **query-response-interval**, and **robust-count** are used:

$(125 * 2) + 10 = 260$

You can display the time remaining in the multicast listener interval before a group times out by using the **show igmp-snooping membership** command.

#### Related Documentation

- [Example: Configuring IGMP Snooping on page 3727](#)
- [Verifying IGMP Snooping \(CLI Procedure\) on page 3826](#)
- [IGMP Snooping on EX Series Switches Overview on page 3717](#)

## Configuring IGMP Snooping (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

IGMP snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, the EX Series switch monitors the IGMP transmissions between a host (a network device) and a multicast router, keeping track of the multicast groups and associated

member interfaces. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces.

You can configure IGMP snooping on one or more VLANs to allow the switch to examine IGMP packets and make forwarding decisions based on packet content. By default, IGMP snooping is enabled on EX Series switches.

To enable IGMP snooping and configure individual options by using the J-Web interface:

1. Select **Configure > Switching > IGMP Snooping**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—Creates an IGMP snooping configuration for the VLAN.
- **Edit**—Modifies an IGMP snooping configuration for the VLAN.
- **Delete**—Deletes a selected VLAN from the IGMP snooping configuration.

When you are adding or editing an IGMP snooping configuration, enter information as described in [Table 388 on page 3736](#).

3. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.

To disable IGMP snooping on a VLAN, select the VLAN from the list and click **Disable**.



**NOTE:** The **Disable** option is not available for EX4300 switches.

**Table 388: IGMP Snooping Configuration Fields**

| Field           | Function                                                                                                                                                                                                                      | Your Action                                                                                    |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
| VLAN Name       | Specifies the VLAN on which to enable IGMP snooping.                                                                                                                                                                          | Select a VLAN from the list to add it to the snooping configuration.                           |
| Immediate Leave | Immediately removes a multicast group membership from an interface when it receives a leave message from that interface without waiting for any other IGMP messages to be exchanged (IGMP version 2 and IGMP version 3 only). | To enable the option, select the check box.<br><br>To disable the option, clear the check box. |
| Robust Count    | Specifies the number of timeout intervals the switch waits before timing out a multicast group.                                                                                                                               | Type a value.                                                                                  |



Table 388: IGMP Snooping Configuration Fields (*continued*)

| Field           | Function                                                                                                                   | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------|----------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interfaces List | Statically configures an interface as a switching interface toward a multicast router or as a member of a multicast group. | <p>Click one of the following options:</p> <ul style="list-style-type: none"> <li>• <b>Add</b>—Adds an interface to the IGMP snooping configuration. <ol style="list-style-type: none"> <li>1. Select an interface from the list. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.</li> <li>2. Select <b>Multicast Router Interface</b>.</li> <li>3. Type the maximum number of groups an interface can join.</li> <li>4. In <b>Static</b>, choose one: <ul style="list-style-type: none"> <li>• Click <b>Add</b>, type a group IP address, and click <b>OK</b>.</li> <li>• Select a group and click <b>Remove</b> to remove the group membership.</li> </ul> </li> </ol> </li> <li>• <b>Edit</b>—Edits the interface settings for the IGMP snooping configuration.</li> <li>• <b>Remove</b>—Deletes an interface configured for IGMP snooping.</li> </ul> |

- Related Documentation**
- *Example: Configuring IGMP Snooping on EX Series Switches*
  - *Configuring IGMP Snooping (CLI Procedure)*
  - [IGMP Snooping on EX Series Switches Overview on page 3717](#)

## Configuring IGMP Snooping Tracing Operations (CLI Procedure)

By enabling tracing operations for IGMP snooping, you can record detailed messages about the operation of the protocol, such as the various types of protocol packets sent and received. [Table 389 on page 3737](#) describes the tracing operations you can enable and the flags used to specify them in the tracing configuration.

Table 389: Supported Tracing Operations for IGMP Snooping

| Tracing Operation                               | Flag           |
|-------------------------------------------------|----------------|
| Trace all (equivalent of including all flags).  | <b>all</b>     |
| Trace general IGMP snooping protocol events.    | <b>general</b> |
| Trace communication over routing socket events. | <b>krt</b>     |
| Trace leave reports (IGMPv2 and IGMPv3 only).   | <b>leave</b>   |

Table 389: Supported Tracing Operations for IGMP Snooping (*continued*)

| Tracing Operation                                                                                                            | Flag           |
|------------------------------------------------------------------------------------------------------------------------------|----------------|
| Trace nexthop-related events.                                                                                                | <b>nexthop</b> |
| Trace normal IGMP snooping protocol events. If you do not specify this flag, only unusual or abnormal operations are traced. | <b>normal</b>  |
| Trace all IGMP packets.                                                                                                      | <b>packets</b> |
| Trace policy processing.                                                                                                     | <b>policy</b>  |
| Trace IGMP membership query messages.                                                                                        | <b>query</b>   |
| Trace membership reports                                                                                                     | <b>report</b>  |
| Trace routing information.                                                                                                   | <b>route</b>   |
| Trace state transitions.                                                                                                     | <b>state</b>   |
| Trace routing protocol task processing.                                                                                      | <b>task</b>    |
| Trace timer processing.                                                                                                      | <b>timer</b>   |
| Trace VLAN-related events.                                                                                                   | <b>vlan</b>    |

This topic covers:

- [Configuring Tracing Operations on page 3738](#)
- [Viewing, Stopping, and Restarting Tracing Operations on page 3739](#)

### Configuring Tracing Operations

To configure tracing operations for IGMP snooping:

1. Configure the filename for the trace file:

```
[edit protocols igmp-snooping ]
user@switch# set traceoptions file filename
```

For example:

```
[edit protocols igmp-snooping ]
user@switch# set traceoptions file mld-snoop-trace
```

2. (Optional) Configure the maximum number of trace files and size of the trace files:

```
[edit protocols igmp-snooping ]
user@switch # set file files number size size
```

For example:

```
[edit protocols igmp-snooping ]
user@switch # set traceoptions file files 5 size 1m
```

causes the contents of the trace file to be emptied and archived in a .gz file when the file reaches 1 MB. Four archive files are maintained, the contents of which are rotated whenever the current active trace file is archived.

If you omit this step, the maximum number of trace files defaults to 10, with the maximum file size defaulting to 128 K.

3. Specify one of the tracing flags shown in [Table 389 on page 3737](#):

```
[edit protocols igmp-snooping ]
user@switch # set traceoptions flag flagname
```

For example, to perform trace operations on VLAN-related events and IGMP query messages:

```
[edit protocols igmp-snooping ]
user@switch# set traceoptions flag vlan

[edit protocols igmp-snooping ]
user@switch# set traceoptions flag query
```

### Viewing, Stopping, and Restarting Tracing Operations

When you commit the configuration, tracing operations begin. You can view the trace file in the `/var/log` directory. For example:

```
user@switch> file show /var/log/igmp-snoop-trace
```

You can stop and restart tracing operations by deactivating and reactivating the configuration:

```
[edit]
user@switch# deactivate protocols igmp-snooping traceoptions

[edit]
user@switch# activate protocols igmp-snooping traceoptions
```

#### Related Documentation

- [Configuring IGMP Snooping \(CLI Procedure\)](#)
- [Tracing and Logging Junos OS Operations](#)

## Configuration Statements

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 3741](#)
- [\[edit protocols igmp-snooping\] Configuration Statement Hierarchy on page 3743](#)
- [accounting \(Protocols IGMP Interface\) on page 3745](#)
- [accounting \(Protocols IGMP\) on page 3745](#)
- [address \(Anycast RPs\) on page 3746](#)
- [address \(Local RPs\) on page 3746](#)
- [anycast-pim on page 3747](#)
- [assert-timeout on page 3748](#)
- [auto-rp on page 3749](#)

- [bootstrap on page 3750](#)
- [bootstrap-export on page 3751](#)
- [bootstrap-import on page 3752](#)
- [bootstrap-priority on page 3753](#)
- [data-forwarding on page 3754](#)
- [dense-groups on page 3755](#)
- [disable \(IGMP Snooping\) on page 3755](#)
- [disable \(Protocols IGMP\) on page 3756](#)
- [disable \(PIM\) on page 3757](#)
- [dr-election-on-p2p on page 3758](#)
- [dr-register-policy on page 3758](#)
- [embedded-rp on page 3759](#)
- [export \(Protocols PIM Bootstrap\) on page 3760](#)
- [family \(Bootstrap\) on page 3761](#)
- [family \(Local RP\) on page 3762](#)
- [graceful-restart \(Protocols PIM\) on page 3763](#)
- [group \(IGMP Snooping\) on page 3763](#)
- [group \(Protocols IGMP\) on page 3764](#)
- [group-ranges on page 3765](#)
- [hello-interval \(Protocols PIM\) on page 3766](#)
- [hold-time \(Protocols PIM\) on page 3767](#)
- [igmp-snooping on page 3768](#)
- [immediate-leave \(Protocols IGMP\) on page 3769](#)
- [immediate-leave \(IGMP Snooping\) on page 3770](#)
- [import \(Protocols PIM Bootstrap\) on page 3771](#)
- [import \(Protocols PIM\) on page 3772](#)
- [infinity on page 3773](#)
- [interface \(IGMP Snooping\) on page 3774](#)
- [interface \(Protocols PIM\) on page 3775](#)
- [interface \(Protocols IGMP\) on page 3777](#)
- [join-load-balance on page 3778](#)
- [local on page 3779](#)
- [local-address \(Protocols PIM\) on page 3780](#)
- [mapping-agent-election on page 3781](#)
- [maximum-rps on page 3782](#)
- [mode \(Protocols PIM\) on page 3783](#)
- [multicast-router-interface \(IGMP Snooping\) on page 3784](#)

- [neighbor-policy](#) on page 3785
- [pim](#) on page 3786
- [priority \(PIM Interfaces\)](#) on page 3790
- [priority \(Bootstrap\)](#) on page 3791
- [priority \(PIM RPs\)](#) on page 3792
- [query-interval \(Protocols IGMP\)](#) on page 3793
- [query-last-member-interval \(Protocols IGMP\)](#) on page 3794
- [query-response-interval \(Protocols IGMP\)](#) on page 3795
- [receiver](#) on page 3796
- [restart-duration \(Protocols PIM\)](#) on page 3797
- [rib-group \(Protocols PIM\)](#) on page 3798
- [robust-count \(IGMP Snooping\)](#) on page 3799
- [robust-count \(Protocols IGMP\)](#) on page 3799
- [rp](#) on page 3800
- [rp-register-policy](#) on page 3802
- [rp-set](#) on page 3803
- [source \(Protocols IGMP\)](#) on page 3804
- [source-vlans](#) on page 3805
- [spt-threshold](#) on page 3806
- [ssm-map \(Protocols IGMP\)](#) on page 3807
- [static \(IGMP Snooping\)](#) on page 3807
- [static \(Protocols PIM\)](#) on page 3808
- [static \(Protocols IGMP\)](#) on page 3809
- [traceoptions \(Protocols PIM\)](#) on page 3810
- [traceoptions \(Protocols IGMP\)](#) on page 3813
- [traceoptions \(IGMP Snooping\)](#) on page 3816
- [version \(Protocols IGMP\)](#) on page 3818
- [version \(IGMP Snooping\)](#) on page 3818
- [version \(PIM\)](#) on page 3819
- [vlan \(IGMP Snooping\)](#) on page 3820

## [\[edit protocols\]](#) Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the [\[edit protocols\]](#) hierarchy:

- [\[edit protocols bfd\]](#) Configuration Statement Hierarchy on EX Series Switches on page 391
- [\[edit protocols bgp\]](#) Configuration Statement Hierarchy on EX Series Switches on page 392

- [\[edit protocols connections\]](#) Configuration Statement Hierarchy on EX Series Switches on page 401
- [\[edit protocols dcbx\]](#) Configuration Statement Hierarchy on EX Series Switches on page 403
- [\[edit protocols dot1x\]](#) Configuration Statement Hierarchy on EX Series Switches on page 404
- [\[edit protocols igmp\]](#) Configuration Statement Hierarchy on EX Series Switches on page 406
- [\[edit protocols igmp-snooping\]](#) Configuration Statement Hierarchy on EX Series Switches on page 407
- [\[edit protocols isis\]](#) Configuration Statement Hierarchy on EX Series Switches on page 408
- [\[edit protocols lacp\]](#) Configuration Statement Hierarchy on EX Series Switches on page 411
- [\[edit protocols link-management\]](#) Configuration Statement Hierarchy on EX Series Switches on page 412
- [\[edit protocols lldp\]](#) Configuration Statement Hierarchy on EX Series Switches on page 413
- [\[edit protocols lldp-med\]](#) Configuration Statement Hierarchy on EX Series Switches on page 415
- [\[edit protocols mld\]](#) Configuration Statement Hierarchy on EX Series Switches on page 416
- [\[edit protocols mld-snooping\]](#) Configuration Statement Hierarchy on EX Series Switches on page 417
- [\[edit protocols mpls\]](#) Configuration Statement Hierarchy on EX Series Switches on page 418
- [\[edit protocols msdp\]](#) Configuration Statement Hierarchy on EX Series Switches on page 429
- [\[edit protocols mstp\]](#) Configuration Statement Hierarchy on EX Series Switches on page 431
- [\[edit protocols mvrp\]](#) Configuration Statement Hierarchy on EX Series Switches on page 433
- [\[edit protocols neighbor-discovery\]](#) Configuration Statement Hierarchy on EX Series Switches on page 434
- [\[edit protocols oam\]](#) Configuration Statement Hierarchy on EX Series Switches on page 435
- [\[edit protocols ospf\]](#) Configuration Statement Hierarchy on EX Series Switches on page 438
- [\[edit protocols ospf3\]](#) Configuration Statement Hierarchy on EX Series Switches on page 441

- [\[edit protocols pim\] Configuration Statement Hierarchy on EX Series Switches on page 444](#)
- [\[edit protocols rip\] Configuration Statement Hierarchy on EX Series Switches on page 447](#)
- [\[edit protocols ripng\] Configuration Statement Hierarchy on EX Series Switches on page 450](#)
- [\[edit protocols router-advertisement\] Configuration Statement Hierarchy on EX Series Switches on page 451](#)
- [\[edit protocols router-discovery\] Configuration Statement Hierarchy on EX Series Switches on page 452](#)
- [\[edit protocols rstp\] Configuration Statement Hierarchy on EX Series Switches on page 453](#)
- [\[edit protocols rsvp\] Configuration Statement Hierarchy on EX Series Switches on page 455](#)
- [\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches on page 459](#)
- [\[edit protocols stp\] Configuration Statement Hierarchy on EX Series Switches on page 460](#)
- [\[edit protocols uplink-failure-detection\] Configuration Statement Hierarchy on EX Series Switches on page 461](#)
- [\[edit protocols vrrp\] Configuration Statement Hierarchy on EX Series Switches on page 462](#)
- [\[edit protocols vstp\] Configuration Statement Hierarchy on EX Series Switches on page 463](#)

**Related  
Documentation**

- *EX Series Switch Software Features Overview*
- *EX Series Virtual Chassis Software Features Overview*

## [\[edit protocols igmp-snooping\] Configuration Statement Hierarchy](#)

This topic lists supported and unsupported configuration statements in the [\[edit protocols igmp-snooping\]](#) hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols igmp-snooping\] Hierarchy Level on page 3744](#)
- [Unsupported Statements in the \[edit protocols igmp-snooping\] Hierarchy Level on page 3744](#)

### Supported Statements in the [edit protocols igmp-snooping] Hierarchy Level

The following hierarchy shows the **[edit protocols igmp-snooping]** configuration statements supported on EX Series switches:

```
protocols {
  igmp-snooping {
    vlan vlan-name {
      immediate-leave;
      interface interface-name {
        group-limit <1..65535>
        host-only-interface
        multicast-router-interface;
        immediate-leave;
        static {
          group multicast-ip-address {
            source <>
          }
        }
      }
    }
  }
  l2-querier {
    source-address ip-address;
  }
  proxy {
    source-address ip-address;
  }
  query-interval number;
  query-last-member-interval number;
  query-response-interval number;
  robust-count number;
  traceoptions {
    file filename <files number> <no-stamp> <replace> <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag <flag-modifier>;
  }
}
```

### Unsupported Statements in the [edit protocols igmp-snooping] Hierarchy Level

All statements in the **[edit protocols igmp-snooping]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)



## accounting (Protocols IGMP Interface)

---

|                                 |                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (accounting   no-accounting);                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp <a href="#">interface</a> <i>interface-name</i> ],<br>[edit protocols igmp <a href="#">interface</a> <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Enable or disable the collection of IGMP join and leave event statistics for an interface.                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Recording IGMP Join and Leave Events</i></li> </ul>                                                                                              |

## accounting (Protocols IGMP)

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | accounting;                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp],<br>[edit protocols igmp]                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Enable the collection of IGMP join and leave event statistics on the system.                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Recording IGMP Join and Leave Events</i></li> </ul>                                                                                            |

## address (Anycast RPs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>address address &lt;forward-msdp-sa&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>pim rp local</code> (inet   inet6) <code>anycast-pim rp-set</code> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim rp local</code> (inet   inet6) <code>anycast-pim rp-set</code> ],<br>[edit protocols <code>pim rp local</code> (inet   inet6) <code>anycast-pim rp-set</code> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim rp local</code> (inet   inet6) <code>anycast-pim rp-set</code> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the anycast rendezvous point (RP) addresses in the RP set. Multiple addresses can be configured in an RP set. If the RP has peer Multicast Source Discovery Protocol (MSDP) connections, then the RP must forward MSDP source active (SA) messages.                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>address</b> —RP address in an RP set.<br><br><b>forward-msdp-sa</b> —(Optional) Forward MSDP SAs to this address.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## address (Local RPs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>address address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>pim rp local family</code> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim rp local family</code> (inet   inet6)],<br>[edit protocols <code>pim rp local family</code> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim rp local family</code> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the local rendezvous point (RP) address.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <b>address</b> —Local RP address.                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Local PIM RPs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                  |

## anycast-pim

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>anycast-pim {   rp-set {     address address &lt;forward-msdp-sa&gt;;   } }</pre>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim rp local family</b> (inet   inet6)],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6)],<br/> [edit protocols <b>pim rp local family</b> (inet   inet6)],<br/> [edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.<br/> Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/> Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Configure properties for anycast RP using PIM.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/> routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring PIM Anycast With or Without MSDP</i></li> </ul>                                                                                                                                                                                                                                                                                                                                         |

## assert-timeout

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>assert-timeout seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ]                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Multicast routing devices running PIM sparse mode often forward the same stream of multicast packets onto the same LAN through the rendezvous-point tree (RPT) and shortest-path tree (SPT). PIM assert messages help routing devices determine which routing device forwards the traffic and prunes the RPT for this group. By default, routing devices enter an assert cycle every 180 seconds. You can configure this assert timeout to be between 5 and 210 seconds. |
| <b>Options</b>                  | <b>seconds</b> —Time for routing device to wait before another assert message cycle.<br><b>Range:</b> 5 through 210 seconds<br><b>Default:</b> 180 seconds                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring the PIM Assert Timeout</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                     |

## auto-rp

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> auto-rp {   (announce   discovery   mapping);   (mapping-agent-election   no-mapping-agent-election); } </pre>                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit protocols <a href="#">pim rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>]</p>                                                                           |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure automatic RP announcement and discovery.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>announce</b>—Configure the routing device to listen only for mapping packets and also to advertise itself if it is an RP.</p> <p><b>discovery</b>—Configure the routing device to listen only for mapping packets.</p> <p><b>mapping</b>—Configures the routing device to announce, listen for and generate mapping packets, and announce that the routing device is eligible to be an RP.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Auto-RP</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                           |

## bootstrap

---

|                                 |                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bootstrap {<br/>    family (inet   inet6) {<br/>        export [ <i>policy-names</i> ];<br/>        import [ <i>policy-names</i> ];<br/>        priority <i>number</i>;<br/>    }<br/>}</pre>                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>pim rp</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim rp</b> ],<br>[edit protocols <b>pim rp</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                      |
| <b>Description</b>              | Configure parameters to control bootstrap routers and messages.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li></ul>                                                                                                                                                   |

## bootstrap-export

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit protocols <a href="#">pim rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more export policies to control outgoing PIM bootstrap messages.                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> <li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li> <li>• <a href="#">bootstrap-import on page 3752</a></li> </ul>                                                                                                                                          |

## bootstrap-import

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ],<br>[edit protocols <a href="#">pim rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more import policies to control incoming PIM bootstrap messages.                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li><li>• <a href="#">bootstrap-export on page 3751</a></li></ul>                                                                                                                               |



## bootstrap-priority

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit protocols <a href="#">pim rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                      |
| <b>Description</b>              | Configure whether this routing device is eligible to be a bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router.                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b><i>number</i></b>—Priority for becoming the bootstrap router. A value of 0 means that the routing device is not eligible to be the bootstrap router.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 0</p>                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> </ul>                                                                                                                                                                                                                                                                           |

## data-forwarding

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>data-forwarding {<br/>  receiver {<br/>    source-vlans <i>vlan-list</i>;<br/>    install;<br/>  }<br/>  source {<br/>    groups <i>group-prefix</i>;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> )]                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Configure the VLAN to be a multicast source VLAN (MVLAN) or a multicast VLAN registration (MVR) receiver VLAN. Each data-forwarding VLAN, which can be a multicast source VLAN (MVLAN) or a multicast receiver VLAN, must have exactly one source statement or exactly one receiver statement. A data-forwarding VLAN can operate only in IGMP version 2 (IGMPv2) mode.</p> <p>The remaining statements are explained separately.</p> |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast VLAN Registration</i></li><li>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i></li></ul>                                                                                                                                                                                                                                                       |

## dense-groups

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dense-groups {<br/>    addresses;<br/>}</code>                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Configure which groups are operating in dense mode.                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>addresses</b> —Address of groups operating in dense mode.                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Sparse-Dense Mode Properties</i></li> </ul>                                                                                                                                                                                                                                                 |

## disable (IGMP Snooping)

|                                 |                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>disable;</code>                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">igmp-snooping</a> <a href="#">vlan</a> (all   <i>vlan-name</i> )]                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2 for EX Series switches.                                                                            |
| <b>Description</b>              | Disable IGMP snooping on the VLAN. Multicast traffic will be flooded to all interfaces on the VLAN except the source interface.                 |
| <b>Default</b>                  | If you do not include this statement in the configuration for a VLAN, IGMP snooping is enabled on the VLAN.                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring IGMP Snooping (CLI Procedure)</i></li> <li>• <i>show igmp-snooping vlans</i></li> </ul> |

## disable (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp <a href="#">interface</a> <i>interface-name</i> ],<br>[edit protocols igmp <a href="#">interface</a> <i>interface-name</i> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Disable IGMP on the system.                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Disabling IGMP</i></li></ul>                                                                                                                        |

## disable (PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim</b> family (inet   inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim rp local</b> family (inet   inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim rp local</b> family (inet   inet6)],</p> <p>[edit protocols <b>pim</b>],</p> <p>[edit protocols <b>pim</b> family (inet   inet6)],</p> <p>[edit protocols <b>pim interface</b> <i>interface-name</i>],</p> <p>[edit protocols <b>pim rp local</b> family (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim</b> family (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp local</b> family (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>disable</b> statement extended to the <b>[family]</b> hierarchy level in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Explicitly disable PIM at the protocol, interface or family hierarchy levels.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Disabling PIM</li> <li>disable (PIM Graceful Restart)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## dr-election-on-p2p

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | dr-election-on-p2p;                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 9.1 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                              |
| <b>Description</b>              | Enable PIM designated router (DR) election on point-to-point (P2P) links.                                                                                                                                                                                                                                                                               |
| <b>Default</b>                  | No PIM DR election is performed on point-to-point links.                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Designated Router Election on Point-to-Point Links</i></li></ul>                                                                                                                                                                                                                             |

## dr-register-policy

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | dr-register-policy [ <i>policy-names</i> ];                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ],<br>[edit protocols <a href="#">pim rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                          |
| <b>Description</b>              | Apply one or more policies to control outgoing PIM register messages.                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Register Message Filters on a PIM RP and DR</i></li><li>• <a href="#">rp-register-policy on page 3802</a></li></ul>                                                                                                                                                                                          |

## embedded-rp

|                                 |                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> embedded-rp {   group-ranges {     destination-ip-prefix &lt;/prefix-length&gt;;   }   maximum-rps limit; } </pre>                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim rp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim rp</b>],</p> <p>[edit protocols <b>pim rp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                  |
| <b>Description</b>              | <p>Configure properties for embedded IP version 6 (IPv6) RPs.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Embedded RP for IPv6</i></li> </ul>                                                                                                                                                                                                                                                |

## export (Protocols PIM Bootstrap)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp bootstrap family</a> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap family</a> (inet   inet6)],<br>[edit protocols <a href="#">pim rp bootstrap family</a> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap family</a> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more export policies to control outgoing PIM bootstrap messages.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li><li>• <a href="#">import (Protocols PIM Bootstrap) on page 3771</a></li></ul>                                                                                                                                                                                                                                           |



## family (Bootstrap)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>family (inet   inet6) {     export [ <i>policy-names</i> ];     import [ <i>policy-names</i> ];     priority <i>number</i>; }</pre>                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp bootstrap</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap</a>],</p> <p>[edit protocols <a href="#">pim rp bootstrap</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure which IP protocol type bootstrap properties to apply.                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>inet</b>—Apply IP version 4 (IPv4) local RP properties.</p> <p><b>inet6</b>—Apply IPv6 local RP properties.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> <li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li> </ul>                                                                                                                                                                                                                                           |

## family (Local RP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>family (inet   inet6) {<br/>    disable;<br/>    address address;<br/>    anycast-pim {<br/>        local-address address;<br/>        rp-set {<br/>            address address &lt;forward-msdp-sa&gt;;<br/>        }<br/>    }<br/>    group-ranges {<br/>        destination-ip-prefix &lt;/prefix-length&gt;;<br/>    }<br/>    hold-time seconds;<br/>    override;<br/>    priority number;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp local</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp local</a> ],<br>[edit protocols <a href="#">pim rp local</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp local</a> ]                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure which IP protocol type local RP properties to apply.                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>inet</b> —Apply IP version 4 (IPv4) local RP properties.<br><br><b>inet6</b> —Apply IPv6 local RP properties.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Local PIM RPs</i></li></ul>                                                                                                                                                                                                                                                                                                                                 |

## graceful-restart (Protocols PIM)


|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | graceful-restart {<br>disable;<br>no-bidirectional-mode;<br>restart-duration seconds;<br>}                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Configure PIM sparse mode graceful restart.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Sparse Mode Graceful Restart</a></li> </ul>                                                                                                                                                                                                                                        |

## group (IGMP Snooping)

|                                 |                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | group <i>ip-address</i> ;                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">igmp-snooping</a> <a href="#">vlan</a> (all   <i>vlan-name</i> ) <a href="#">interface</a> (all   <i>interface-name</i> ) <a href="#">static</a> ]      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                |
| <b>Description</b>              | Configure a static multicast group on an interface.                                                                                                                                 |
| <b>Options</b>                  | <i>ip-address</i> —Valid IP multicast address for the multicast group.                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring IGMP Snooping (CLI Procedure)</a></li> <li>• <a href="#">show igmp-snooping membership on page 3871</a></li> </ul> |

## group (Protocols IGMP)

---

|                                                                                                                                                                     |                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                       | <pre>group <i>multicast-group-address</i> {<br/>  exclude;<br/>  group-count <i>number</i>;<br/>  group-increment <i>increment</i>;<br/>  source <i>ip-address</i> {<br/>    source-count <i>number</i>;<br/>    source-increment <i>increment</i>;<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>                                                                                                                                              | [edit logical-systems <i>logical-system-name</i> protocols igmp <b>interface</b> <i>interface-name</i> <b>static</b> ],<br>[edit protocols igmp <b>interface</b> <i>interface-name</i> <b>static</b> ]                                                                  |
| <b>Release Information</b>                                                                                                                                          | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                          |
| <b>Description</b>                                                                                                                                                  | Specify the IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.                                                                                                                       |
| <hr/>                                                                                                                                                               |                                                                                                                                                                                                                                                                         |
| <div> <b>NOTE:</b> You must specify a unique address for each group.</div> <hr/> |                                                                                                                                                                                                                                                                         |
| The remaining statements are explained separately.                                                                                                                  |                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                     | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                        | <ul style="list-style-type: none"><li>• <i>Enabling IGMP Static Group Membership</i></li></ul>                                                                                                                                                                          |

## group-ranges

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>group-ranges {     destination-ip-prefix&lt;/prefix-length&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp embedded-rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp embedded-rp</a>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols <a href="#">pim rp embedded-rp</a>],</p> <p>[edit protocols <a href="#">pim rp local family</a> (inet   inet6)],</p> <p>[edit protocols <a href="#">pim rp static</a> address <i>address</i>],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp embedded-rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp local family</a> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp static</a> address <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 13.3 for the PTX5000 router.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Configure the address ranges of the multicast groups for which this routing device can be a rendezvous point (RP).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | The routing device is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b><i>destination-ip-prefix&lt;/prefix-length&gt;</i></b> —Addresses or address ranges for which this routing device can be an RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Local PIM RPs in the Multicast Protocols Feature Guide for Routing Devices</i></li> <li>• <i>Configuring PIM Embedded RP for IPv6 in the Multicast Protocols Feature Guide for Routing Devices</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## hello-interval (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hello-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br><a href="#">pim interface interface-name</a> ],<br>[edit protocols <a href="#">pim interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Specify how often the routing device sends PIM hello packets out of an interface.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>seconds</b> —Length of time between PIM hello packets.<br><b>Range:</b> 0 through 255<br><b>Default:</b> 30 seconds                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">hold-time on page 3767</a></li><li>• <i>Modifying the PIM Hello Interval</i></li></ul>                                                                                                                                                                                                                                                                                                                     |

## hold-time (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hold-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim<br/>         rp bidirectional address <i>address</i>],<br/>         [edit protocols pim rp bidirectional address <i>address</i>],<br/>         [edit protocols <b>pim rp local family</b> (inet   inet6)],<br/>         [edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/>         Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>         Statement introduced in Junos OS Release 11.3 for the QFX Series.<br/>         Support for bidirectional RP addresses introduced in Junos OS Release 12.1.<br/>         Statement introduced in Junos OS Release 13.3 for the PTX5000 router.</p>                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Specify the time period for which a neighbor is to consider the sending routing device (this routing device) to be operative (up).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>seconds</b>—Hold time.<br/> <b>Range:</b> 0 through 255<br/> <b>Default:</b> 150 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/>         routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Local PIM RPs in the Multicast Protocols Feature Guide for Routing Devices</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## igmp-snooping

---

**Syntax** `igmp-snooping {  
 traceoptions {  
 file filename <files number> <no-stamp> <replace> <size size> <world-readable |  
 no-world-readable>;  
 flag flag <flag-modifier>;  
 }  
 vlan (all | vlan-name) {  
 data-forwarding {  
 source {  
 groups group-prefix;  
 }  
 receiver {  
 source-vlans vlan-list;  
 install;  
 }  
 }  
 disable;  
 immediate-leave;  
 interface (all | interface-name) {  
 multicast-router-interface;  
 static {  
 group ip-address;  
 }  
 }  
 proxy {  
 source-address ip-address;  
 }  
 robust-count number;  
 version version;  
 }  
}`

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 9.1 for EX Series switches.

**Description** Configure IGMP snooping. The factory default configuration enables IGMP snooping on all VLANs.

The remaining statements are explained separately.


**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring IGMP Snooping on EX Series Switches*
- *Configuring IGMP Snooping (CLI Procedure)*



## immediate-leave (Protocols IGMP)


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>immediate-leave;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp <b>interface</b> <i>interface-name</i> ],<br>[edit protocols igmp <b>interface</b> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.</p> <p>Starting in Junos OS Release 9.3, both IGMP version 2 and IGMP version 3 do host tracking when the <b>immediate-leave</b> statement is configured. This means that the multicast group leaves only when the last host leaves. The routing device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</p> <p>When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.</p> <p>When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.</p> |
|                                 | <p> <b>NOTE:</b> Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the routing device only knows about the one interested host and does not know about the others.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

- Related Documentation**
- *Specifying Immediate-Leave Host Removal for IGMP*

---

## immediate-leave (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | immediate-leave;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols <b>igmp-snooping</b> <b>vlan</b> (all   <i>vlan-name</i> )]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure IGMP snooping immediate leave for the specified VLAN. When you configure immediate leave, host tracking is enabled, which allows the switch to track the hosts that send membership reports. The switch can then determine when the last host on an interface leaves the multicast group and immediately stop forwarding multicast traffic to the interface.</p> <p>Configuring immediate leave reduces the amount of time it takes for the switch to stop sending multicast traffic to an interface when the last host leaves the group. When immediate leave is disabled, the switch no longer tracks hosts. Instead, whenever it receives a leave report from a host, it sends out a group-specific query to all hosts. If it does not receive any membership reports on the interface in response to the group-specific query within a set interval, it stops forwarding multicast traffic to the interface.</p> |
|                                 | <div> <b>NOTE:</b> Immediate leave is supported for both IGMP version 2 (IGMPv2) and IGMPv3. However, with IGMPv2, we recommend that you configure immediate leave only when there is only one IGMP host on an interface. In IGMPv2, only one host on a interface sends a membership report in response to a general query—any other interested hosts suppress their reports. Report suppression avoids a flood of reports for the same group, but it also interferes with host tracking because the switch knows only about one interested host on the interface at any given time.</div>                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | The immediate-leave feature is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping on EX Series Switches</i></li><li>• <i>Configuring IGMP Snooping (CLI Procedure)</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## import (Protocols PIM Bootstrap)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp bootstrap</a> (inet   inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap</a> (inet   inet6)],</p> <p>[edit protocols <a href="#">pim rp bootstrap</a> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap</a> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Apply one or more import policies to control incoming PIM bootstrap messages.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> <li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li> <li>• <a href="#">export (Protocols PIM Bootstrap) on page 3760</a></li> </ul>                                                                                                                                                                                                                              |

## import (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Apply one or more policies to routes being imported into the routing table from PIM. Use the <b>import</b> statement to filter PIM join messages and prevent them from entering the network.                                                                                                                                                            |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Filtering Incoming PIM Join Messages</i></li></ul>                                                                                                                                                                                                                                                           |

## infinity

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>infinity [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim spt-threshold</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim spt-threshold</a>],</p> <p>[edit protocols <a href="#">pim spt-threshold</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim spt-threshold</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the <b>infinity</b> statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair.                                                                                                                                  |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the PIM SPT Threshold Policy</i></li> </ul>                                                                                                                                                                                                                                                                                                                   |

## interface (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface (all   <i>interface-name</i>) {<br/>    multicast-router-interface;<br/>    static {<br/>        group <i>ip-address</i>;<br/>    }<br/>}</pre>                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <b>igmp-snooping</b> <b>vlan</b> (all   <i>vlan-name</i> )]                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                   |
| <b>Description</b>              | For IGMP snooping, configure an interface as either a multicast-router interface or as a static member of a multicast group.                                                                                           |
| <b>Options</b>                  | <p><b>all</b>—All interfaces in the VLAN.</p> <p><b><i>interface-name</i></b>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>                                                     |
| <b>Required Privilege Level</b> | <p><b>routing</b>—To view this statement in the configuration.</p> <p><b>routing-control</b>—To add this statement to the configuration.</p>                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping on EX Series Switches</i></li><li>• <i>Configuring IGMP Snooping (CLI Procedure)</i></li><li>• <i>show igmp-snooping vlans</i></li></ul> |

## interface (Protocols PIM)

```
Syntax interface (Protocols PIM) (all | interface-name) {
    accept-remote-source;
    disable;
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (0 | 1 | automatic);
    }
    bidirectional {
        df-election {
            backoff-period milliseconds;
            offer-period milliseconds;
            robustness-count number;
        }
    }
    family (inet | inet6) {
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (0 | 1 | automatic);
        }
        disable;
    }
    hello-interval seconds;
}
```

```
mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse | sparse-dense);
neighbor-policy [ policy-names ];
override-interval milliseconds;
priority number;
propagation-delay milliseconds;
reset-tracking-bit;
version version;
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [pim](#)],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
[pim](#)],  
[edit protocols [pim](#)],  
[edit routing-instances *routing-instance-name* protocols [pim](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Enable PIM on an interface and configure interface-specific properties.

**Options** *interface-name*—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify [all](#).

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [PIM on Aggregated Interfaces](#)



## interface (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interface <i>interface-name</i> {   disable;   (accounting   no-accounting);   group-limit <i>limit</i>;   group-policy [ <i>policy-names</i> ];   immediate-leave;   oif-map <i>map-name</i>;   passive;   promiscuous-mode;   ssm-map <i>ssm-map-name</i>;   ssm-map-policy <i>ssm-map-policy-name</i>;   static {     group <i>multicast-group-address</i> {       exclude;       group-count <i>number</i>;       group-increment <i>increment</i>;       source <i>ip-address</i> {         source-count <i>number</i>;         source-increment <i>increment</i>;       }     }   }   version <i>version</i>; } </pre> |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp],<br>[edit protocols igmp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Enable IGMP on an interface and configure interface-specific properties.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Enabling IGMP</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## join-load-balance

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>join-load-balance {<br/>    automatic;<br/>}</pre>                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                              |
| <b>Description</b>              | Enable load balancing of PIM join messages across interfaces and routing devices.                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>automatic</b> —Enables automatic load balancing of PIM join messages. When a new interface or neighbor is introduced into the network, ECMP joins are redistributed with minimal disruption to traffic.                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM Make-Before-Break Join Load Balancing</i></li><li>• <i>Configuring PIM Join Load Balancing</i></li><li>• <i>clear pim join-distribution</i> in the <a href="#">CLI Explorer</a></li></ul>                                                                                           |

## local

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> local {   disable;   address address;   family (inet   inet6) {     disable;     address address;     anycast-pim {       local-address address;       rp-set {         address address &lt;forward-msdp-sa&gt;;       }     }     group-ranges {       destination-ip-prefix&lt;/prefix-length&gt;;     }     hold-time seconds;     override;     priority number;   }   group-ranges {     destination-ip-prefix&lt;/prefix-length&gt;;   }   hold-time seconds;   override;   priority number; } </pre> |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit protocols <a href="#">pim rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>]</p>                                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure the routing device's RP properties.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring Local PIM RPs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                |

## local-address (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>local-address address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <b>pim rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code><br><code>[edit protocols <b>pim rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6) <b>anycast-pim</b>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the routing device local address for the anycast rendezvous point (RP). If this statement is omitted, the router ID is used as this address.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>address</b> —Anycast RP IPv4 or IPv6 address, depending on <b>family</b> configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM Anycast With or Without MSDP</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## mapping-agent-election

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (mapping-agent-election   no-mapping-agent-election);                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp auto-rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp auto-rp</a> ],<br>[edit protocols <a href="#">pim rp auto-rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp auto-rp</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                          |
| <b>Description</b>              | Configure the routing device mapping announcements as a mapping agent.                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>mapping-agent-election</b>—Mapping agents do not announce mappings when receiving mapping messages from a higher-addressed mapping agent.</p> <p><b>no-mapping-agent-election</b>—Mapping agents always announce mappings and do not perform mapping agent election.</p> <p><b>Default:</b> mapping-agent-election</p>                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Auto-RP</i></li> </ul>                                                                                                                                                                                                                                                                                                                  |

## maximum-rps

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-rps <i>limit</i>;</code>                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp embedded-rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp embedded-rp</a> ],<br>[edit protocols <a href="#">pim rp embedded-rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp embedded-rp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                      |
| <b>Description</b>              | Limit the number of RPs that the routing device acknowledges.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>limit</i> —Number of RPs.<br><b>Range:</b> 1 through 500<br><b>Default:</b> 100                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Embedded RP for IPv6</i></li></ul>                                                                                                                                                                                                                                                                                                                       |

## mode (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mode (bidirectional-sparse   bidirectional-sparse-dense   dense   sparse   sparse-dense);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>pim interface</b> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i> ],<br>[edit protocols <b>pim interface</b> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br><b>bidirectional-sparse</b> and <b>bidirectional-sparse-dense</b> options introduced in Junos OS Release 12.1.<br>Statement introduced in Junos OS Release 13.3 for the PTX5000 router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the PIM mode on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p>The choice of PIM mode is closely tied to controlling how groups are mapped to PIM modes, as follows:</p> <ul style="list-style-type: none"> <li>• <b>bidirectional-sparse</b>—Use if all multicast groups are operating in bidirectional, sparse, or SSM mode.</li> <li>• <b>bidirectional-sparse-dense</b>—Use if multicast groups, except those that are specified in the <b>dense-groups</b> statement, are operating in bidirectional, sparse, or SSM mode.</li> <li>• <b>dense</b>—Use if all multicast groups are operating in dense mode.</li> <li>• <b>sparse</b>—Use if all multicast groups are operating in sparse mode or SSM mode.</li> <li>• <b>sparse-dense</b>—Use if multicast groups, except those that are specified in the <b>dense-groups</b> statement, are operating in sparse mode or SSM mode.</li> </ul> <p><b>Default:</b> Sparse mode</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Dense Mode Properties in the Multicast Protocols Feature Guide for Routing Devices</i></li> <li>• <i>Configuring PIM Sparse-Dense Mode Properties in the Multicast Protocols Feature Guide for Routing Devices</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## multicast-router-interface (IGMP Snooping)

---

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | multicast-router-interface;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i> ],<br>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp-snooping</b> vlan (all   <i>vlan-name</i> ) <b>interface</b> (all   <i>interface-name</i> )],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping interface <i>interface-name</i> ] |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Statically configure the interface as an IGMP snooping multicast-router interface—that is, an interface that faces toward a multicast router or other IGMP querier.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <div> <b>NOTE:</b> If the specified interface is a trunk port, the interface becomes a multicast-routing device interface for all VLANs configured on the trunk port. In addition, all unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded to the multicast routing device interface, even if the interface is configured as a multicast routing device interface only for IGMP snooping.</div> <div>Configure an interface as a bridge interface toward other multicast routing devices.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | The interface can either be a host-side or multicast-routing device interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping on EX Series Switches</i></li><li>• <i>Example: Configuring IGMP Snooping</i></li><li>• <i>Configuring IGMP Snooping (CLI Procedure)</i></li><li>• <i>IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview</i></li><li>• <i>host-only-interface</i></li><li>• <a href="#">show igmp-snooping membership on page 3871</a></li></ul>                                                                                                                                                                                                                                                                                |



## neighbor-policy

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>neighbor-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>pim interface interface-name</code> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br><code>pim interface interface-name</code> ],<br>[edit protocols <code>pim interface interface-name</code> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim interface interface-name</code> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Apply a PIM interface-level policy to filter neighbor IP addresses.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <i>policy-name</i> —Name of the policy that filters neighbor IP addresses.                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Interface-Level PIM Neighbor Policies</i></li> </ul>                                                                                                                                                                                                                                                                                                                                       |

## pim

---

```
Syntax  pim {
    disable;
    assert-timeout seconds;
    dense-groups {
        addresses;
    }
    dr-election-on-p2p;
    export;
    family (inet | inet6) {
        disable;
    }
    graceful-restart {
        disable;
        no-bidirectional-mode;
        restart-duration seconds;
    }
    import [ policy-names ];
    interface interface-name {
        family (inet | inet6) {
            disable;
        }
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
            }
            loose-check;
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (0 | 1 | automatic);
        }
        accept-remote-source;
        disable;
        bidirectional {
            df-election {
                backoff-period milliseconds;
                offer-period milliseconds;
                robustness-count number;
            }
        }
        family (inet | inet6) {
            disable;
        }
        hello-interval seconds;
```

```

mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse |
sparse-dense);
neighbor-policy [ policy-names ];
override-interval milliseconds;
priority number;
propagation-delay milliseconds;
reset-tracking-bit;
version version;
}
join-load-balance;
join-prune-timeout;
mdt {
  data-mdt-reuse;
  group-range multicast-prefix;
  threshold {
    group group-address {
      source source-address {
        rate threshold-rate;
      }
    }
    tunnel-limit limit;
  }
}
mvpn {
  autodiscovery {
    inet-mdt;
  }
}
nonstop-routing;
override-interval milliseconds;
propagation-delay milliseconds;
reset-tracking-bit;
rib-group group-name;
rp {
  auto-rp {
    (announce | discovery | mapping);
    (mapping-agent-election | no-mapping-agent-election);
  }
  bidirectional {
    address address {
      group-ranges {
        destination-ip-prefix </prefix-length>;
      }
      hold-time seconds;
      priority number;
    }
  }
  bootstrap {
    family (inet | inet6) {
      export [ policy-names ];
      import [ policy-names ];
      priority number;
    }
  }
  bootstrap-import [ policy-names ];
  bootstrap-export [ policy-names ];
}

```

```
bootstrap-priority number;  
dr-register-policy [ policy-names ];  
embedded-rp {  
  group-ranges {  
    destination-ip-prefix </prefix-length>;  
  }  
  maximum-rps limit;  
}  
group-rp-mapping {  
  family (inet | inet6) {  
    log-interval seconds;  
    maximum limit;  
    threshold value;  
  }  
}  
log-interval seconds;  
maximum limit;  
threshold value;  
}  
local {  
  family (inet | inet6) {  
    address address;  
    anycast-pim {  
      rp-set {  
        address address <forward-msdp-sa>;  
      }  
      disable;  
      local-address address;  
    }  
    group-ranges {  
      destination-ip-prefix </prefix-length>;  
    }  
    hold-time seconds;  
    override;  
    priority number;  
  }  
}  
register-limit {  
  family (inet | inet6) {  
    log-interval seconds;  
    maximum limit;  
    threshold value;  
  }  
}  
log-interval seconds;  
maximum limit;  
threshold value;  
}  
rp-register-policy [ policy-names ];  
spt-threshold {  
  infinity [ policy-names ];  
}  
static {  
  address address {
```

```

        override;
        version version;
        group-ranges {
            destination-ip-prefix </prefix-length>;
        }
    }
}
rpf-selection {
    group group-address {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
    prefix-list prefix-list-addresses {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
}
sglimit {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
    log-interval seconds;
    maximum limit;
    threshold value;
}
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
tunnel-devices [ mt-fpc/pic/port ];
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],  
 [edit protocols],  
 [edit routing-instances *routing-instance-name* protocols]

**Release Information** Statement introduced before Junos OS Release 7.4.  
**family** statement introduced in Junos OS Release 9.6.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.

|                                 |                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>              | Enable PIM on the routing device.<br><br>The remaining statements are explained separately.                                                                                                                                                                              |
| <b>Default</b>                  | PIM is disabled on the routing device.                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode</i></li><li>• <i>Configuring PIM Dense Mode Properties</i></li><li>• <i>Configuring PIM Sparse-Dense Mode Properties</i></li></ul> |

---

## priority (PIM Interfaces)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority number;</code>                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>pim interface interface-name</code> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim interface interface-name</code> ],<br>[edit protocols <code>pim interface interface-name</code> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim interface interface-name</code> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the routing device's likelihood to be elected as the designated router.                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>number</b> —Routing device's priority for becoming the designated router. A higher value corresponds to a higher priority.<br><b>Range:</b> 0 through 4294967295<br><b>Default:</b> 1 (Each routing device has an equal probability of becoming the DR.)                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Interface Priority for PIM Designated Router Selection</i></li></ul>                                                                                                                                                                                                                                                                                                                     |

## priority (Bootstrap)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp bootstrap</a> (inet   inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap</a> (inet   inet6)],</p> <p>[edit protocols <a href="#">pim rp bootstrap</a> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap</a> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the routing device's likelihood to be elected as the bootstrap router.                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b><i>number</i></b>—Routing device's priority for becoming the bootstrap router. A higher value corresponds to a higher priority.</p> <p><b>Range:</b> 0 through a 32-bit number</p> <p><b>Default:</b> 0 (The routing device has the least likelihood of becoming the bootstrap router and sends packets with a priority of 0.)</p>                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> <li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li> <li>• <a href="#">bootstrap-priority on page 3753</a></li> </ul>                                                                                                                                                                                                                                            |

## priority (PIM RPs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim</code><br><code>rp bidirectional address <i>address</i>],</code><br><code>[edit protocols pim rp bidirectional address <i>address</i>],</code><br><code>[edit protocols <b>pim rp local family</b> (inet   inet6)],</code><br><code>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6)]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.<br>Statement introduced in Junos OS Release 13.3 for the PTX5000 router.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | For PIM-SM, configure this routing device's priority for becoming an RP.<br><br>For bidirectional PIM, configure this RP address' priority for becoming an RP.<br><br>The bootstrap router uses this field when selecting the list of candidate rendezvous points to send in the bootstrap message. A smaller number increases the likelihood that the routing device or RP address becomes the RP. A priority value of 0 means that bootstrap router can override the group range being advertised by the candidate RP.                                                                                                                                                                                       |
| <b>Options</b>                  | <b><i>number</i></b> —Priority for becoming an RP. A lower value corresponds to a higher priority.<br><b>Range:</b> 0 through 255<br><b>Default:</b> 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Local PIM RPs in the Multicast Protocols Feature Guide for Routing Devices</i></li><li>• <i>Example: Configuring Bidirectional PIM</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



## query-interval (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | query-interval <i>seconds</i> ;                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp],<br>[edit protocols igmp]                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                 |
| <b>Description</b>              | Specify how often the querier routing device sends general host-query messages.                                                                                                                                                                                                |
| <b>Options</b>                  | <i>seconds</i> —Time interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 125 seconds                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Modifying the IGMP Host-Query Message Interval</i></li> <li>• <a href="#">query-last-member-interval (Protocols IGMP) on page 3794</a></li> <li>• <a href="#">query-response-interval (Protocols IGMP) on page 3795</a></li> </ul> |

## query-last-member-interval (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | query-last-member-interval <i>seconds</i> ;                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp],<br>[edit protocols igmp]                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                |
| <b>Description</b>              | Specify how often the querier routing device sends group-specific query messages.                                                                                                                                                                             |
| <b>Options</b>                  | <b>seconds</b> —Time interval, in fractions of a second or seconds.<br><b>Range:</b> 0.1 through 0.9, then in 1-second intervals 1 through 999999<br><b>Default:</b> 1 second                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Modifying the IGMP Last-Member Query Interval</i></li><li>• <a href="#">query-interval (Protocols IGMP) on page 3793</a></li><li>• <a href="#">query-response-interval (Protocols IGMP) on page 3795</a></li></ul> |

## query-response-interval (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-response-interval <i>seconds</i>;</code>                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp],<br>[edit protocols igmp]                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                    |
| <b>Description</b>              | Specify how long the querier routing device waits to receive a response to a host-query message from a host.                                                                                                                                                      |
| <b>Options</b>                  | <b><i>seconds</i></b> —The query response interval must be less than the query interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 10 seconds                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Modifying the IGMP Query Response Interval</i></li> <li>• <a href="#">query-interval (Protocols IGMP) on page 3793</a></li> <li>• <a href="#">query-last-member-interval (Protocols IGMP) on page 3794</a></li> </ul> |

## receiver

---

|                                 |                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>receiver {<br/>    source-vlans <i>vlan-list</i>;<br/>    install;<br/>}</pre>                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding]                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                          |
| <b>Description</b>              | Configure a VLAN as a multicast receiver VLAN of the multicast VLAN (MVLAN).<br><br>The remaining statements are explained separately.                                             |
| <b>Default</b>                  | Disabled                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast VLAN Registration</i></li><li>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i></li></ul> |

## restart-duration (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>restart-duration <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim graceful-restart</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim graceful-restart</a> ],<br>[edit protocols <a href="#">pim graceful-restart</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim graceful-restart</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the duration of the graceful restart interval.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b><i>seconds</i></b> —Time that the routing device waits (in seconds) to complete PIM sparse mode graceful restart.<br><b>Range:</b> 30 through 300<br><b>Default:</b> 60                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring PIM Sparse Mode Graceful Restart</i></li> </ul>                                                                                                                                                                                                                                                                                                                       |

## rib-group (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rib-group {<br/>    inet <i>group-name</i>;<br/>    inet6 <i>group-name</i>;<br/>}</pre>                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Associate a routing table group with PIM.                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b><i>table-name</i></b> —Name of the routing table. The name must be one that you defined with the <b>rib-groups</b> statement at the <b>[edit routing-options]</b> hierarchy level.                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring a Dedicated PIM RPF Routing Table</i></li></ul>                                                                                                                                                                                                                                         |

## robust-count (IGMP Snooping)

|                                 |                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>robust-count <i>number</i>;</code>                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols <b>igmp-snooping</b> <b>vlan</b> (all   <i>vlan-name</i> )]                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                                   |
| <b>Description</b>              | Configure the number of queries the switch sends before removing a multicast group from the multicast forwarding table. We recommend that the robust count be set to the same value on all multicast routers and switches in the VLAN. |
| <b>Default</b>                  | The default is the value of the <b>robust-count</b> statement configured for IGMP. The default for the IGMP <b>robust-count</b> statement is 2.                                                                                        |
| <b>Options</b>                  | <b><i>number</i></b> —Number of queries the switch sends before timing out a multicast group.<br><b>Range:</b> 2 through 10                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring IGMP Snooping (CLI Procedure)</i></li> </ul>                                                                                                                                   |

## robust-count (Protocols IGMP)

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>robust-count <i>number</i>;</code>                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp],<br>[edit protocols igmp]                                                                                                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Tune the expected packet loss on a subnet. This factor is used to calculate the group member interval, other querier present interval, and last-member query count.                            |
| <b>Options</b>                  | <b><i>number</i></b> —Robustness variable.<br><b>Range:</b> 2 through 10<br><b>Default:</b> 2                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Modifying the IGMP Robustness Variable</i></li> </ul>                                                                                              |

## rp

---

```
Syntax  rp {
    auto-rp {
        (announce | discovery | mapping);
        (mapping-agent-election | no-mapping-agent-election);
    }
    bidirectional {
        address address {
            group-ranges {
                destination-ip-prefix </prefix-length>;
            }
            hold-time seconds;
            priority number;
        }
    }
    bootstrap {
        family (inet | inet6) {
            export [ policy-names ];
            import [ policy-names ];
            priority number;
        }
    }
    bootstrap-export [ policy-names ];
    bootstrap-import [ policy-names ];
    bootstrap-priority number;
    dr-register-policy [ policy-names ];
    embedded-rp {
        group-ranges {
            destination-ip-prefix </prefix-length>;
        }
        maximum-rps limit;
    }
    group-rp-mapping {
        family (inet | inet6) {
            log-interval seconds;
            maximum limit;
            threshold value;
        }
    }
    log-interval seconds;
    maximum limit;
    threshold value;
}
local {
    family (inet | inet6) {
        disable;
        address address;
        anycast-pim {
            local-address address;
            address address <forward-msdp-sa>;
            rp-set {
            }
        }
    }
}
```



```

    }
    group-ranges {
        destination-ip-prefix</prefix-length>;
    }
    hold-time seconds;
    override;
    priority number;
}
}
register-limit {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
}
log-interval seconds;
maximum limit;
threshold value;
}
}
register-probe-time register-probe-time;
}
rp-register-policy [ policy-names ];
static {
    address address {
        override;
        version version;
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
    }
}
}
}

```

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a>],</p> <p>[edit protocols <a href="#">pim</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                          |
| <b>Description</b>              | <p>Configure the routing device as an actual or potential RP. A routing device can be an RP for more than one group.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                     |
| <b>Default</b>                  | If you do not include the <b>rp</b> statement, the routing device can never become the RP.                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                         |

**Related Documentation**

- [Understanding PIM Sparse Mode](#)

## **rp-register-policy**

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rp-register-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ],<br>[edit protocols <a href="#">pim rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                          |
| <b>Description</b>              | Apply one or more policies to control incoming PIM register messages.                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Register Message Filters on a PIM RP and DR</a></li><li>• <a href="#">dr-register-policy on page 3758</a></li></ul>                                                                                                                                                                                 |

## rp-set

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rp-set {   address address &lt;forward-msdp-sa&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit protocols <b>pim local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim local family</b> (inet   inet6) <b>anycast-pim</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring PIM Anycast With or Without MSDP</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                               |

## source (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>source <i>ip-address</i> {<br/>    source-count <i>number</i>;<br/>    source-increment <i>increment</i>;<br/>}</pre>                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp <b>interface</b> <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ],<br>[edit protocols igmp <b>interface</b> <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                   |
| <b>Description</b>              | Specify the IP version 4 (IPv4) unicast source address for the multicast group being statically configured on an interface.                                                                                                                                                      |
| <b>Options</b>                  | <i>ip-address</i> —IPv4 unicast address.<br><br>The remaining statements are explained separately.                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling IGMP Static Group Membership</i></li></ul>                                                                                                                                                                                   |

---

## source-vlans

---

|                                 |                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-vlans <i>vlan-list</i>;</code>                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding receiver]                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                             |
| <b>Description</b>              | Specify a list of multicast VLANs (MVLANS) from which this multicast receiver VLAN receives multicast traffic. Either all of these MVLANS must be in proxy mode or none of them can be in proxy mode. |
| <b>Default</b>                  | Disabled                                                                                                                                                                                              |
| <b>Options</b>                  | <i>vlan-list</i> —Names of the MVLANS.                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast VLAN Registration</i></li><li>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i></li></ul>                    |

## spt-threshold

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>spt-threshold {<br/>    infinfinity [ <i>policy-names</i> ];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>pim</code> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim</code> ],<br>[edit protocols <code>pim</code> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim</code> ]                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Set the SPT threshold to infinity for a source-group address pair. Last-hop multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or an SPT rooted at the source. By default, last-hop routing devices transition to a direct SPT to the source. You can configure this routing device to set the SPT transition value to infinity to prevent this transition for any source-group address pair.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring the PIM SPT Threshold Policy</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## ssm-map (Protocols IGMP)

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-map ssm-map-name;</code>                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp <a href="#">interface interface-name</a> ],<br>[edit protocols igmp <a href="#">interface interface-name</a> ]             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Apply an SSM map to an IGMP interface.                                                                                                                                                     |
| <b>Options</b>                  | <i>ssm-map-name</i> —Name of SSM map.                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring SSM Mapping</i></li> </ul>                                                                                                |

## static (IGMP Snooping)

|                                 |                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>static {   group ip-address; }</pre>                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">igmp-snooping vlan</a> (all   <i>vlan-name</i> ) <a href="#">interface</a> (all   <i>interface-name</i> )]                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                       |
| <b>Description</b>              | <p>Statically define multicast groups on an interface.</p> <p>The remaining statement is explained separately.</p>                                                         |
| <b>Default</b>                  | No multicast groups are statically defined.                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring IGMP Snooping (CLI Procedure)</i></li> <li>• <a href="#">show igmp-snooping membership on page 3871</a></li> </ul> |

## static (Protocols PIM)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>static {<br/>  address address {<br/>    group-ranges {<br/>      destination-ip-prefix&lt;/prefix-length&gt;;<br/>    }<br/>    override;<br/>    version version;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                            |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ],<br>[edit protocols <a href="#">pim rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ]                                                                                                                                                     |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                          |
| Description              | <p>Configure static RP addresses. The default static RP address is 224.0.0.0/4. To configure other addresses, include one or more <b>address</b> statements. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.</p> <p>For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                     |
| Related Documentation    | <ul style="list-style-type: none"><li>Configuring the Static PIM RP Address on the Non-RP Routing Device</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                      |



## static (Protocols IGMP)

**Syntax**

```
static {
    group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
            source-count number;
            source-increment increment;
        }
    }
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols igmp **interface** *interface-name*],  
[edit protocols igmp **interface** *interface-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Test multicast forwarding on an interface without a receiver host.

The **static** statement simulates IGMP joins on a routing device statically on an interface without any IGMP hosts. It is supported for both IGMPv2 and IGMPv3 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.



**NOTE:** To prevent joining too many groups accidentally, the **static** statement is not supported with the **interface all** statement.

The remaining statements are explained separately.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- *Enabling IGMP Static Group Membership*

## traceoptions (Protocols PIM)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>         | <p>Configure PIM tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>             | The default PIM trace options are those inherited from the routing protocol's <b>traceoptions</b> statement included at the [edit <b>routing-options</b> ] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the <b>pim-log</b> file.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>PIM Tracing Flags</b></p> <ul style="list-style-type: none"><li>• <b>assert</b>—Assert messages</li><li>• <b>bidirectional-df-election</b>—Bidirectional PIM designated-forwarder (DF) election events</li></ul> |

- **bootstrap**—Bootstrap messages
- **cache**—Packets in the PIM sparse mode routing cache
- **graft**—Graft and graft acknowledgment messages
- **hello**—Hello packets
- **join**—Join messages
- **mt**—Multicast tunnel messages
- **nsr-synchronization**—Nonstop active routing (NSR) synchronization messages
- **packets**—All PIM packets
- **prune**—Prune messages
- **register**—Register and register stop messages
- **rp**—Candidate RP advertisements
- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow users to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 0 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.                          |
|                                 | routing-control and trace-control—To add this statement to the configuration.           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Trace Options</i></li></ul>  |
|                                 | <ul style="list-style-type: none"><li>• <i>Tracing DVMRP Protocol Traffic</i></li></ul> |
|                                 | <ul style="list-style-type: none"><li>• <i>Tracing MSDP Protocol Traffic</i></li></ul>  |
|                                 | <ul style="list-style-type: none"><li>• <i>Configuring PIM Trace Options</i></li></ul>  |

## traceoptions (Protocols IGMP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;<i>flag-modifier</i>&gt; &lt;disable&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols igmp],<br>[edit protocols igmp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>         | <p>Configure IGMP tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>To trace the paths of multicast packets, use the <b>mtrace</b> command.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>             | The default IGMP trace options are those inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the file <b>igmp-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>IGMP Tracing Flags</b></p> <ul style="list-style-type: none"> <li><b>leave</b>—Leave group messages (for IGMP version 2 only).</li> <li><b>mtrace</b>—Mtrace packets. Use the <b>mtrace</b> command to troubleshoot the software.</li> <li><b>packets</b>—All IGMP packets.</li> </ul> |

- **query**—IGMP membership query messages, including general and group-specific queries.
- **report**—Membership report messages.

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow users to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                              |                                                                                        |
|------------------------------|----------------------------------------------------------------------------------------|
| <b>Required Privilege</b>    | routing and trace—To view this statement in the configuration.                         |
| <b>Level</b>                 | routing-control and trace-control—To add this statement to the configuration.          |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Tracing IGMP Protocol Traffic</i></li></ul> |

## traceoptions (IGMP Snooping)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable  <br/>    no-world-readable&gt;;<br/>    flag <i>flag</i> &lt;flag-modifier&gt;;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>     | [edit protocols <a href="#">igmp-snooping</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.1 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>         | Define tracing operations for IGMP snooping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>             | The <b>traceoptions</b> feature is disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files, including the active trace file. When a trace file reaches its maximum size, its contents are archived into a compressed file named <b><i>filename.0</i></b> and the trace file is emptied. When the trace file reaches its maximum size again, the <b><i>filename.0</i></b> archive file is renamed <b><i>filename.1</i></b> and a new <b><i>filename.0</i></b> archive file is created from the contents of the trace file. This process continues until the maximum number of trace files is reached, at which point the system starts overwriting the oldest archive file each time the trace file is archived. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000<br/><b>Default:</b> 10 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—All tracing operations.</li><li>• <b>general</b>—Trace general IGMP snooping protocol events.</li><li>• <b>krt</b>—Trace communication over routing socket.</li><li>• <b>leave</b>—Trace leave group messages (IGMPv2 and IGMPv3 only).</li><li>• <b>nexthop</b>—Trace nexthop-related events.</li><li>• <b>normal</b>—Trace normal IGMP snooping protocol events. If you do not specify this flag, only unusual or abnormal operations are traced.</li><li>• <b>packets</b>—Trace all IGMP packets.</li><li>• <b>policy</b>—Trace policy processing.</li><li>• <b>query</b>—Trace IGMP membership query messages.</li></ul> |



- **report**—Trace membership report messages.
- **route**—Trace routing information.
- **state**—Trace IGMP state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.
- **vlan**—Trace VLAN-related events.

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers per flag:

- **detail**—Provide detailed trace information
- **disable**—Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

**no-stamp**—(Optional) Omit the timestamp at the beginning of each line in the trace file.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**replace**—(Optional) Replace an existing trace file if there is one. If you do not include this option, tracing output is appended to an existing trace file.

**size size** —(Optional) Maximum size of each trace file, in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is zipped and renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum size, you also must specify a maximum number of files with the **files** option.

**Syntax:** *x* to specify bytes, *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10240 through 4294967295 bytes

**Default:** 128 KB

**world-readable**—(Optional) Allow unrestricted file access.

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                                                               |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Configuring IGMP Snooping Tracing Operations (CLI Procedure) on page 3737</a></li> </ul> |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|

## version (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version version;</code>                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp <a href="#">interface interface-name</a> ],<br>[edit protocols igmp <a href="#">interface interface-name</a> ]                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Specify the version of IGMP.                                                                                                                                                                   |
| <b>Options</b>                  | <b>version</b> —IGMP version number.<br><b>Range:</b> 1, 2, or 3<br><b>Default:</b> IGMP version 2                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Changing the IGMP Version</i></li></ul>                                                                                                             |

## version (IGMP Snooping)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version number;</code>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">igmp-snooping vlan</a> (all   <i>vlan-name</i> )]                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Specify the IGMP version for the IGMP general query that the switch sends to hosts when an interface comes up. The configured IGMP version affects only the version of the general queries sent by a switch. It does not affect the version of IGMP messages that the switch can snoop. For example, If the switch is configured for IGMP version 1 (IGMPv1), it can snoop IGMPv2 and IGMPv3 messages. |
| <b>Default</b>                  | If you do not configure the <b>version</b> statement, the default is IGMPv2.                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>version</b> —IGMP version number.<br><b>Range:</b> 1 and 2.                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring IGMP Snooping (CLI Procedure)</i></li><li>• <i>Configuring IGMP Snooping</i></li></ul>                                                                                                                                                                                                                                                          |

## version (PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version <i>version</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim rp static</b> address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim rp static</b> address <i>address</i>],</p> <p>[edit protocols <b>pim interface</b> <i>interface-name</i>],</p> <p>[edit protocols <b>pim rp static</b> address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp static</b> address <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Specify the version of PIM.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>version</b>—PIM version number.</p> <p><b>Range:</b> 1 or 2</p> <p><b>Default:</b> PIMv1 for rendezvous point (RP) mode (at the [edit protocols <b>pim rp static</b> address <i>address</i>] hierarchy level). PIMv2 for interface mode (at the [edit protocols <b>pim interface</b> <i>interface-name</i>] hierarchy level).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Enabling PIM Sparse Mode</i></li> <li>• <i>Configuring PIM Dense Mode Properties</i></li> <li>• <i>Configuring PIM Sparse-Dense Mode Properties</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## vlan (IGMP Snooping)

```
Syntax  vlan (all | vlan-name) {
        data-forwarding {
            source {
                groups group-prefix;
            }
            receiver {
                source-vlans vlan-list;
                install;
            }
        }
        disable;
        immediate-leave;
        interface (all | interface-name) {
            multicast-router-interface;
            static {
                group ip-address;
            }
        }
        proxy {
            source-address ip-address;
        }
        robust-count number;
        version version;
    }
```

Hierarchy Level [edit protocols **igmp-snooping**]

**Release Information** Statement introduced in Junos OS Release 9.1 for EX Series switches.  
Statement updated with enhanced ? (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.

**Description** Configure IGMP snooping parameters for a VLAN.

When the **vlan** configuration statement is used without the **disable** statement, IGMP snooping is enabled on the specified VLAN or on all VLANs.



**NOTE:** You cannot configure IGMP snooping on a secondary VLAN.

**Default** If the **vlan** statement is not included in the configuration, IGMP snooping is disabled.

- Options**
- **all**—All VLANs on the switch
  - ***vlan-name***—Name of a VLAN.



**TIP:** When you configure IGMP snooping parameters using the **vlan all** statement, any VLAN that is not individually configured for IGMP snooping

inherits the `vlan all` configuration. Any VLAN that is individually configured for IGMP snooping, on the other hand, inherits none of its configuration from `vlan all`. Any parameters that are not explicitly defined for the individual VLAN assume their default values, not the values specified in the `vlan all` configuration.

For example, in the following configuration:

```
protocols {
  igmp-snooping {
    vlan all {
      robust-count 8;
    }
    vlan employee {
      interface ge-0/0/8.0 {
        static {
          group 239.0.10.3
        }
      }
    }
  }
}
```

all VLANs, except `employee`, have a robust count of 8. Because `employee` has been individually configured, its robust count value is not determined by the value set under `vlan all`. Instead, its robust count is the default value of 2.

.....

The remaining statements are explained separately.

|                              |                                                                                                                                                                                                                            |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege</b>    | routing—To view this statement in the configuration.                                                                                                                                                                       |
| <b>Level</b>                 | routing-control—To add this statement to the configuration.                                                                                                                                                                |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping on EX Series Switches</i></li> <li>• <i>Configuring IGMP Snooping (CLI Procedure)</i></li> <li>• <i>show igmp-snooping vlans</i></li> </ul> |



# Administration

- [Routine Monitoring on page 3823](#)
- [Operational Commands on page 3828](#)

## Routine Monitoring

---

- [Monitoring IGMP Snooping on page 3823](#)
- [Verifying IGMP Snooping \(CLI Procedure\) on page 3826](#)

## Monitoring IGMP Snooping

**Purpose**



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring feature to view status and information about IGMP snooping configuration on your EX Series switch.

**Action**

To display IGMP snooping details in the J-Web interface, select **Monitor > Switching > IGMP Snooping**.

To display IGMP snooping details in the CLI, enter the following commands:

- `show igmp-snooping route`
- `show igmp-snooping statistics`
- `show igmp-snooping vlans`



**NOTE:** On EX4300 switches, to display IGMP snooping details in the CLI, enter the following commands:

- `show igmp snooping interface`
- `show igmp snooping statistics`
- `show multicast snooping next-hops`
- `show multicast snooping route`

**Meaning** Table 390 on page 3824 summarizes the IGMP snooping details displayed.

**Table 390: Summary of IGMP Snooping Output Fields**

| Field                                                          | Values                                                                                                     |
|----------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| IGMP Snooping Monitor                                          |                                                                                                            |
| VLAN                                                           | The VLAN for which IGMP snooping is enabled.                                                               |
| Interfaces                                                     | Indicates the interfaces configured as switching interfaces that are associated with the multicast router. |
| Groups                                                         | Indicates the number of the multicast groups learned by the VLAN.                                          |
| Learning Domain                                                | Learning domain for snooping.                                                                              |
| <b>NOTE:</b> This option is supported only on EX4300 switches. |                                                                                                            |
| Query Interval                                                 | Frequency (in seconds) with which the router sends membership queries when it is the querier.              |
| <b>NOTE:</b> This option is supported only on EX4300 switches. |                                                                                                            |
| Query Response Interval                                        | Time (in seconds) that the router waits for a response to a general query.                                 |
| <b>NOTE:</b> This option is supported only on EX4300 switches. |                                                                                                            |
| Last Member Query Interval                                     | Time (in seconds) that the router waits for a report in response to a group-specific query.                |
| <b>NOTE:</b> This option is supported only on EX4300 switches. |                                                                                                            |
| Robustness Count                                               | Number of times the router retries a query.                                                                |
| <b>NOTE:</b> This option is supported only on EX4300 switches. |                                                                                                            |
| MRouters                                                       | Specifies the multicast router.                                                                            |
| <b>NOTE:</b> This option is not supported on EX4300 switches.  |                                                                                                            |
| Receivers                                                      | Specifies the multicast receiver.                                                                          |
| <b>NOTE:</b> This option is not supported on EX4300 switches.  |                                                                                                            |
| Interface Details                                              |                                                                                                            |
| Interfaces                                                     | Name of the interface.                                                                                     |
| <b>NOTE:</b> This option is only supported on EX4300 switches. |                                                                                                            |
| State                                                          | Operating state of the interface. Values are Up or Down.                                                   |
| <b>NOTE:</b> This option is only supported on EX4300 switches. |                                                                                                            |



Table 390: Summary of IGMP Snooping Output Fields (*continued*)

| Field                                                                                                                                                         | Values                                                                                                          |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Group count                                                                                                                                                   | Number of groups on the interface.                                                                              |
| NOTE: This option is only supported on EX4300 switches.                                                                                                       |                                                                                                                 |
| Immediate Leave                                                                                                                                               | State of immediate leave. Values are On or Off.                                                                 |
| NOTE: This option is only supported on EX4300 switches.                                                                                                       |                                                                                                                 |
| Router Interface                                                                                                                                              | Router interfaces that are part of this learning domain.                                                        |
| NOTE: This option is only supported on EX4300 switches.                                                                                                       |                                                                                                                 |
| IGMP Route Information                                                                                                                                        |                                                                                                                 |
| VLAN                                                                                                                                                          | The VLAN for which IGMP snooping is enabled.                                                                    |
| NOTE: This option is not supported on EX4300 switches.                                                                                                        |                                                                                                                 |
| Group                                                                                                                                                         | Indicates the multicast groups learned by the VLAN.                                                             |
| NOTE: This option is not supported on EX4300 switches.                                                                                                        |                                                                                                                 |
| Next-Hop                                                                                                                                                      | Specifies the next hop assigned by the switch after performing the route lookup.                                |
| NOTE: This option is not supported on EX4300 switches.                                                                                                        |                                                                                                                 |
| Statistics                                                                                                                                                    |                                                                                                                 |
| Packets per Vlan                                                                                                                                              | Displays the number of packets sent or received, or the number of received errors.                              |
| NOTE: This option is supported only on EX4300 switches.                                                                                                       |                                                                                                                 |
| Global                                                                                                                                                        | Displays the statistics name and statistics count. The statistics names are:                                    |
| NOTE: This option is supported only on EX4300 switches.                                                                                                       |                                                                                                                 |
| <ul style="list-style-type: none"> <li>• Bad Length</li> <li>• Bad Checksum</li> <li>• Bad Receive If</li> <li>• Rx non-local</li> <li>• Timed out</li> </ul> |                                                                                                                 |
| Multicast Snooping Nexthops                                                                                                                                   |                                                                                                                 |
| ID                                                                                                                                                            | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine. |
| NOTE: This option is supported only on EX4300 switches.                                                                                                       |                                                                                                                 |
| Reference Count                                                                                                                                               | Number of cache entries that are using this next hop.                                                           |
| NOTE: This option is supported only on EX4300 switches.                                                                                                       |                                                                                                                 |

Table 390: Summary of IGMP Snooping Output Fields (*continued*)

| Field                                                          | Values                                                      |
|----------------------------------------------------------------|-------------------------------------------------------------|
| Kernel Reference Count                                         | Kernel reference count for the next hop.                    |
| <b>NOTE:</b> This option is supported only on EX4300 switches. |                                                             |
| Downstream Interface                                           | Interface names associated with each multicast next-hop ID. |
| <b>NOTE:</b> This option is supported only on EX4300 switches. |                                                             |

- Related Documentation**
- [show igmp-snooping route on page 3874](#)
  - [show igmp-snooping statistics on page 3876](#)
  - [show igmp-snooping vlans](#)
  - [Configuring IGMP Snooping \(CLI Procedure\)](#)
  - [Example: Configuring IGMP Snooping on EX Series Switches](#)

## Verifying IGMP Snooping (CLI Procedure)

Internet Group Management Protocol (IGMP) snooping constrains the flooding of IPv4 multicast traffic on VLANs on a switch. This topic describes how to verify IGMP snooping operation on the switch.

It covers:

- [Verifying IGMP Snooping Memberships on page 3826](#)
- [Viewing IGMP Snooping Statistics on page 3827](#)
- [Viewing IGMP Snooping Routing Information on page 3828](#)

### Verifying IGMP Snooping Memberships

**Purpose** Determine group memberships, multicast-router interfaces, host IGMP versions, and the current values of timeout counters.

**Action** Enter the following command:

```
user@switch> show igmp-snooping membership detail
VLAN: vlan2 Tag: 2 (Index: 3)
  Router interfaces:
    ge-1/0/0.0 dynamic Uptime: 00:14:24 timeout: 253
  Group: 225.0.0.1
    ge-1/0/17.0 259 Last reporter: 13.0.0.90 Receiver count: 1
    Uptime: 00:00:19 timeout: 259 Flags: <V3-hosts>
    Include source: 10.2.11.5, 10.2.11.12
```

**Meaning** The switch has multicast membership information for one VLAN on the switch, **vlan2**. IGMP snooping might be enabled on other VLANs, but the switch does not have any multicast membership information for them. The following information is provided:

- Information on the multicast-router interfaces for the VLAN—in this case, **ge-1/0/0.0**. The multicast-router interface has been learned by IGMP snooping, as indicated by the dynamic value. The timeout value shows how many seconds from now the interface will be removed from the multicast forwarding table if the switch does not receive IGMP queries or Protocol Independent Multicast (PIM) updates on the interface.
- Information about the group memberships for the VLAN:
  - Currently, the VLAN has membership in only one multicast group, **225.0.0.1**.
  - The host or hosts that have reported membership in the group are on interface **ge-1/0/17.0**. The last host that reported membership in the group has address **13.0.0.90**. The number of hosts belonging to the group on the interface is shown in the Receiver count field, which is displayed only when host tracking is enabled if immediate leave is configured on the VLAN.
  - The Uptime field shows that the multicast group has been active on the interface for 19 seconds. The interface group membership will time out in 259 seconds if no hosts respond to membership queries during this interval. The Flags field shows the lowest version of IGMP used by a host that is currently a member of the group, which in this case is IGMP version 3 (IGMPv3).
  - Because the interface has IGMPv3 hosts on it, the source addresses from which the IGMPv3 hosts want to receive group multicast traffic are shown (addresses **10.2.11.5** and **10.2.11.12**). The timeout value for the interface group membership is derived from the largest timeout value for all sources addresses for the group.

### Viewing IGMP Snooping Statistics

**Purpose** Display IGMP snooping statistics, such as number of IGMP queries, reports, and leaves received and how many of these IGMP messages contained errors.

**Action** Enter the following command:

```
user@switch> show igmp-snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 0
```

| IGMP Type | Received | Transmitted | Recv Errors |
|-----------|----------|-------------|-------------|
| Queries:  | 74295    | 0           | 0           |
| Reports:  | 18148423 | 0           | 16333523    |
| Leaves:   | 0        | 0           | 0           |
| Other:    | 0        | 0           | 0           |

**Meaning** The output shows how many IGMP messages of each type—**Queries**, **Reports**, **Leaves**—the switch received or transmitted on interfaces on which IGMP snooping is enabled. For each message type, it also shows the number of IGMP packets the switch received that had errors—for example, packets that do not conform to the IGMPv1, IGMPv2, or IGMPv3 standards. If the **Recv Errors** count increases, verify that the hosts are compliant with

IGMP standards. If the switch is unable to recognize the IGMP message type for a packet, it counts the packet under **Receive unknown**.

### Viewing IGMP Snooping Routing Information

---

**Purpose** Display the next-hop information maintained in the multicast forwarding table.

**Action** Enter the following command:

```
user@switch> show igmp-snooping route detail
VLAN          Group          Next-hop
v100          224.0.0.0, *    1323
               Interfaces: ge-0/0/0.0
VLAN          Group          Next-hop
v100          226.0.0.1, *    1322
               Interfaces: ge-0/0/0.0, ge-0/0/1.0, ge-0/0/47.0
```

**Meaning** The output shows the next-hop interfaces for a given multicast group on a VLAN. For example, route **226.0.0.1** on **v100** has next-hop interfaces **ge-0/0/0.0**, **ge-0/0/1.0**, and **ge-0/0/47.0**.

**Related Documentation**

- [clear igmp-snooping membership on page 3835](#)
- [clear igmp-snooping statistics on page 3836](#)
- *Example: Configuring IGMP Snooping on EX Series Switches*
- *Configuring IGMP Snooping (CLI Procedure)*

## Operational Commands

---

- [clear igmp membership](#)
- [clear igmp statistics](#)
- [clear igmp-snooping membership](#)
- [clear igmp-snooping statistics](#)
- [clear multicast bandwidth-admission](#)
- [clear multicast scope](#)
- [clear multicast sessions](#)
- [clear multicast statistics](#)
- [clear pim join](#)
- [clear pim register](#)
- [clear pim statistics](#)
- [mtrace](#)
- [mtrace from-source](#)
- [mtrace monitor](#)
- [mtrace to-gateway](#)

- `show igmp group`
- `show igmp interface`
- `show igmp statistics`
- `show igmp-snooping membership`
- `show igmp-snooping route`
- `show igmp-snooping statistics`
- `show multicast flow-map`
- `show multicast interface`
- `show multicast minfo`
- `show multicast next-hops`
- `show multicast pim-to-igmp-proxy`
- `show multicast pim-to-mld-proxy`
- `show multicast route`
- `show multicast rpf`
- `show multicast scope`
- `show multicast sessions`
- `show multicast usage`
- `show pim bootstrap`
- `show pim interfaces`
- `show pim join`
- `show pim neighbors`
- `show pim rps`
- `show pim source`
- `show pim statistics`

## clear igmp membership

---

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                               | <a href="#">Syntax on page 3830</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3830</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Syntax                                       | <code>clear igmp membership</code><br><code>&lt;group address-range&gt;</code><br><code>&lt;interface interface-name&gt;</code><br><code>&lt;logical-system (all   logical-system-name)&gt;</code>                                                                                                                                                                                                                                                                                                                                                                      |
| Syntax (EX Series Switch and the QFX Series) | <code>clear igmp membership</code><br><code>&lt;group address-range&gt;</code><br><code>&lt;interface interface-name&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Release Information                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                |
| Description                                  | Clear Internet Group Management Protocol (IGMP) group members.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Options                                      | <b>none</b> —Clear all IGMP members on all interfaces and for all address ranges.<br><br><b>group address-range</b> —(Optional) Clear all IGMP members that are in a particular address range. An example of a range is <b>224.2/16</b> . If you omit the destination prefix length, the default is <b>/32</b> .<br><br><b>interface interface-name</b> —(Optional) Clear all IGMP group members on an interface.<br><br><b>logical-system (all   logical-system-name)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Related Documentation                        | <ul style="list-style-type: none"><li>• <a href="#">show igmp group on page 3860</a></li><li>• <a href="#">show igmp interface on page 3864</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                               |
| List of Sample Output                        | <a href="#">clear igmp membership on page 3830</a><br><a href="#">clear igmp membership interface on page 3831</a><br><a href="#">clear igmp membership group on page 3832</a>                                                                                                                                                                                                                                                                                                                                                                                          |
| Output Fields                                | See <a href="#">show igmp group</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## Sample Output

### clear igmp membership

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

```
user@host> show igmp group
```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0  | 224.2.127.253   | 10.1.128.1    | 186     |
| so-0/0/0  | 224.2.127.254   | 10.1.128.1    | 186     |
| so-0/0/0  | 239.255.255.255 | 10.1.128.1    | 187     |
| so-0/0/0  | 224.1.127.255   | 10.1.128.1    | 188     |
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |

```

user@host> clear igmp membership
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0

```

```

user@host> show igmp group

```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |

### clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```

user@host> show igmp group

```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0  | 224.2.127.253   | 10.1.128.1    | 210     |
| so-0/0/0  | 239.255.255.255 | 10.1.128.1    | 210     |
| so-0/0/0  | 224.1.127.255   | 10.1.128.1    | 215     |
| so-0/0/0  | 224.2.127.254   | 10.1.128.1    | 216     |
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |

```

user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0

```

```

user@host> show igmp group

```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |

## clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0  | 224.2.127.253   | 10.1.128.1    | 210     |
| so-0/0/0  | 239.255.255.255 | 10.1.128.1    | 210     |
| so-0/0/0  | 224.1.127.255   | 10.1.128.1    | 215     |
| so-0/0/0  | 224.2.127.254   | 10.1.128.1    | 216     |
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |

```
user@host> clear igmp membership group 239.225/16
```

```
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0  
Clearing Group Membership Range 239.225.0.0/16 on so-1/0/0  
Clearing Group Membership Range 239.225.0.0/16 on so-2/0/0
```

```
user@host> show igmp group
```

| Interface | Group           | Last Reported | Timeout |
|-----------|-----------------|---------------|---------|
| so-0/0/0  | 224.1.127.255   | 10.1.128.1    | 231     |
| so-0/0/0  | 224.2.127.254   | 10.1.128.1    | 233     |
| so-0/0/0  | 224.2.127.253   | 10.1.128.1    | 236     |
| local     | 224.0.0.6       | (null)        | 0       |
| local     | 224.0.0.5       | (null)        | 0       |
| local     | 224.2.127.254   | (null)        | 0       |
| local     | 239.255.255.255 | (null)        | 0       |
| local     | 224.0.0.2       | (null)        | 0       |
| local     | 224.0.0.13      | (null)        | 0       |



## clear igmp statistics

|                                    |                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3833</a><br><a href="#">Syntax (EX Series Switches) on page 3833</a>                                                                                                                                                                                                                                      |
| <b>Syntax</b>                      | clear igmp statistics<br><interface <i>interface-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                  |
| <b>Syntax (EX Series Switches)</b> | clear igmp statistics<br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                             |
| <b>Description</b>                 | Clear Internet Group Management Protocol (IGMP) statistics.                                                                                                                                                                                                                                                                          |
| <b>Options</b>                     | <b>none</b> —Clear IGMP statistics on all interfaces.<br><br><b>interface <i>interface-name</i></b> —(Optional) Clear IGMP statistics for the specified interface only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | clear                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li>• <a href="#">show igmp statistics on page 3868</a></li> </ul>                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>       | <a href="#">clear igmp statistics on page 3833</a>                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>               | See <a href="#">show igmp statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                        |

## Sample Output

### clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        8883         459      0
V1 Membership Report    0            0        0
DVMRP                   19784        35476     0
PIM V1                  18310         0        0
Cisco Trace             0            0        0
V2 Membership Report    0            0        0
Group Leave             0            0        0
Mtrace Response         0            0        0

```

|                                     |   |   |   |
|-------------------------------------|---|---|---|
| Mtrace Request                      | 0 | 0 | 0 |
| Domain Wide Report                  | 0 | 0 | 0 |
| V3 Membership Report                | 0 | 0 | 0 |
| Other Unknown types                 |   |   | 0 |
| IGMP v3 unsupported type            |   |   | 0 |
| IGMP v3 source required for SSM     |   |   | 0 |
| IGMP v3 mode not applicable for SSM |   |   | 0 |

## IGMP Global Statistics

|                |      |
|----------------|------|
| Bad Length     | 0    |
| Bad Checksum   | 0    |
| Bad Receive If | 0    |
| Rx non-local   | 1227 |

```
user@host> clear igmp statistics
```

```
user@host> show igmp statistics
```

IGMP packet statistics for all interfaces

| IGMP Message type                   | Received | Sent | Rx errors |
|-------------------------------------|----------|------|-----------|
| Membership Query                    | 0        | 0    | 0         |
| V1 Membership Report                | 0        | 0    | 0         |
| DVMRP                               | 0        | 0    | 0         |
| PIM V1                              | 0        | 0    | 0         |
| Cisco Trace                         | 0        | 0    | 0         |
| V2 Membership Report                | 0        | 0    | 0         |
| Group Leave                         | 0        | 0    | 0         |
| Mtrace Response                     | 0        | 0    | 0         |
| Mtrace Request                      | 0        | 0    | 0         |
| Domain Wide Report                  | 0        | 0    | 0         |
| V3 Membership Report                | 0        | 0    | 0         |
| Other Unknown types                 |          |      | 0         |
| IGMP v3 unsupported type            |          |      | 0         |
| IGMP v3 source required for SSM     |          |      | 0         |
| IGMP v3 mode not applicable for SSM |          |      | 0         |
| IGMP Global Statistics              |          |      |           |
| Bad Length                          | 0        |      |           |
| Bad Checksum                        | 0        |      |           |
| Bad Receive If                      | 0        |      |           |
| Rx non-local                        | 0        |      |           |

---

## clear igmp-snooping membership

---

|                                 |                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear igmp-snooping membership</code><br><code>&lt;vlan <i>vlan-name</i>&gt;</code>                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                 |
| <b>Description</b>              | Clear IGMP snooping dynamic membership information from the multicast forwarding table.                                                                                            |
| <b>Options</b>                  | <b>none</b> —Clear dynamic membership information for all VLANs.<br><br><b>vlan <i>vlan-name</i></b> —(Optional) Clear dynamic membership information for the specified VLAN.      |
| <b>Required Privilege Level</b> | view                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping membership on page 3871</a></li><li>• <a href="#">clear igmp-snooping statistics on page 3836</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">clear igmp-snooping membership on page 3835</a>                                                                                                                        |

### Sample Output

#### clear igmp-snooping membership

```
user@switch> clear igmp-snooping membership vlan employee-vlan
```

## clear igmp-snooping statistics

---

|                                 |                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear igmp-snooping statistics                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                 |
| <b>Description</b>              | Clear IGMP snooping statistics.                                                                                                                                                    |
| <b>Required Privilege Level</b> | view                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping statistics on page 3876</a></li><li>• <a href="#">clear igmp-snooping membership on page 3835</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">clear igmp-snooping statistics on page 3836</a>                                                                                                                        |

### Sample Output

#### clear igmp-snooping statistics

```
user@switch> clear igmp-snooping statistics
```

## clear multicast bandwidth-admission

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>clear multicast bandwidth-admission &lt;group <i>group-address</i>&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;source <i>source-address</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 8.3.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Reapply IP multicast bandwidth admissions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>none</b>—Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.</p> <p><b>group <i>group-address</i></b>—(Optional) Reapply multicast bandwidth admissions for the specified group.</p> <p><b>inet</b>—(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.</p> <p><b>inet6</b>—(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Reapply multicast bandwidth admission settings for the specified instance. If you do not specify an instance, the command applies to the master routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows:</p> <ul style="list-style-type: none"> <li>• If the interface is congested, and it was admitted previously, it is removed.</li> <li>• If the interface was rejected previously, the <b>clear multicast bandwidth-admission</b> command enables the interface to be admitted as long as enough bandwidth exists on the interface.</li> <li>• If you do not specify an interface, issuing the <b>clear multicast bandwidth-admission</b> command readmits any previously rejected interface for the relevant entries as long as enough bandwidth exists on the interface.</li> </ul> <p>To manually reject previously admitted outbound interfaces, you must specify the interface.</p> <p><b>source <i>source-address</i></b>—(Optional) Use with the <b>group</b> option to reapply multicast bandwidth admission settings for the specified (source, group) entry.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Related Documentation** • [show multicast interface on page 3880](#)

**List of Sample Output** [clear multicast bandwidth-admission on page 3838](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear multicast bandwidth-admission](#)

```
user@host> clear multicast bandwidth-admission
```

## clear multicast scope

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3839</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3839</a>                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax</b>                                       | <pre>clear multicast scope &lt;inet   inet6&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear multicast scope &lt;inet   inet6&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 7.6.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                   |
| <b>Description</b>                                  | Clear IP multicast scope statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                      | <p><b>none</b>—(Same as <b>logical-system all</b>) Clear multicast scope statistics.</p> <p><b>inet</b>—(Optional) Clear multicast scope statistics for IPv4 family addresses.</p> <p><b>inet6</b>—(Optional) Clear multicast scope statistics for IPv6 family addresses.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear multicast scope statistics on a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">show multicast scope on page 3902</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>                        | <a href="#">clear multicast scope on page 3839</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## Sample Output

### clear multicast scope

```
user@host> clear multicast scope
```

## clear multicast sessions

---

|                                              |                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                               | <a href="#">Syntax on page 3840</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3840</a>                                                                                                                                                                                                                                                          |
| Syntax                                       | clear multicast sessions<br><logical-system (all   <i>logical-system-name</i> )><br>< <i>regular-expression</i> >                                                                                                                                                                                                                                                         |
| Syntax (EX Series Switch and the QFX Series) | clear multicast sessions<br>< <i>regular-expression</i> >                                                                                                                                                                                                                                                                                                                 |
| Release Information                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                  |
| Description                                  | Clear IP multicast sessions.                                                                                                                                                                                                                                                                                                                                              |
| Options                                      | <b>none</b> —(Same as <b>logical-system all</b> ) Clear multicast sessions.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>regular-expression</i></b> —(Optional) Clear only multicast sessions that contain the specified regular expression. |
| Required Privilege Level                     | clear                                                                                                                                                                                                                                                                                                                                                                     |
| Related Documentation                        | <ul style="list-style-type: none"><li>• <a href="#">show multicast sessions on page 3904</a></li></ul>                                                                                                                                                                                                                                                                    |
| List of Sample Output                        | <a href="#">clear multicast sessions on page 3840</a>                                                                                                                                                                                                                                                                                                                     |
| Output Fields                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                     |

## Sample Output

### clear multicast sessions

```
user@host> clear multicast sessions
```



## clear multicast statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3841</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3841</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                       | clear multicast statistics<br><inet   inet6><br><instance <i>instance-name</i> ><br><interface <i>interface-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | clear multicast statistics<br><inet   inet6><br><instance <i>instance-name</i> ><br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>                                  | Clear IP multicast statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                                      | <b>none</b> —Clear multicast statistics for all supported address families on all interfaces.<br><br><b>inet</b> —(Optional) Clear multicast statistics for IPv4 family addresses.<br><br><b>inet6</b> —(Optional) Clear multicast statistics for IPv6 family addresses.<br><br><b>instance <i>instance-name</i></b> —(Optional) Clear multicast statistics for the specified instance.<br><br><b>interface <i>interface-name</i></b> —(Optional) Clear multicast statistics on a specific interface.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li><a href="#">show multicast statistics</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>                        | <a href="#">clear multicast statistics on page 3841</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Sample Output

### clear multicast statistics

```
user@host> clear multicast statistics
```

## clear pim join

---

**List of Syntax**    [Syntax on page 3842](#)  
                          [Syntax \(EX Series Switch and the QFX Series\) on page 3842](#)

**Syntax**    clear pim join  
              <group-address>  
              <bidirectional | dense | sparse>  
              <exact>  
              <inet | inet6>  
              <instance *instance-name*>  
              <logical-system (all | *logical-system-name*)>  
              <rp *ip-address/prefix* | source *ip-address/prefix*>  
              <sg | star-g>

**Syntax (EX Series Switch and the QFX Series)**    clear pim join  
                                                          <group-address>  
                                                          <dense | sparse>  
                                                          <exact>  
                                                          <inet | inet6>  
                                                          <instance *instance-name*>  
                                                          <rp *ip-address/prefix* | source *ip-address/prefix*>  
                                                          <sg | star-g>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              Command introduced in Junos OS Release 9.0 for EX Series switches.  
                              **inet6** and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.  
                              Command introduced in Junos OS Release 11.3 for the QFX Series.  
                              Multiple new filter options introduced in Junos OS Release 13.2.

**Description**    Clear the Protocol Independent Multicast (PIM) join and prune states.

**Options**    **none**—Clear the PIM join and prune states for all groups, family addresses, and instances.

**group-address**—(Optional) Clear the PIM join and prune states for a group address.

**bidirectional | dense | sparse**—(Optional) Clear PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.

**exact**—(Optional) Clear only the group that exactly matches the specified group address.

**inet | inet6**—(Optional) Clear the PIM entries for IPv4 or IPv6 family addresses, respectively.

**instance *instance-name***—(Optional) Clear the entries for a specific PIM-enabled routing instance.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**rp *ip-address/prefix* | source *ip-address/prefix***—(Optional) Clear the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

**sg | star-g**—(Optional) Clear PIM (S,G) or (\*,G) entries.

**Additional Information** The `clear pim join` command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.

**Required Privilege Level** clear

**Related Documentation**

- [show pim join on page 3915](#)

**List of Sample Output**

- [clear pim join on page 3843](#)
- [clear pim join inet6 on page 3843](#)
- [clear pim join inet6 star-g on page 3843](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

`clear pim join`

```
user@host> clear pim join
Cleared 8 Join/Prune states
```

`clear pim join inet6`

```
user@host> clear pim join inet6
Cleared 4 Join/Prune states
```

`clear pim join inet6 star-g`

```
user@host> clear pim join inet6 star-g
Cleared 1 Join/Prune states
```

## clear pim register

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3844</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3844</a><br><a href="#">Syntax (PTX Series) on page 3844</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                       | <pre>clear pim register &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear pim register &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (PTX Series)</b>                          | <pre>clear pim register &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 7.6.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                  | Clear Protocol Independent Multicast (PIM) register message counters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                      | <p>none—Clear PIM register message counters for all family addresses, instances, and interfaces.</p> <p><b>inet   inet6</b>—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear register message counters for a specific PIM-enabled routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear PIM register message counters for a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>                       | The <b>clear pim register</b> command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Related Documentation** • [show pim statistics on page 3943](#)

**List of Sample Output** [clear pim register on page 3845](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear pim register](#)

```
user@host> clear pim register
```

## clear pim statistics

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3846</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3846</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax</b>                                       | <pre>clear pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>clear pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                                  | Clear Protocol Independent Multicast (PIM) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                                      | <p><b>none</b>—Clear PIM statistics for all family addresses, instances, and interfaces.</p> <p><b>inet   inet6</b>—(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear statistics for a specific PIM-enabled routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear PIM statistics for a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>                       | The <b>clear pim statistics</b> command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"><li>• <a href="#">show pim statistics on page 3943</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>                        | <a href="#">clear pim statistics on page 3847</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>                                | See <a href="#">show pim statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Sample Output

### clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type      Received      Sent  Rx errors
Hello                  0             0       0
Register               0             0       0
Register Stop          0             0       0
Join Prune             0             0       0
Bootstrap              0             0       0
Assert                 0             0       0
Graft                  0             0       0
Graft Ack              0             0       0
Candidate RP           0             0       0
V1 Query               2111          4222       0
V1 Register            0             0       0
V1 Register Stop       0             0       0
V1 Join Prune          14200         13115       0
V1 RP Reachability     0             0       0
V1 Assert              0             0       0
V1 Graft               0             0       0
V1 Graft Ack           0             0       0
PIM statistics summary for all interfaces:
Unknown type           0
V1 Unknown type        0
Unknown Version        0
Neighbor unknown       0
Bad Length             0
Bad Checksum           0
Bad Receive If         0
Rx Intf disabled       2007
Rx V1 Require V2       0
Rx Register not RP     0
RP Filtered Source     0
Unknown Reg Stop       0
Rx Join/Prune no state 1040
Rx Graft/Graft Ack no state 0
...
```

```
user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type      Received      Sent  Rx errors
Hello                  0             0       0
Register               0             0       0
Register Stop          0             0       0
Join Prune             0             0       0
Bootstrap              0             0       0
Assert                 0             0       0
Graft                  0             0       0
Graft Ack              0             0       0
Candidate RP           0             0       0
V1 Query               1             0       0
V1 Register            0             0       0
...
```





## mtrace

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mtrace source</code><br><logical-system <i>logical-system-name</i> ><br><routing-instance <i>routing-instance-name</i> >                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 9.5 for SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.<br>Command introduced in Junos OS Release 12.3 for the PTX Series. |
| <b>Description</b>              | Display trace information about an IP multicast path.                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <i>source</i> —Source hostname or address.<br><br><i>logical-system (logical-system-name)</i> —(Optional) Perform this operation on a logical system.<br><br><i>routing-instance routing-instance-name</i> —(Optional) Trace a particular routing instance.                                                                                                            |
| <b>Additional Information</b>   | The <b>mtrace</b> command for multicast traffic is similar to the <b>traceroute</b> command used for unicast traffic. Unlike <b>traceroute</b> , <b>mtrace</b> traces traffic backwards, from the receiver to the source.                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">mtrace source on page 3851</a>                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 391 on page 3849</a> describes the output fields for the <b>mtrace</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                              |

**Table 391: mtrace Output Fields**

| Field Name                        | Field Description                                             |
|-----------------------------------|---------------------------------------------------------------|
| <b>Mtrace from</b>                | IP address of the receiver.                                   |
| <b>to</b>                         | IP address of the source.                                     |
| <b>via group</b>                  | IP address of the multicast group (if any).                   |
| <b>Querying full reverse path</b> | Indicates the full reverse path query has begun.              |
| <i>number-of-hops</i>             | Number of hops from the source to the named router or switch. |
| <i>router-name</i>                | Name of the router or switch for this hop.                    |
| <i>address</i>                    | Address of the router or switch for this hop.                 |

Table 391: mtrace Output Fields (*continued*)

| Field Name      | Field Description                              |
|-----------------|------------------------------------------------|
| <i>protocol</i> | Protocol used (for example, PIM).              |
| Round trip time | Average round-trip time, in milliseconds (ms). |
| total ttl of    | Time-to-live (TTL) threshold.                  |

## Sample Output

### mtrace source

```
user@host> mtrace 192.1.4.2
Mtrace from 192.1.4.2 to 192.1.1.2 via group 0.0.0.0
Querying full reverse path... * *
 0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-2 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
-3 hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.
```

## mtrace from-source

---

**Syntax** `mtrace from-source source source`  
`<brief | detail>`  
`<extra-hops extra-hops>`  
`<group group>`  
`<interval interval>`  
`<loop>`  
`<max-hops max-hops>`  
`<max-queries max-queries>`  
`<multicast-response | unicast-response>`  
`<no-resolve>`  
`<no-router-alert>`  
`<response response>`  
`<routing-instance routing-instance-name>`  
`<ttl tll>`  
`<wait-time wait-time>`

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Display trace information about an IP multicast path from a source to this router or switch. If you specify a group address with this command, Junos OS returns additional information, such as packet rates and losses.

**Options** **brief | detail**—(Optional) Display the specified level of output.

**extra-hops *extra-hops***—(Optional) Number of hops to take after reaching a nonresponsive router. You can specify a number between **0** and **255**.

**group *group***—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

**interval *interval***—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10** seconds.

**loop**—(Optional) Loop indefinitely, displaying rate and loss statistics.

**max-hops *max-hops***—(Optional) Maximum hops to trace toward the source. The range of values is **0** through **255**. The default value is **32** hops.

**max-queries *max-queries***—(Optional) Maximum number of query attempts for any hop. The range of values is 1 through **32**. The default is **3**.

**multicast-response**—(Optional) Always request the response using multicast.

**no-resolve**—(Optional) Do not attempt to display addresses symbolically.

**no-router-alert**—(Optional) Do not use the router-alert IP option.

**response *response***—(Optional) Send trace response to a host or multicast address.

**routing-instance** *routing-instance-name*—(Optional) Trace a particular routing instance.

**source** *source*—Source hostname or address.

**ttl** *ttl*—(Optional) IP time-to-live (TTL) value. You can specify a number between 0 and 255. Local queries to the multicast group use a value of 1. Otherwise, the default value is 127.

**unicast-response**—(Optional) Always request the response using unicast.

**wait-time** *wait-time*—(Optional) Number of seconds to wait for a response. The default value is 3.

**Required Privilege Level**

view

**List of Sample Output** [mtrace from-source on page 3854](#)

**Output Fields** [Table 392 on page 3853](#) describes the output fields for the **mtrace from-source** command. Output fields are listed in the approximate order in which they appear.

**Table 392: mtrace from-source Output Fields**

| Field Name                        | Field Description                                             |
|-----------------------------------|---------------------------------------------------------------|
| <b>Mtrace from</b>                | IP address of the receiver.                                   |
| <b>to</b>                         | IP address of the source.                                     |
| <b>via group</b>                  | IP address of the multicast group (if any).                   |
| <b>Querying full reverse path</b> | Indicates the full reverse path query has begun.              |
| <b>number-of-hops</b>             | Number of hops from the source to the named router or switch. |
| <b>router-name</b>                | Name of the router or switch for this hop.                    |
| <b>address</b>                    | Address of the router or switch for this hop.                 |
| <b>protocol</b>                   | Protocol used (for example, PIM).                             |
| <b>Round trip time</b>            | Average round-trip time, in milliseconds (ms).                |
| <b>total ttl of</b>               | Time-to-live (TTL) threshold.                                 |
| <b>source</b>                     | Source address.                                               |
| <b>Response Dest</b>              | Response destination address.                                 |
| <b>Overall</b>                    | Average packet rate for all traffic at each hop.              |

Table 392: mtrace from-source Output Fields (*continued*)

| Field Name                                | Field Description                                                                                                |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Packet Statistics for Traffic From</b> | Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop. |
| <b>Receiver</b>                           | IP address receiving the multicast.                                                                              |
| <b>Query source</b>                       | IP address sending the mtrace query.                                                                             |

## Sample Output

### mtrace from-source

```

user@host> mtrace from-source source 192.1.4.2 group 225.1.1.1
Mtrace from 192.1.4.2 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
 0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-2 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
-3 hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.

Waiting to accumulate statistics...Results after 10 seconds:

Source      Response Dest    Overall    Packet Statistics For Traffic From
192.1.4.2   192.1.1.2  Packet    192.1.4.2 To 225.1.1.1
      v      ___/ rtt    2 ms      Rate      Lost/Sent = Pct  Rate
192.1.2.1
192.1.3.2   routerC.lab.mycompany.net
      v      ^      ttl    2          0/0    = --    0 pps
192.1.4.1
192.1.2.2   routerB.lab.mycompany.net
      v      \__  ttl    3          ?/0          0 pps
192.1.1.2   192.1.1.2
Receiver      Query Source

```

## mtrace monitor

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mtrace monitor                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series. |
| <b>Description</b>              | Listen passively for IP multicast responses. To exit the <b>mtrace monitor</b> command, type Ctrl+c.                                                                                     |
| <b>Options</b>                  | <b>none</b> —Trace the master instance.                                                                                                                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">mtrace monitor on page 3856</a>                                                                                                                                              |
| <b>Output Fields</b>            | <a href="#">Table 393 on page 3855</a> describes the output fields for the <b>mtrace monitor</b> command. Output fields are listed in the approximate order in which they appear.        |

**Table 393: mtrace monitor Output Fields**

| Field Name              | Field Description                                             |
|-------------------------|---------------------------------------------------------------|
| <b>Mtrace query at</b>  | Date and time of the query.                                   |
| <b>by</b>               | Address of the host issuing the query.                        |
| <b>resp to</b>          | Response destination.                                         |
| <b>qid</b>              | Query ID number.                                              |
| <b>packet from...to</b> | IP address of the query source and default group destination. |
| <b>from...to</b>        | IP address of the multicast source and the response address.  |
| <b>via group</b>        | IP address of the group to trace.                             |
| <b>mxhop</b>            | Maximum hop setting.                                          |

## Sample Output

### mtrace monitor

```
user@host> mtrace monitor
Mtrace query at Oct 22 13:36:14 by 192.1.3.2, resp to 224.0.1.32, qid 74a5b8
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:17 by 192.1.3.2, resp to 224.0.1.32, qid 1d07ba
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:20 by 192.1.3.2, resp to same, qid 2fea1d
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:30 by 192.1.3.2, resp to same, qid 7c88ad
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)
```



## mtrace to-gateway

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> mtrace to-gateway gateway gateway &lt;brief   detail&gt; &lt;extra-hops extra-hops&gt; &lt;group group&gt; &lt;interface interface-name&gt; &lt;interval interval&gt; &lt;loop&gt; &lt;max-hops max-hops&gt; &lt;max-queries max-queries&gt; &lt;multicast-response   unicast-response&gt; &lt;no-resolve&gt; &lt;no-router-alert&gt; &lt;response response&gt; &lt;routing-instance routing-instance-name&gt; &lt;tll ttl&gt; &lt;unicast-response&gt; &lt;wait-time wait-time&gt; </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>         | Display trace information about a multicast path from this router or switch to a gateway router or switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>             | <p><b>gateway gateway</b>—Send the trace query to a gateway multicast address.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>extra-hops extra-hops</b>—(Optional) Number of hops to take after reaching a nonresponsive router or switch. You can specify a number between <b>0</b> and <b>255</b>.</p> <p><b>group group</b>—(Optional) Group address for which to trace the path. The default group address is <b>0.0.0.0</b>.</p> <p><b>interface interface-name</b>—(Optional) Source address for sending the trace query.</p> <p><b>interval interval</b>—(Optional) Number of seconds to wait before gathering statistics again. The default value is <b>10</b>.</p> <p><b>loop</b>—(Optional) Loop indefinitely, displaying rate and loss statistics.</p> <p><b>max-hops max-hops</b>—(Optional) Maximum hops to trace toward the source. You can specify a number between <b>0</b> and <b>255</b>. The default value is <b>32</b>.</p> <p><b>max-queries max-queries</b>—(Optional) Maximum number of query attempts for any hop. You can specify a number between <b>0</b> and <b>255</b>. The default value is <b>3</b>.</p> <p><b>multicast-response</b>—(Optional) Always request the response using multicast.</p> <p><b>no-resolve</b>—(Optional) Do not attempt to display addresses symbolically.</p> |

**no-router-alert**—(Optional) Do not use the router-alert IP option.

**response *response***—(Optional) Send trace response to a host or multicast address.

**routing-instance *routing-instance-name***—(Optional) Trace a particular routing instance.

**ttl *tll***—(Optional) IP time-to-live value. You can specify a number between 0 and 225. Local queries to the multicast group use TTL 1. Otherwise, the default value is 127.

**unicast-response**—(Optional) Always request the response using unicast.

**wait-time *wait-time***—(Optional) Number of seconds to wait for a response. The default value is 3.

Required Privilege Level

List of Sample Output [mtrace to-gateway on page 3858](#)

Output Fields [Table 394 on page 3858](#) describes the output fields for the **mtrace to-gateway** command. Output fields are listed in the approximate order in which they appear.

**Table 394: mtrace to-gateway Output Fields**

| Field Name                        | Field Description                                             |
|-----------------------------------|---------------------------------------------------------------|
| <b>Mtrace from</b>                | IP address of the receiver.                                   |
| <b>to</b>                         | IP address of the source.                                     |
| <b>via group</b>                  | IP address of the multicast group (if any).                   |
| <b>Querying full reverse path</b> | Indicates the full reverse path query has begun.              |
| <b><i>number-of-hops</i></b>      | Number of hops from the source to the named router or switch. |
| <b><i>router-name</i></b>         | Name of the router or switch for this hop.                    |
| <b><i>address</i></b>             | Address of the router or switch for this hop.                 |
| <b><i>protocol</i></b>            | Protocol used (for example, PIM).                             |
| <b>Round trip time</b>            | Average round-trip time, in milliseconds (ms).                |
| <b>total ttl of</b>               | Time-to-live (TTL) threshold.                                 |

## Sample Output

### mtrace to-gateway

```
user@host> mtrace to-gateway gateway 192.1.3.2 group 225.1.1.1 interface 192.1.1.73 brief
```

```
Mtrace from 192.1.1.73 to 192.1.1.2 via group 225.1.1.1
```

```
Querying full reverse path... * *  
 0 routerA.lab.mycompany.net (192.1.1.2)  
-1 routerA.lab.mycompany.net (192.1.1.2) PIM thresh^ 1  
-2 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1  
-3 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1  
Round trip time 2 ms; total ttl of 3 required.
```

## show igmp group

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3860</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3860</a>                                                                                                                                                                                                                                                                                                          |
| <b>Syntax</b>                                       | <pre>show igmp group &lt;brief   detail&gt; &lt;group-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show igmp group &lt;brief   detail&gt; &lt;group-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                   |
| <b>Description</b>                                  | Display Internet Group Management Protocol (IGMP) group membership information.                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about membership for all IGMP groups.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>group-name</b>—(Optional) Display group membership for the specified IP address only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li><a href="#">clear igmp membership on page 3830</a></li> </ul>                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>                        | <a href="#">show igmp group (Include Mode) on page 3861</a><br><a href="#">show igmp group (Exclude Mode) on page 3862</a><br><a href="#">show igmp group brief on page 3862</a><br><a href="#">show igmp group detail on page 3862</a>                                                                                                                                                                                   |
| <b>Output Fields</b>                                | <p><a href="#">Table 395 on page 3860</a> describes the output fields for the <b>show igmp group</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                 |

**Table 395: show igmp group Output Fields**

| Field Name        | Field Description                                                                                                                                       | Level of Output |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>  | Name of the interface that received the IGMP membership report. A name of <b>local</b> indicates that the local routing device joined the group itself. | All levels      |
| <b>Group</b>      | Group address.                                                                                                                                          | All levels      |
| <b>Group Mode</b> | Mode the SSM group is operating in: <b>Include</b> or <b>Exclude</b> .                                                                                  | All levels      |

Table 395: show igmp group Output Fields (*continued*)

| Field Name       | Field Description                                                                                                                                                                                                         | Level of Output |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Source           | Source address.                                                                                                                                                                                                           | All levels      |
| Source timeout   | Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.      | detail          |
| Last reported by | Address of the host that last reported membership in this group.                                                                                                                                                          | All levels      |
| Timeout          | Time remaining until the group membership is removed.                                                                                                                                                                     | brief none      |
| Group timeout    | Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer. | detail          |
| Type             | Type of group membership: <ul style="list-style-type: none"> <li>• <b>Dynamic</b>—Host reported the membership.</li> <li>• <b>Static</b>—Membership is configured.</li> </ul>                                             | All levels      |

## Sample Output

### show igmp group (Include Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.2
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.3
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.4
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.2
    Group mode: Include
    Source: 10.0.0.4
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:      0 Type: Dynamic

```

```
Group: 224.0.0.22
Source: 0.0.0.0
Last reported by: Local
Timeout: 0 Type: Dynamic
```

### show igmp group (Exclude Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Source: 0.0.0.0
    Last reported by: Local
    Timeout: 0 Type: Dynamic
  Group: 224.0.0.22
    Source: 0.0.0.0
    Last reported by: Local
    Timeout: 0 Type: Dynamic
```

### show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

### show igmp group detail

```
user@host> show igmp group detail
Interface: t1-0/1/0.0
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.2
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.3
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.4
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
  Group: 232.1.1.2
    Group mode: Include
    Source: 10.0.0.4
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout: 0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Group mode: Exclude
```

```
Source: 0.0.0.0
Source timeout: 0
Last reported by: Local
Group timeout:      0 Type: Dynamic
Group: 224.0.0.22
Group mode: Exclude
Source: 0.0.0.0
Source timeout: 0
Last reported by: Local
Group timeout:      0 Type: Dynamic
```

## show igmp interface

|                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                 | <a href="#">Syntax on page 3864</a><br><a href="#">Syntax (EX Series Switches and the QFX Series) on page 3864</a>                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax</b>                                         | <pre>show igmp interface &lt;brief   detail&gt; &lt;interface-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches and the QFX Series)</b> | <pre>show igmp interface &lt;brief   detail&gt; &lt;interface-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                            | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                             |
| <b>Description</b>                                    | Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                        | <p><b>none</b>—Display standard information about all IGMP-enabled interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface-name</b>—(Optional) Display information about the specified IGMP-enabled interface only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                       | view                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                          | <ul style="list-style-type: none"> <li>• <a href="#">clear igmp membership on page 3830</a></li> </ul>                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>                          | <a href="#">show igmp interface on page 3866</a><br><a href="#">show igmp interface brief on page 3867</a><br><a href="#">show igmp interface detail on page 3867</a><br><a href="#">show igmp interface &lt;interface-name&gt; on page 3867</a>                                                                                                                                                                                    |
| <b>Output Fields</b>                                  | <p><a href="#">Table 396 on page 3864</a> describes the output fields for the <b>show igmp interface</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                       |

**Table 396: show igmp interface Output Fields**

| Field Name | Field Description                                                               | Level of Output |
|------------|---------------------------------------------------------------------------------|-----------------|
| Interface  | Name of the interface.                                                          | All levels      |
| Querier    | Address of the routing device that has been elected to send membership queries. | All levels      |



Table 396: show igmp interface Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>State</b>              | State of the interface: <b>Up</b> or <b>Down</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |
| <b>SSM Map Policy</b>     | Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| <b>Timeout</b>            | How long until the IGMP querier is declared to be unreachable, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels      |
| <b>Version</b>            | IGMP version being used on the interface: <b>1</b> , <b>2</b> , or <b>3</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels      |
| <b>Groups</b>             | Number of groups on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |
| <b>Group limit</b>        | Maximum number of groups allowed on the interface. Any joins requested after the limit is reached are rejected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>Group threshold</b>    | Configured threshold at which a warning message is generated.<br><br>This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>Group log-interval</b> | Time (in seconds) between consecutive log messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| <b>Immediate Leave</b>    | State of the immediate leave option: <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface.</li> <li>• <b>Off</b>—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels      |
| <b>Promiscuous Mode</b>   | State of the promiscuous mode option: <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces.</li> <li>• <b>Off</b>—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels      |
| <b>Passive</b>            | State of the passive mode option: <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves.</li> <li>• <b>Off</b>—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves.</li> </ul> <p>The <b>passive</b> statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the <b>on</b> state declaration:</p> <ul style="list-style-type: none"> <li>• <b>send-general-query</b>—The interface sends general queries.</li> <li>• <b>send-group-query</b>—The interface sends group-specific and group-source-specific queries.</li> <li>• <b>allow-receive</b>—The interface receives control traffic.</li> </ul> | All levels      |

Table 396: show igmp interface Output Fields (*continued*)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| OIF map               | Name of the OIF map (if configured) associated with the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels      |
| SSM map               | Name of the source-specific multicast (SSM) map (if configured) used on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels      |
| Configured Parameters | Information configured by the user: <ul style="list-style-type: none"> <li>• <b>IGMP Query Interval</b>—Interval (in seconds) at which this router sends membership queries when it is the querier.</li> <li>• <b>IGMP Query Response Interval</b>—Time (in seconds) that the router waits for a report in response to a general query.</li> <li>• <b>IGMP Last Member Query Interval</b>—Time (in seconds) that the router waits for a report in response to a group-specific query.</li> <li>• <b>IGMP Robustness Count</b>—Number of times the router retries a query.</li> </ul> | All levels      |
| Derived Parameters    | Derived information: <ul style="list-style-type: none"> <li>• <b>IGMP Membership Timeout</b>—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed.</li> <li>• <b>IGMP Other Querier Present Timeout</b>—Time (in seconds) that the router waits for the IGMP querier to send a query.</li> </ul>                                                                                                                                                                                  | All levels      |

## Sample Output

### show igmp interface

```

user@host> show igmp interface
Interface: at-0/3/1.0
  Querier: 10.111.30.1
  State:      Up Timeout:  None Version:  2 Groups:      4
  SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
  Querier: 10.111.10.1
  State:      Up Timeout:  None Version:  2 Groups:      2
  SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
  Querier: 10.111.20.1
  State:      Up Timeout:  None Version:  2 Groups:      4
  SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0

```

### show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 3866](#).

### show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 3866](#).

### show igmp interface <interface-name>

```
user@host# show igmp interface ge-3/2/0.0
Interface: ge-3/2/0.0
Querier: 20.1.1.1
State: Up Timeout:      None Version: 3 Groups:      1
Group limit: 8
Group threshold: 60
Group log-interval: 10
Immediate leave: Off
Promiscuous mode: Off
```

## show igmp statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3868</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3868</a>                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | <pre>show igmp statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show igmp statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                        |
| <b>Description</b>                                  | Display Internet Group Management Protocol (IGMP) statistics.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                      | <p><b>none</b>—Display IGMP statistics for all interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display IGMP statistics about the specified interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">clear igmp statistics on page 3833</a></li> </ul>                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>                        | <a href="#">show igmp statistics on page 3869</a><br><a href="#">show igmp statistics interface on page 3870</a>                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>                                | <p><a href="#">Table 397 on page 3868</a> describes the output fields for the <b>show igmp statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                 |

**Table 397: show igmp statistics Output Fields**

| Field Name             | Field Description                                                                          |
|------------------------|--------------------------------------------------------------------------------------------|
| IGMP packet statistics | Heading for IGMP packet statistics for all interfaces or for the specified interface name. |

Table 397: show igmp statistics Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IGMP Message type      | <p>Summary of IGMP statistics:</p> <ul style="list-style-type: none"> <li>• <b>Membership Query</b>—Number of membership queries sent and received.</li> <li>• <b>V1 Membership Report</b>—Number of version 1 membership reports sent and received.</li> <li>• <b>DVMRP</b>—Number of DVMRP messages sent or received.</li> <li>• <b>PIM V1</b>—Number of PIM version 1 messages sent or received.</li> <li>• <b>Cisco Trace</b>—Number of Cisco trace messages sent or received.</li> <li>• <b>V2 Membership Report</b>—Number of version 2 membership reports sent or received.</li> <li>• <b>Group Leave</b>—Number of group leave messages sent or received.</li> <li>• <b>Mtrace Response</b>—Number of Mtrace response messages sent or received.</li> <li>• <b>Mtrace Request</b>—Number of Mtrace request messages sent or received.</li> <li>• <b>Domain Wide Report</b>—Number of domain-wide reports sent or received.</li> <li>• <b>V3 Membership Report</b>—Number of version 3 membership reports sent or received.</li> <li>• <b>Other Unknown types</b>—Number of unknown message types received.</li> <li>• <b>IGMP v3 unsupported type</b>—Number of messages received with unknown and unsupported IGMP version 3 message types.</li> <li>• <b>IGMP v3 source required for SSM</b>—Number of IGMP version 3 messages received that contained no source.</li> <li>• <b>IGMP v3 mode not applicable for SSM</b>—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).</li> </ul> |
| Received               | Number of messages received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Sent                   | Number of messages sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Rx errors              | Number of received packets that contained errors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| IGMP Global Statistics | <p>Summary of IGMP statistics for all interfaces.</p> <ul style="list-style-type: none"> <li>• <b>Bad Length</b>—Number of messages received with length errors so severe that further classification could not occur.</li> <li>• <b>Bad Checksum</b>—Number of messages received with a bad IP checksum. No further classification was performed.</li> <li>• <b>Bad Receive If</b>—Number of messages received on an interface not enabled for IGMP.</li> <li>• <b>Rx non-local</b>—Number of messages received from senders that are not local.</li> <li>• <b>Timed out</b>—Number of groups that timed out as a result of not receiving an explicit leave message.</li> <li>• <b>Rejected Report</b>—Number of reports dropped because of the IGMP group policy.</li> <li>• <b>Total Interfaces</b>—Number of interfaces configured to support IGMP.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### show igmp statistics

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        8883         459      0
V1 Membership Report    0            0        0

```

|                                     |      |   |   |
|-------------------------------------|------|---|---|
| DVMRP                               | 0    | 0 | 0 |
| PIM V1                              | 0    | 0 | 0 |
| Cisco Trace                         | 0    | 0 | 0 |
| V2 Membership Report                | 0    | 0 | 0 |
| Group Leave                         | 0    | 0 | 0 |
| Mtrace Response                     | 0    | 0 | 0 |
| Mtrace Request                      | 0    | 0 | 0 |
| Domain Wide Report                  | 0    | 0 | 0 |
| V3 Membership Report                | 0    | 0 | 0 |
| Other Unknown types                 |      |   | 0 |
| IGMP v3 unsupported type            |      |   | 0 |
| IGMP v3 source required for SSM     |      |   | 0 |
| IGMP v3 mode not applicable for SSM |      |   | 0 |
| IGMP Global Statistics              |      |   |   |
| Bad Length                          | 0    |   |   |
| Bad Checksum                        | 0    |   |   |
| Bad Receive If                      | 0    |   |   |
| Rx non-local                        | 1227 |   |   |
| Timed out                           | 0    |   |   |
| Rejected Report                     | 0    |   |   |
| Total Interfaces                    | 2    |   |   |

#### show igmp statistics interface

```
user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type      Received      Sent  Rx errors
Membership Query        0           230      0
V1 Membership Report    0           0        0
```

## show igmp-snooping membership

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show igmp-snooping membership &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;vlan (<i>vlan-id</i>   <i>vlan-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display the multicast group membership information maintained by IGMP snooping.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>none</b>—Display the multicast group membership information about all VLANs on which IGMP snooping is enabled.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output. The default is <b>brief</b>.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display the multicast group membership information about the specified interface.</p> <p><b>vlan (<i>vlan-id</i>   <i>vlan-name</i>)</b>—(Optional) Display the multicast group membership for the specified VLAN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show igmp-snooping route on page 3874</a></li> <li>• <a href="#">show igmp-snooping statistics on page 3876</a></li> <li>• <a href="#">show igmp-snooping vlans</a></li> <li>• <a href="#">Verifying IGMP Snooping (CLI Procedure) on page 3826</a></li> <li>• <a href="#">Configuring IGMP Snooping (CLI Procedure)</a></li> </ul>                                                                                                                         |
| <b>List of Sample Output</b>    | <p><a href="#">show igmp-snooping membership on page 3872</a></p> <p><a href="#">show igmp-snooping membership detail on page 3872</a></p> <p><a href="#">show igmp-snooping membership vlan detail on page 3873</a></p>                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | Table 398 on page 3871 lists the output fields for the <b>show igmp-snooping membership</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                     |

Table 398: show igmp-snooping membership Output Fields

| Field Name | Field Description                                          | Level of Output |
|------------|------------------------------------------------------------|-----------------|
| VLAN       | Name of the VLAN.                                          | All             |
| Interfaces | Interfaces that are members of the listed multicast group. | All             |
| Tag        | Numerical identifier of the VLAN.                          | detail          |

Table 398: show igmp-snooping membership Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Router interfaces</b> | <p>List of information about multicast-router interfaces:</p> <ul style="list-style-type: none"> <li>Name of the multicast-router interface.</li> <li><b>static</b> or <b>dynamic</b>—Whether the multicast-router interface is static or dynamic.</li> <li><b>Uptime</b>—For static interfaces, amount of time since the interface was configured as a multicast-router interface or since the interface last flapped. For dynamic interfaces, amount of time since the first query was received on the interface or since the interface last flapped.</li> <li><b>timeout</b>—Seconds remaining before a dynamic multicast-router interface times out.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail</b>   |
| <b>Group</b>             | <p>IP multicast address of the multicast group.</p> <p>The following information is provided for the multicast group:</p> <ul style="list-style-type: none"> <li>Name of the interface belonging to the multicast group.</li> <li><b>Last reporter</b>—Last host to report membership for the multicast group.</li> <li><b>Receiver count</b>—Number of hosts on the interface that are members of the multicast group. This field appears only if <b>immediate-leave</b> is configured on the VLAN.</li> <li><b>Uptime</b>—Length of time (in hours, minutes, and seconds) a multicast group has been active on the interface.</li> <li><b>timeout</b>—Time (in seconds) left until the entry for the multicast group is removed from the multicast group if no membership reports are received on the interface. This counter is reset to its maximum value when a membership report is received.</li> <li><b>Flags</b>—The lowest IGMP version in use by a host that is a member of the group on the interface.<br/>If the flag <b>static</b> is included, the interface has been configured as static member of the multicast group.</li> <li><b>Include source</b>—Multicast source addresses of all IGMPv3 membership reports received for the group on the interface.</li> </ul> | <b>detail</b>   |

## Sample Output

### show igmp-snooping membership

```

user@switch> show igmp-snooping membership
VLAN: vlan24
  224.1.1.1      *
    Interfaces: ge-0/0/0.0
  224.1.1.100   *
    Interfaces: ge-0/0/0.0
  225.1.1.100   *
    Interfaces: ge-0/0/0.0

```

### show igmp-snooping membership detail

```

user@switch> show igmp-snooping membership detail

VLAN: vlan2 Tag: 2 (Index: 3)
Router interfaces:

```



```
ge-1/0/0.0 dynamic Uptime: 00:14:24 timeout: 253
Group: 225.0.0.1
ge-1/0/17.0 259 Last reporter: 13.0.0.90 Receiver count: 1
Uptime: 00:00:19 timeout: 259 Flags: <V3-hosts>
Include source: 10.2.11.5, 10.2.11.12
```

#### show igmp-snooping membership vlan detail

```
user@switch> show igmp-snooping membership vlan vlan700 detail
VLAN: vlan700 Tag: 700 (Index: 52)
Router interfaces:
  ae2.0 dynamic Uptime: 16:53:13 timeout: 245
Group: 230.150.10.1
  ge-0/0/1.0 Last reporter: 100.2.188.201
  Uptime: 17:00:52 timeout: 237 Flags: <V2-hosts>
  ge-0/0/0.0 Last reporter: 100.2.188.202
  Uptime: 17:00:50 timeout: 243 Flags: <V2-hosts>
```

## show igmp-snooping route

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show igmp-snooping route &lt;brief   detail&gt; &lt;ethernet-switching   inet&gt; &lt;vlan (vlan-id   vlan-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Display IGMP snooping route information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>none</b>—Display route information for all VLANs on which IGMP snooping is enabled.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output. The default is <b>brief</b>.</p> <p><b>ethernet-switching</b>—(Optional) Display information on Layer 2 multicast routes. This is the default.</p> <p><b>inet</b>—(Optional) Display information for Layer 3 multicast routes.</p> <p><b>vlan (vlan-id   vlan-name)</b>—(Optional) Display route information for the specified VLAN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show igmp-snooping membership on page 3871</a></li> <li>• <a href="#">show igmp-snooping statistics on page 3876</a></li> <li>• <a href="#">show igmp-snooping vlans</a></li> <li>• <a href="#">Verifying IGMP Snooping (CLI Procedure) on page 3826</a></li> <li>• <a href="#">Configuring IGMP Snooping (CLI Procedure)</a></li> </ul>                                                                                                                     |
| <b>List of Sample Output</b>    | <p><a href="#">show igmp-snooping route vlan v18 on page 3875</a></p> <p><a href="#">show igmp-snooping route detail on page 3875</a></p> <p><a href="#">show igmp-snooping route inet detail on page 3875</a></p>                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | Table 399 on page 3874 lists the output fields for the <b>show igmp-snooping route</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                           |

**Table 399: show igmp-snooping route Output Fields**

| Field Name    | Field Description                                   |
|---------------|-----------------------------------------------------|
| Table         | Routing table ID for virtual routing instances.     |
| Routing Table | Routing table ID for virtual routing instances.     |
| VLAN          | Name of the VLAN on which IGMP snooping is enabled. |
| Group         | Multicast IPv4 group address.                       |

Table 399: show igmp-snooping route Output Fields (*continued*)

| Field Name                     | Field Description                                                                    |
|--------------------------------|--------------------------------------------------------------------------------------|
| <b>Next-hop</b>                | ID associated with the next-hop device.                                              |
| <b>Routing next-hop</b>        | ID associated with the Layer 3 next-hop device.                                      |
| <b>Interface or Interfaces</b> | Name of the interface or interfaces in the VLAN associated with the multicast group. |
| <b>Layer 2 next-hop</b>        | ID associated with the Layer 2 next-hop device.                                      |

## Sample Output

### show igmp-snooping route vlan v18

```

user@switch> show igmp-snooping route vlan v18
VLAN      Group      Next-hop
v1an18    224.0.0.0, *
v1an18    225.20.20.1, *    1539

```

### show igmp-snooping route detail

```

user@switch> show igmp-snooping route detail
VLAN      Group      Next-hop
default   224.0.0.0, *
v1an100    224.0.0.0, *    1332
          Interfaces: ge-1/0/1.0
VLAN      Group      Next-hop
v1an100    226.0.0.1, *    1334
          Interfaces: ge-1/0/1.0, ge-5/0/30.0

```

### show igmp-snooping route inet detail

```

user@switch> show igmp-snooping route inet detail
Routing table: 0
Group: 229.0.0.1, 171.2.60.100
  Routing next-hop: 3448
  vlan.100
  Interface: vlan.100, VLAN: vlan100, Layer 2 next-hop: 3343

```

## show igmp-snooping statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show igmp-snooping statistics</b>                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Display IGMP snooping statistics.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear igmp-snooping statistics on page 3836</a></li> <li>• <a href="#">show igmp-snooping membership on page 3871</a></li> <li>• <a href="#">show igmp-snooping route on page 3874</a></li> <li>• <a href="#">show igmp-snooping vlans</a></li> <li>• <a href="#">Verifying IGMP Snooping (CLI Procedure) on page 3826</a></li> <li>• <a href="#">Configuring IGMP Snooping (CLI Procedure)</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show igmp-snooping statistics on page 3877</a>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | Table 400 on page 3876 lists the output fields for the <b>show igmp-snooping statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                    |

**Table 400: show igmp-snooping statistics Output Fields**

| Field Name        | Field Description                                                                                            |
|-------------------|--------------------------------------------------------------------------------------------------------------|
| Bad length        | IGMP packet has illegal or bad length.                                                                       |
| Bad checksum      | IGMP or IP checksum is incorrect.                                                                            |
| Invalid interface | Packet was received through an invalid interface.                                                            |
| Not local         | Not used—always 0.                                                                                           |
| Receive unknown   | Unknown IGMP type.                                                                                           |
| Timed out         | Not used—always 0.                                                                                           |
| IGMP Type         | Type of IGMP message (Query, Report, Leave, or Other).                                                       |
| Received          | Number of IGMP packets received.                                                                             |
| Transmitted       | Number of IGMP packets transmitted.                                                                          |
| Recv Errors       | Number of packets received that did not conform to the IGMP version 1 (IGMPv1), IGMPv2, or IGMPv3 standards. |

## Sample Output

### show igmp-snooping statistics

```
user@switch> show igmp-snooping statistics
```

```
Bad length: 0 Bad checksum: 0 Invalid interface: 0
```

```
Not local: 0 Receive unknown: 0 Timed out: 0
```

| IGMP Type | Received | Transmitted | Recv Errors |
|-----------|----------|-------------|-------------|
| Queries:  | 74295    | 0           | 0           |
| Reports:  | 18148423 | 0           | 16333523    |
| Leaves:   | 0        | 0           | 0           |
| Other:    | 0        | 0           | 0           |

## show multicast flow-map

|                                                     |                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3878</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3878</a>                                                                                                                                                                                                                             |
| <b>Syntax</b>                                       | <pre>show multicast flow-map &lt;brief   detail&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                          |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast flow-map &lt;brief   detail&gt;</pre>                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 8.2.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                          |
| <b>Description</b>                                  | Display configuration information about IP multicast flow maps.                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display configuration information about IP multicast flow maps on all systems.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>                        | <a href="#">show multicast flow-map on page 3879</a><br><a href="#">show multicast flow-map detail on page 3879</a>                                                                                                                                                                                                                          |
| <b>Output Fields</b>                                | <p><a href="#">Table 401 on page 3878</a> describes the output fields for the <b>show multicast flow-map</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                            |

**Table 401: show multicast flow-map Output Fields**

| Field Name           | Field Description                                         | Levels of Output |
|----------------------|-----------------------------------------------------------|------------------|
| <b>Name</b>          | Name of the flow map.                                     | All levels       |
| <b>Policy</b>        | Name of the policy associated with the flow map.          | All levels       |
| <b>Cache-timeout</b> | Cache timeout value assigned to the flow map.             | All levels       |
| <b>Bandwidth</b>     | Bandwidth setting associated with the flow map.           | All levels       |
| <b>Adaptive</b>      | Whether or not adaptive mode is enabled for the flow map. | none             |
| <b>Flow-map</b>      | Name of the flow map.                                     | <b>detail</b>    |

Table 401: show multicast flow-map Output Fields (*continued*)

| Field Name                | Field Description                                         | Levels of Output |
|---------------------------|-----------------------------------------------------------|------------------|
| <b>Adaptive Bandwidth</b> | Whether or not adaptive mode is enabled for the flow map. | <b>detail</b>    |
| <b>Redundant Sources</b>  | Redundant sources defined for the same destination group. | <b>detail</b>    |

## Sample Output

### show multicast flow-map

```

user@host> show multicast flow-map
Instance: master
Name          Policy          Cache timeout    Bandwidth Adaptive
map2          policy2         never            2000000 no
map1          policy1         60 seconds      2000000 no

```

## Sample Output

### show multicast flow-map detail

```

user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
  Policy:          policy1
  Cache Timeout:   600 seconds
  Bandwidth:       2000000
  Adaptive Bandwidth: yes
  Redundant Sources: 11.11.11.11
  Redundant Sources: 11.11.11.12
  Redundant Sources: 11.11.11.13

```

## show multicast interface

|                                                     |                                                                                                                                                                                                                                       |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3880</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3880</a>                                                                                                                      |
| <b>Syntax</b>                                       | <pre>show multicast interface &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                         |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show multicast interface                                                                                                                                                                                                              |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 8.3.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                   |
| <b>Description</b>                                  | Display bandwidth information about IP multicast interfaces.                                                                                                                                                                          |
| <b>Options</b>                                      | <p><b>none</b>—Display all interfaces that have multicast configured.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>                        | <a href="#">show multicast interface on page 3881</a>                                                                                                                                                                                 |
| <b>Output Fields</b>                                | <a href="#">Table 402 on page 3880</a> describes the output fields for the <b>show multicast interface</b> command. Output fields are listed in the approximate order in which they appear.                                           |

**Table 402: show multicast interface Output Fields**

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b>                        | Name of the multicast interface.                                                                                                                                                                                                                                                                                                                                   |
| <b>Maximum bandwidth (bps)</b>          | Maximum bandwidth setting, in bits per second, for this interface.                                                                                                                                                                                                                                                                                                 |
| <b>Remaining bandwidth (bps)</b>        | Amount of bandwidth, in bits per second, remaining on the interface.                                                                                                                                                                                                                                                                                               |
| <b>Mapped bandwidth deduction (bps)</b> | <p>Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface.</p> <p><b>NOTE:</b> Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.</p> <p>This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |



Table 402: show multicast interface Output Fields (*continued*)

| Field Name                                   | Field Description                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Local bandwidth deduction (bps)</b>       | <p>Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface.</p> <p><b>NOTE:</b> Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.</p> <p>This field does not appear in the output when the no QoS adjustment feature is disabled.</p> |
| <b>Reverse OIF mapping</b>                   | <p>State of the reverse OIF mapping feature (<b>on</b> or <b>off</b>).</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                                                                    |
| <b>Reverse OIF mapping no QoS adjustment</b> | <p>State of the no QoS adjustment feature (<b>on</b> or <b>off</b>) for interfaces that are using reverse OIF mapping.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                    |
| <b>Leave timer</b>                           | <p>Amount of time a mapped interface remains active after the last mapping ends.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                                                          |
| <b>No QoS adjustment</b>                     | <p>State (<b>on</b>) of the no QoS adjustment feature when this feature is enabled.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>                                                                                                                                                                       |

## Sample Output

### show multicast interface

```

user@host> show multicast interface
Interface           Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3            10000000                  0
fe-0/0/3.210        10000000                  -2000000
fe-0/0/3.220        100000000                 100000000
fe-0/0/3.230        20000000                 18000000
fe-0/0/2.200        100000000                 100000000

```

## show multicast mrinfo

|                                 |                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show multicast mrinfo</code><br><code>&lt;host&gt;</code>                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                  |
| <b>Description</b>              | Display configuration information about IP multicast networks, including neighboring multicast router addresses.                                                                                                          |
| <b>Options</b>                  | <b>none</b> —Display configuration information about all multicast networks.<br><br><b>host</b> —(Optional) Display configuration information about a particular host. Replace <i>host</i> with a hostname or IP address. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show multicast mrinfo on page 3883</a>                                                                                                                                                                        |
| <b>Output Fields</b>            | <a href="#">Table 403 on page 3882</a> describes the output fields for the <b>show multicast mrinfo</b> command. Output fields are listed in the approximate order in which they appear.                                  |

Table 403: show multicast mrinfo Output Fields

| Field Name                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>source-address</i>                | Query address, hostname (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <i>ip-address-1—&gt;ip-address-2</i> | Queried router interface address and directly attached neighbor interface address, respectively.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <i>(name or ip-address)</i>          | Name or IP address of neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <i>[metric/threshold/type/flags]</i> | Neighbor's multicast profile: <ul style="list-style-type: none"> <li><b>metric</b>—Always has a value of 1, because <b>mrinfo</b> queries the directly connected interfaces of a device.</li> <li><b>threshold</b>—Multicast threshold time-to-live (TTL). The range of values is 0 through 255.</li> <li><b>type</b>—Multicast connection type: <b>pim</b> or <b>tunnel</b>.</li> <li><b>flags</b>—Flags for this route: <ul style="list-style-type: none"> <li><b>querier</b>—Queried router is the designated router for the neighboring session.</li> <li><b>leaf</b>—Link is a leaf in the multicast network.</li> <li><b>down</b>—Link status indicator.</li> </ul> </li> </ul> |

## Sample Output

show multicast mrinfo

```
user@host> show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
  192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
  0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```

## show multicast next-hops

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3884</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3884</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                       | <pre>show multicast next-hops &lt;brief   detail&gt; &lt;identifier-number&gt; &lt;inet   inet6&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast next-hops &lt;brief   detail&gt; &lt;identifier-number&gt; &lt;inet   inet6&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p><b>detail</b> option display of next-hop ID number introduced in Junos OS Release 11.1 for M Series and T Series routers and EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                                  | Display the entries in the IP multicast next-hop table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about all entries in the multicast next-hop table for all supported address families.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p>When you include the <b>detail</b> option on M Series and T Series routers and EX Series switches, the downstream interface name includes the next-hop ID number in parentheses, in the form <b>fe-0/1/2.0-(1048574)</b> where <b>1048574</b> is the next-hop ID number.</p> <p><b>identifier-number</b>—(Optional) Show a particular next hop by ID number. The range of values is 1 through <b>65,535</b>.</p> <p><b>inet   inet6</b>—(Optional) Display entries for IPv4 or IPv6 family addresses, respectively.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                        | <a href="#">show multicast next-hops on page 3885</a><br><a href="#">show multicast next-hops (Bidirectional PIM on page 3885</a><br><a href="#">show multicast next-hops brief on page 3886</a><br><a href="#">show multicast next-hops detail on page 3886</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Output Fields** Table 404 on page 3885 describes the output fields for the **show multicast next-hops** command. Output fields are listed in the approximate order in which they appear.

**Table 404: show multicast next-hops Output Fields**

| Field Name                     | Field Description                                                                                                                                      |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Family</b>                  | Protocol family (such as <b>INET</b> ).                                                                                                                |
| <b>ID</b>                      | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine.                                        |
| <b>Refcount</b>                | Number of cache entries that are using this next hop.                                                                                                  |
| <b>KRefcount</b>               | Kernel reference count for the next hop.                                                                                                               |
| <b>Downstream interface</b>    | Interface names associated with each multicast next-hop ID.                                                                                            |
| <b>Incoming interface list</b> | List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM. |

## Sample Output

### show multicast next-hops

```
user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefcount  Downstream interface
262142      4          2  so-1/0/0.0
262143      2          1  mt-1/1/0.49152
262148      2          1  mt-1/1/0.32769
```

### show multicast next-hops (Bidirectional PIM)

```
user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefcount  Downstream interface
2097151      8          4  ge-0/0/1.0

Family: INET6
ID      Refcount  KRefcount  Downstream interface
2097157      2          1  ge-0/0/1.0

Family: Incoming interface list
ID      Refcount  KRefcount  Downstream interface
513      5          2  lo0.0
                    ge-0/0/1.0
514      5          2  lo0.0
                    ge-0/0/1.0
                    xe-4/1/0.0
515      3          1  lo0.0
                    ge-0/0/1.0
                    xe-4/1/0.0
544      1          0  lo0.0
                    xe-4/1/0.0
```

### show multicast next-hops brief

The output for the **show multicast next-hops brief** command is identical to that for the **show multicast next-hops** command. For sample output, see [show multicast next-hops on page 3885](#).

### show multicast next-hops detail

```
user@host> show multicast next-hops detail
Family: INET
ID          Refcount KRefCount Downstream interface
1048577      2          1 fe-0/1/2.0-(1048574)
              ge-0/2/3.0-(1048576)
```

## show multicast pim-to-igmp-proxy

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3887</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3887</a>                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax</b>                                       | <pre>show multicast pim-to-igmp-proxy &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast pim-to-igmp-proxy &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 9.6 for EX Series switches.</p> <p><b>instance</b> option introduced in Junos OS Release 10.3.</p> <p><b>instance</b> option introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                    |
| <b>Description</b>                                  | Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                      | <p><b>none</b>—Display configuration information about PIM-to-IGMP message translation for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li><a href="#">Configuring PIM-to-IGMP and PIM-to-MLD Message Translation</a></li> </ul>                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>                        | <a href="#">show multicast pim-to-igmp-proxy on page 3888</a><br><a href="#">show multicast pim-to-igmp-proxy instance on page 3888</a>                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>                                | <p><a href="#">Table 405 on page 3887</a> describes the output fields for the <b>show multicast pim-to-igmp-proxy</b> command. Output fields are listed in the order in which they appear.</p>                                                                                                                                                                                                                                                      |

**Table 405: show multicast pim-to-igmp-proxy Output Fields**

| Field Name         | Field Description                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b>    | Routing instance. Default instance is <b>master</b> (inet.0 routing table).                                                                           |
| <b>Proxy state</b> | State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: <b>enabled</b> or <b>disabled</b> . |

Table 405: show multicast pim-to-igmp-proxy Output Fields (*continued*)

| Field Name            | Field Description                                                                                             |
|-----------------------|---------------------------------------------------------------------------------------------------------------|
| <i>interface-name</i> | Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured. |

## Sample Output

### show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

### show multicast pim-to-igmp-proxy instance

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```



## show multicast pim-to-mld-proxy

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3889</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3889</a>                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax</b>                                       | show multicast pim-to-mld-proxy<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show multicast pim-to-mld-proxy<br><instance <i>instance-name</i> >                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 9.6.<br>Command introduced in Junos OS Release 9.6 for EX Series switches.<br><b>instance</b> option introduced in Junos OS Release 10.3.<br><b>instance</b> option introduced in Junos OS Release 10.3 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                     |
| <b>Description</b>                                  | Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                      | <b>none</b> —Display configuration information about PIM-to-MLD message translation for all routing instances.<br><br><b>instance</b> <i>instance-name</i> —(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance.<br><br><b>logical-system</b> (all   <i>logical-system-name</i> )—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>                        | <a href="#">show multicast pim-to-mld-proxy on page 3890</a><br><a href="#">show multicast pim-to-mld-proxy instance on page 3890</a>                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>                                | <a href="#">Table 406 on page 3889</a> describes the output fields for the <b>show multicast pim-to-mld-proxy</b> command. Output fields are listed in the order in which they appear.                                                                                                                                                                                                                                                        |

**Table 406: show multicast pim-to-mld-proxy Output Fields**

| Field Name            | Field Description                                                                                                                                   |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Proxy state</b>    | State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: <b>enabled</b> or <b>disabled</b> . |
| <i>interface-name</i> | Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured.                                        |

## Sample Output

### show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

### show multicast pim-to-mld-proxy instance

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```

## show multicast route

**List of Syntax**   [Syntax on page 3891](#)  
[Syntax \(EX Series Switch and the QFX Series\) on page 3891](#)

**Syntax**   `show multicast route`  
                   `<brief | detail | extensive | summary>`  
                   `<active | all | inactive>`  
                   `<group group>`  
                   `<inet | inet6>`  
                   `<instance instance name>`  
                   `<logical-system (all | logical-system-name)>`  
                   `<regular-expression>`  
                   `<source-prefix source-prefix>`

**Syntax (EX Series Switch and the QFX Series)**   `show multicast route`  
                   `<brief | detail | extensive | summary>`  
                   `<active | all | inactive>`  
                   `<group group>`  
                   `<inet | inet6>`  
                   `<instance instance name>`  
                   `<regular-expression>`  
                   `<source-prefix source-prefix>`

**Release Information**   Command introduced before Junos OS Release 7.4.  
                                   Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                   inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.  
                                   Command introduced in Junos OS Release 11.3 for the QFX Series.  
                                   Support for bidirectional PIM added in Junos OS Release 12.1.

**Description**   Display the entries in the IP multicast forwarding table. You can display similar information with the **show route table inet.1** command.

**Options**   **none**—Display standard information about all entries in the multicast forwarding table for all routing instances.

**brief | detail | extensive | summary**—(Optional) Display the specified level of output.

**active | all | inactive**—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.

**group group**—(Optional) Display the cache entries for a particular group.

**inet | inet6**—(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.

**instance instance-name**—(Optional) Display entries in the multicast forwarding table for a specific multicast instance.

**logical-system (all | logical-system-name)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**regular-expression**—(Optional) Display information about the multicast forwarding table entries that match a UNIX OS-style regular expression.

**source-prefix source-prefix**—(Optional) Display the cache entries for a particular source prefix.

**Required Privilege Level** view

**Related Documentation** • *Example: Configuring Multicast-Only Fast Reroute in a PIM Domain*

**List of Sample Output** [show multicast route on page 3893](#)  
[show multicast route \(Bidirectional PIM\) on page 3894](#)  
[show multicast route brief on page 3894](#)  
[show multicast route detail on page 3895](#)  
[show multicast route extensive \(Bidirectional PIM\) on page 3895](#)  
[show multicast route extensive \(Multicast-Only Fast Reroute\) on page 3896](#)  
[show multicast route instance <instance-name> on page 3896](#)  
[show multicast route summary on page 3897](#)

**Output Fields** [Table 407 on page 3892](#) describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

**Table 407: show multicast route Output Fields**

| Field Name                    | Field Description                                                                                                                                                                               | Level of Output  |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| family                        | IPv4 address family ( <b>INET</b> ) or IPv6 address family ( <b>INET6</b> ).                                                                                                                    | All levels       |
| Group                         | Group address.<br><br>For any-source multicast routes, for example for bidirectional PIM, the group address includes the prefix length.                                                         | All levels       |
| Source                        | Prefix and length of the source as it is in the multicast forwarding table.                                                                                                                     | All levels       |
| Incoming interface list       | List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.                                          | All levels       |
| Upstream interface            | Name of the interface on which the packet with this source prefix is expected to arrive.                                                                                                        | All levels       |
| Upstream rpf interface list   | When multicast-only fast reroute (MoFRR) is enabled, a PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. | All levels       |
| Downstream interface list     | List of interface names to which the packet with this source prefix is forwarded.                                                                                                               | All levels       |
| Number of outgoing interfaces | Total number of outgoing interfaces for each (S,G) entry.                                                                                                                                       | <b>extensive</b> |

Table 407: show multicast route Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output   |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| Session description                    | Name of the multicast session.                                                                                                                                                                                                                                                                                                                                                                                                        | detail extensive  |
| Statistics                             | Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobits per second packet forwarding statistic queries fails or times out, the statistics field displays <b>Forwarding statistics are not available</b> .<br><br><b>NOTE:</b> On QFX Series switches, this field does not report valid statistics. | detail extensive  |
| Next-hop ID                            | Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the <b>show multicast nexthops</b> command.                                                                                                                                                                                                                                     | detail extensive  |
| Incoming interface list ID             | For bidirectional PIM, incoming interface list identifier.<br><br>Identifiers for interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.                                                                                                                                                                                                      | detail extensive  |
| Upstream protocol                      | The protocol that maintains the active multicast forwarding route for this group or source.<br><br>When the <b>show multicast route extensive</b> command is used with the <b>display-origin-protocol</b> option, the field name is only <b>Protocol</b> and not <b>Upstream Protocol</b> . However, this field also displays the protocol that installed the active route.                                                           | detail extensive  |
| Route type                             | Type of multicast route. Values can be (S,G) or (*G).                                                                                                                                                                                                                                                                                                                                                                                 | summary           |
| Route state                            | Whether the group is <b>Active</b> or <b>Inactive</b> .                                                                                                                                                                                                                                                                                                                                                                               | summary extensive |
| Route count                            | Number of multicast routes.                                                                                                                                                                                                                                                                                                                                                                                                           | summary           |
| Forwarding state                       | Whether the prefix is pruned or forwarding.                                                                                                                                                                                                                                                                                                                                                                                           | extensive         |
| Cache lifetime/timeout                 | Number of seconds until the prefix is removed from the multicast forwarding table. A value of <b>never</b> indicates a permanent forwarding entry. A value of <b>forever</b> indicates routes that do not have keepalive times.                                                                                                                                                                                                       | extensive         |
| Wrong incoming interface notifications | Number of times that the upstream interface was not available.                                                                                                                                                                                                                                                                                                                                                                        | extensive         |
| Uptime                                 | Time since the creation of a multicast route.                                                                                                                                                                                                                                                                                                                                                                                         | extensive         |

## Sample Output

### show multicast route

```
user@host> show multicast route
Family: INET
```

```
Group: 228.0.0.0
Source: 10.255.14.144/32
Upstream interface: local
Downstream interface list:
    so-1/0/0.0

Group: 239.1.1.1
Source: 10.255.14.144/32
Upstream interface: local
Downstream interface list:
    so-1/0/0.0

Group: 239.1.1.1
Source: 10.255.70.15/32
Upstream interface: so-1/0/0.0
Downstream interface list:
    mt-1/1/0.1081344

Family: INET6
```

#### show multicast route (Bidirectional PIM)

```
user@host> show multicast route
Family: INET

Group: 224.1.1.0/24
Source: *
Incoming interface list:
    lo0.0 ge-0/0/1.0
Downstream interface list:
    ge-0/0/1.0

Group: 224.1.3.0/24
Source: *
Incoming interface list:
    lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
    ge-0/0/1.0

Group: 225.1.1.0/24
Source: *
Incoming interface list:
    lo0.0 ge-0/0/1.0
Downstream interface list:
    ge-0/0/1.0

Group: 225.1.3.0/24
Source: *
Incoming interface list:
    lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
    ge-0/0/1.0
Family: INET6
```

#### show multicast route brief

The output for the **show multicast route brief** command is identical to that for the **show multicast route** command. For sample output, see [show multicast route on page 3893](#) or [show multicast route \(Bidirectional PIM\) on page 3894](#).

**show multicast route detail**

```

user@host> show multicast route detail
Family: INET

Group: 228.0.0.0
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Unknown
  Statistics: 8 kbps, 100 pps, 45272 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Administratively Scoped
  Statistics: 0 kbps, 0 pps, 13404 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.70.15/32
  Upstream interface: so-1/0/0.0
  Downstream interface list:
    mt-1/1/0.1081344
  Session description: Administratively Scoped
  Statistics: 46 kbps, 1000 pps, 921077 packets

  Next-hop ID: 262143
  Upstream protocol: PIM

Family: INET6

```

**show multicast route extensive (Bidirectional PIM)**

```

user@host> show multicast route extensive
Family: INET

Group: 224.1.1.0/24
  Source: *
  Incoming interface list:
    lo0.0 ge-0/0/1.0
  Downstream interface list:
    ge-0/0/1.0
  Number of outgoing interfaces: 1
  Session description: NOB Cross media facilities
  Statistics: 0 kbps, 0 pps, 0 packets
  Next-hop ID: 2097153
  Incoming interface list ID: 585
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0

Group: 224.1.3.0/24

```

```
Source: *
Incoming interface list:
  lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
  ge-0/0/1.0
Number of outgoing interfaces: 1
Session description: NOB Cross media facilities
Statistics: 0 kbps, 0 pps, 0 packets
Next-hop ID: 2097153
Incoming interface list ID: 589
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0
```

Family: INET6

### show multicast route extensive (Multicast-Only Fast Reroute)

```
user@host> show multicast route extensive
```

Instance: master Family: INET

```
Group: 225.1.1.1
Source: 10.0.0.1/32
Upstream rpf interface list:
  fe-1/2/13.0 (P) fe-1/2/14.0 (B)
Downstream interface list:
  fe-1/2/15.0
Session description: Unknown
Forwarding statistics are not available
RPF Next-hop ID: 836
Next-hop ID: 1048585
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: 171 seconds
Wrong incoming interface notifications: 0
Uptime: 00:03:09
```

### show multicast route instance <instance-name>

```
user@host> show multicast route instance v1 extensive
```

Instance: v1 Family: INET

```
Group: 224.1.1.1
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3
```

```
Group: 224.1.1.2
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3
```

```
Group: 224.1.1.3
```



```
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  1t-0/3/0.42 1t-0/3/0.46 1t-0/3/0.43
Number of outgoing interfaces: 3
```

```
Instance: v1 Family: INET6
```

#### show multicast route summary

```
user@host>show multicast route summary
Instance: master Family: INET
```

| Route type | Route state | Route count |
|------------|-------------|-------------|
| (S,G)      | Active      | 2           |
| (S,G)      | Inactive    | 3           |

```
Instance: master Family: INET6
```

## show multicast rpf

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3898</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3898</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>                                       | <pre>show multicast rpf &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;prefix&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast rpf &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;prefix&gt; &lt;summary&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>                                  | Display information about multicast reverse-path-forwarding (RPF) calculations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                      | <p><b>none</b>—Display RPF calculation information for all supported address families.</p> <p><b>inet   inet6</b>—(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about multicast RPF calculations for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>prefix</b>—(Optional) Display the RPF calculation information for the specified prefix.</p> <p><b>summary</b>—(Optional) Display a summary of all multicast RPF information.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>                        | <a href="#">show multicast rpf on page 3899</a><br><a href="#">show multicast rpf inet6 on page 3900</a><br><a href="#">show multicast rpf prefix on page 3901</a><br><a href="#">show multicast rpf summary on page 3901</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Output Fields** Table 408 on page 3899 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

**Table 408: show multicast rpf Output Fields**

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b>      | Name of the routing instance. (Displayed when multicast is configured within a routing instance.)                                                                                                                                                                                                                                                                                                             |
| <b>Source prefix</b> | Prefix and length of the source as it exists in the multicast forwarding table.                                                                                                                                                                                                                                                                                                                               |
| <b>Protocol</b>      | How the route was learned.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Interface</b>     | Upstream RPF interface.<br><br><b>NOTE:</b> The displayed interface information does not apply to bidirectional PIM RP addresses. This is because the <b>show multicast rpf</b> command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF interface information, always use the <b>show pim join extensive</b> command when bidirectional PIM is configured. |
| <b>Neighbor</b>      | Upstream RPF neighbor.<br><br><b>NOTE:</b> The displayed neighbor information does not apply to bidirectional PIM. This is because the <b>show multicast rpf</b> command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF neighbor information, always use the <b>show pim join extensive</b> command when bidirectional PIM is configured.                 |

## Sample Output

### show multicast rpf

```

user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
  Protocol: Static

10.255.14.132/32
  Protocol: Direct
  Interface: lo0.0

10.255.245.91/32
  Protocol: IS-IS
  Interface: so-1/1/1.0
  Neighbor: 192.168.195.21

127.0.0.1/32
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0

```

Neighbor: 192.168.14.254

192.168.0.0/16  
Protocol: Static  
Interface: fxp0.0  
Neighbor: 192.168.14.254

192.168.14.0/24  
Protocol: Direct  
Interface: fxp0.0

192.168.14.132/32  
Protocol: Local

192.168.195.20/30  
Protocol: Direct  
Interface: so-1/1/1.0

192.168.195.22/32  
Protocol: Local

192.168.195.36/30  
Protocol: IS-IS  
Interface: so-1/1/1.0  
Neighbor: 192.168.195.21

#### show multicast rpf inet6

```
user@host> show multicast rpf inet6
```

Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128  
Protocol: Direct  
Interface: lo0.0

::10.255.245.91/128  
Protocol: IS-IS  
Interface: so-1/1/1.0  
Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.20/126  
Protocol: Direct  
Interface: so-1/1/1.0

::192.168.195.22/128  
Protocol: Local

::192.168.195.36/126  
Protocol: IS-IS  
Interface: so-1/1/1.0  
Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.76/126  
Protocol: Direct  
Interface: fe-2/2/0.0

::192.168.195.77/128  
Protocol: Local

```
fe80::/64
Protocol: Direct
Interface: so-1/1/1.0

fe80::290:69ff:fe0c:993a/128
Protocol: Local

fe80::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0

ff02::2/128
Protocol: PIM

ff02::d/128
Protocol: PIM
```

#### show multicast rpf prefix

```
user@host> show multicast rpf ff02::/16

Multicast RPF table: inet6.0, 13 entries

ff02::2/128
    Protocol: PIM

ff02::d/128
    Protocol: PIM

...
```

#### show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```

## show multicast scope

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3902</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3902</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | <pre>show multicast scope &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast scope &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                            |
| <b>Description</b>                                  | Display administratively scoped IP multicast information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about administratively scoped multicast information for all supported address families in all routing instances.</p> <p><b>inet   inet6</b>—(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display administratively scoped information for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                        | <a href="#">show multicast scope on page 3903</a><br><a href="#">show multicast scope inet on page 3903</a><br><a href="#">show multicast scope inet6 on page 3903</a>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>                                | <p><a href="#">Table 409 on page 3902</a> describes the output fields for the <b>show multicast scope</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                 |

**Table 409: show multicast scope Output Fields**

| Field Name   | Field Description                                           |
|--------------|-------------------------------------------------------------|
| Scope name   | Name of the multicast scope.                                |
| Group Prefix | Range of multicast groups that are scoped.                  |
| Interface    | Interface that is the boundary of the administrative scope. |

Table 409: show multicast scope Output Fields (*continued*)

| Field Name      | Field Description                 |
|-----------------|-----------------------------------|
| Resolve Rejects | Number of kernel resolve rejects. |

## Sample Output

### show multicast scope

```
user@host> show multicast scope
```

| Scope name | Group Prefix   | Interface  | Resolve Rejects |
|------------|----------------|------------|-----------------|
| 232-net    | 232.232.0.0/16 | fe-0/0/0.1 | 0               |
| local      | 239.255.0.0/16 | fe-0/0/0.1 | 0               |
| local      | ff05::/16      | fe-0/0/0.1 | 0               |
| larry      | ff05::1234/128 | fe-0/0/0.1 | 0               |

### show multicast scope inet

```
user@host> show multicast scope inet
```

| Scope name | Group Prefix   | Interface  | Resolve Rejects |
|------------|----------------|------------|-----------------|
| 232-net    | 232.232.0.0/16 | fe-0/0/0.1 | 0               |
| local      | 239.255.0.0/16 | fe-0/0/0.1 | 0               |

### show multicast scope inet6

```
user@host> show multicast scope inet6
```

| Scope name | Group Prefix   | Interface  | Resolve Rejects |
|------------|----------------|------------|-----------------|
| local      | ff05::/16      | fe-0/0/0.1 | 0               |
| larry      | ff05::1234/128 | fe-0/0/0.1 | 0               |

## show multicast sessions

---

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| List of Syntax                               | <a href="#">Syntax on page 3904</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3904</a>                                                                                                                                                                                                                                                                                                                                                                                        |
| Syntax                                       | <code>show multicast sessions</code><br><code>&lt;brief   detail   extensive&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;<i>regular-expression</i>&gt;</code>                                                                                                                                                                                                                                                                                     |
| Syntax (EX Series Switch and the QFX Series) | <code>show multicast sessions</code><br><code>&lt;brief   detail   extensive&gt;</code><br><code>&lt;<i>regular-expression</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                               |
| Release Information                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                |
| Description                                  | Display information about announced IP multicast sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Options                                      | <b>none</b> —Display standard information about all multicast sessions for all routing instances.<br><br><b>brief   detail   extensive</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>regular-expression</i></b> —(Optional) Display information about announced sessions that match a UNIX-style regular expression. |
| Required Privilege Level                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| List of Sample Output                        | <a href="#">show multicast sessions on page 3905</a><br><a href="#">show multicast sessions regular-expression detail on page 3905</a>                                                                                                                                                                                                                                                                                                                                                                  |
| Output Fields                                | <a href="#">Table 410 on page 3904</a> describes the output fields for the <b>show multicast sessions</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                              |

**Table 410: show multicast sessions Output Fields**

| Field Name          | Field Description                               |
|---------------------|-------------------------------------------------|
| <i>session-name</i> | Name of the known announced multicast sessions. |

---



## Sample Output

### show multicast sessions

```

user@host> show multicast sessions
1-Department of Biological Sciences, LSU
...
Monterey Bay - DockCam
Monterey Bay - JettyCam
Monterey Bay - StandCam
Monterey DockCam
Monterey DockCam / ROV cam
...
NASA TV (MPEG-1)
...
UO Broadcast - NASA Videos - 25 Years of Progress
UO Broadcast - NASA Videos - Journey through the Solar System
UO Broadcast - NASA Videos - Life in the Universe
UO Broadcast - NASA Videos - Nasa and the Airplane
UO Broadcasts OPB's Oregon Story
UO DOD News Clips
UO Medical Management of Biological Casualties (1)
UO Medical Management of Biological Casualties (2)
UO Medical Management of Biological Casualties (3)
...
376 active sessions.

```

### show multicast sessions regular-expression detail

```

user@host> show multicast sessions "NASA TV" detail
SDP Version: 0  Originated by: -@128.223.83.33
Session: NASA TV (MPEG-1)
Description: NASA television in MPEG-1 format, provided by Private University.
Please contact the UO if you have problems with this feed.
Email: Your Name Here <multicast@lists.private.edu>
Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 224.2.231.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 128.223.91.191 live
Attribute: fmtp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106
Connection Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
Attribute: rtpmap:103 L16/22050/2

```

Attribute: rtpmap:104 L16/22050

1 matching sessions.

## show multicast usage

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3907</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3907</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax</b>                                       | <pre>show multicast usage &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast usage &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                  | Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                      | <p><b>none</b>—Display multicast usage information for all supported address families for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>                        | <a href="#">show multicast usage on page 3908</a><br><a href="#">show multicast usage brief on page 3908</a><br><a href="#">show multicast usage instance on page 3908</a><br><a href="#">show multicast usage detail on page 3909</a>                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>                                | <p><a href="#">Table 411 on page 3908</a> describes the output fields for the <b>show multicast usage</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 411: show multicast usage Output Fields

| Field Name      | Field Description                                                                                                                                                                        |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b> | Name of the routing instance. (Displayed when multicast is configured within a routing instance.)                                                                                        |
| <b>Group</b>    | Group address.                                                                                                                                                                           |
| <b>Sources</b>  | Number of sources.                                                                                                                                                                       |
| <b>Packets</b>  | Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays <b>unavailable</b> . |
| <b>Bytes</b>    | Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays <b>unavailable</b> .     |
| <b>Prefix</b>   | IP address.                                                                                                                                                                              |
| <b>/len</b>     | Prefix length.                                                                                                                                                                           |
| <b>Groups</b>   | Number of multicast groups.                                                                                                                                                              |

## Sample Output

### show multicast usage

```

user@host> show multicast usage
Group          Sources  Packets      Bytes
228.0.0.0      1        52847      4439148
239.1.1.1      2        13450      1125530

Prefix         /len  Groups  Packets      Bytes
10.255.14.144  /32   2        66254      5561304
10.255.70.15   /32   1         43        3374...
```

### show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see [show multicast usage on page 3908](#).

### show multicast usage instance

```

user@host> show multicast usage instance VPN-A
Group          Sources  Packets      Bytes
224.2.127.254  1        5538      509496
224.0.1.39     1         13         624
224.0.1.40     1         13         624

Prefix         /len  Groups  Packets      Bytes
192.168.195.34 /32   1        5538      509496
10.255.14.30   /32   1         13         624
```

```
10.255.245.91 /32 1 13 624
...
```

### show multicast usage detail

```
user@host> show multicast usage detail
Group          Sources Packets          Bytes
228.0.0.0      1          53159          4465356
  Source: 10.255.14.144 /32 Packets: 53159 Bytes: 4465356
239.1.1.1      2          13450          1125530
  Source: 10.255.14.144 /32 Packets: 13407 Bytes: 1122156
  Source: 10.255.70.15  /32 Packets: 43 Bytes: 3374
```

```
Prefix          /len Groups Packets          Bytes
10.255.14.144   /32 2          66566          5587512
  Group: 228.0.0.0      Packets: 53159 Bytes: 4465356
  Group: 239.1.1.1      Packets: 13407 Bytes: 1122156
10.255.70.15    /32 1          43             3374
  Group: 239.1.1.1      Packets: 43 Bytes: 3374
```

## show pim bootstrap

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3910</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3910</a>                                                                                                                                                                                                                                                                                   |
| <b>Syntax</b>                                       | <pre>show pim bootstrap &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim bootstrap &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>instance</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                  |
| <b>Description</b>                                  | For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers.                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                                      | <p><b>none</b>—Display PIM bootstrap router information for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>                        | <a href="#">show pim bootstrap on page 3911</a><br><a href="#">show pim bootstrap instance on page 3911</a>                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>                                | <p><a href="#">Table 412 on page 3910</a> describes the output fields for the <b>show pim bootstrap</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                       |

**Table 412: show pim bootstrap Output Fields**

| Field Name           | Field Description                                                            |
|----------------------|------------------------------------------------------------------------------|
| <b>Instance</b>      | Name of the routing instance.                                                |
| <b>BSR</b>           | Bootstrap router.                                                            |
| <b>Pri</b>           | Priority of the routing device as elected to be the bootstrap router.        |
| <b>Local address</b> | Local routing device address.                                                |
| <b>Pri</b>           | Local routing device address priority to be elected as the bootstrap router. |

Table 412: show pim bootstrap Output Fields (*continued*)

| Field Name     | Field Description                                                                                    |
|----------------|------------------------------------------------------------------------------------------------------|
| <b>State</b>   | Local routing device election state: <b>Candidate</b> , <b>Elected</b> , or <b>Ineligible</b> .      |
| <b>Timeout</b> | How long until the local routing device declares the bootstrap router to be unreachable, in seconds. |

## Sample Output

### show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master
```

| BSR                     | Pri | Local address           | Pri | State      | Timeout |
|-------------------------|-----|-------------------------|-----|------------|---------|
| None                    | 0   | 10.255.71.46            | 0   | InEligible | 0       |
| feco:1:1:1:1:0:aff:785c | 34  | feco:1:1:1:1:0:aff:7c12 | 0   | InEligible | 0       |

### show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A
```

| BSR  | Pri | Local address   | Pri | State      | Timeout |
|------|-----|-----------------|-----|------------|---------|
| None | 0   | 192.168.196.105 | 0   | InEligible | 0       |

## show pim interfaces

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3912</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3912</a>                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax</b>                                       | <pre>show pim interfaces &lt;inet   inet6&gt; &lt;instance (instance-name   all)&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim interfaces &lt;inet   inet6&gt; &lt;instance (instance-name   all)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Support for the <b>instance all</b> option added in Junos OS Release 12.1.</p>                                                                        |
| <b>Description</b>                                  | Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                      | <p><b>none</b>—Display interface information for all family addresses for the main instance.</p> <p><b>inet   inet6</b>—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance (instance-name   all)</b>—(Optional) Display information about interfaces for a specific PIM-enabled routing instance or for all routing instances.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>                        | <a href="#">show pim interfaces on page 3913</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>                                | <p><a href="#">Table 413 on page 3912</a> describes the output fields for the <b>show pim interfaces</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                     |

**Table 413: show pim interfaces Output Fields**

| Field Name      | Field Description                                                                          |
|-----------------|--------------------------------------------------------------------------------------------|
| <b>Instance</b> | Name of the routing instance.                                                              |
| <b>Name</b>     | Interface name.                                                                            |
| <b>State</b>    | State of the interface. The state also is displayed in the <b>show interfaces</b> command. |



Table 413: show pim interfaces Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Mode</b>         | <p>PIM mode running on the interface:</p> <ul style="list-style-type: none"> <li>• <b>B</b>—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers.</li> <li>• <b>S</b>—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic.</li> <li>• <b>Dense</b>—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.)</li> <li>• <b>Sparse-Dense</b>—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as <b>dense</b> is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as <b>sparse</b> is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules.</li> </ul> <p>When sparse-dense mode is configured, the output includes both <b>S</b> and <b>D</b>. When bidirectional-sparse mode is configured, the output includes <b>S</b> and <b>B</b>. When bidirectional-sparse-dense mode is configured, the output includes <b>B</b>, <b>S</b>, and <b>D</b>.</p> |
| <b>IP</b>           | Version number of the address family on the interface: <b>4</b> (IPv4) or <b>6</b> (IPv6).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>V</b>            | PIM version running on the interface: 1 or 2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>State</b>        | <p>State of PIM on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—Bidirectional mode is enabled on the interface and on all PIM neighbors.</li> <li>• <b>DR</b>—Designated router.</li> <li>• <b>NotCap</b>—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol.</li> <li>• <b>NotDR</b>—Not the designated router.</li> <li>• <b>P2P</b>—Point to point.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>NbrCnt</b>       | Number of neighbors that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>JoinCnt(sg)</b>  | Number of (s,g) join messages that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>JointCnt(*g)</b> | Number of (*g) join messages that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>DR address</b>   | Address of the designated router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## Sample Output

### show pim interfaces

```

user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,

```

Active = Bidirectional is active, NotCap = Not Bidirectional Capable

| Name           | Stat | Mode | IP | V | State        | NbrCnt | JoinCnt(sg/*g) | DR address |
|----------------|------|------|----|---|--------------|--------|----------------|------------|
| ge-0/3/0.0     | Up   | S    | 4  | 2 | NotDR,NotCap | 1      | 0/0            | 40.0.0.3   |
| ge-0/3/3.50    | Up   | S    | 4  | 2 | DR,NotCap    | 1      | 9901/100       | 50.0.0.2   |
| ge-0/3/3.51    | Up   | S    | 4  | 2 | DR,NotCap    | 1      | 0/0            | 51.0.0.2   |
| pe-1/2/0.32769 | Up   | S    | 4  | 2 | P2P,NotCap   | 0      | 0/0            |            |

## show pim join

**List of Syntax**   [Syntax on page 3915](#)  
[Syntax \(EX Series Switch and the QFX Series\) on page 3915](#)

**Syntax**   `show pim join`  
                   `<brief | detail | extensive | summary>`  
                   `<bidirectional | dense | sparse>`  
                   `<exact>`  
                   `<inet | inet6>`  
                   `<instance instance-name>`  
                   `<logical-system (all | logical-system-name)>`  
                   `<range>`  
                   `<rp ip-address/prefix | source ip-address/prefix>`  
                   `<sg | star-g>`

**Syntax (EX Series Switch and the QFX Series)**   `show pim join`  
                   `<brief | detail | extensive | summary>`  
                   `<dense | sparse>`  
                   `<exact>`  
                   `<inet | inet6>`  
                   `<instance instance-name>`  
                   `<range>`  
                   `<rp ip-address/prefix | source ip-address/prefix>`  
                   `<sg | star-g>`

**Release Information**   Command introduced before Junos OS Release 7.4.  
                               Command introduced in Junos OS Release 9.0 for EX Series switches.  
                               **summary** option introduced in Junos OS Release 9.6.  
                               **inet6** and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.  
                               Support for bidirectional PIM added in Junos OS Release 12.1.  
                               Command introduced in Junos OS Release 11.3 for the QFX Series.  
                               Multiple new filter options introduced in Junos OS Release 13.2.

**Description**   Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.

For bidirectional PIM, display information about PIM group ranges (\*G-range) for each active bidirectional RP group range, in addition to each of the joined (\*G) routes.

**Options**   **none**—Display the standard information about PIM groups for all supported family addresses for all routing instances.

**brief | detail | extensive | summary**—(Optional) Display the specified level of output.

**bidirectional | dense | sparse**—(Optional) Display information about PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.

**exact**—(Optional) Display information about only the group that exactly matches the specified group address.

**inet | inet6**—(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.

**instance *instance-name***—(Optional) Display information about groups for the specified PIM-enabled routing instance only.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**range**—(Optional) Address range of the group, specified as *prefix/prefix-length*.

**rp *ip-address/prefix* | source *ip-address/prefix***—(Optional) Display information about the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

**sg | star-g**—(Optional) Display information about PIM (S,G) or (\*,G) entries.

**Required Privilege Level**

view

**Related Documentation**

- [clear pim join on page 3842](#)
- *Example: Configuring Multicast-Only Fast Reroute in a PIM Domain*
- *Example: Configuring Bidirectional PIM*
- *Example: Configuring PIM State Limits*

**List of Sample Output**

[show pim join summary on page 3920](#)  
[show pim join \(PIM Sparse Mode\) on page 3920](#)  
[show pim join \(Bidirectional PIM\) on page 3920](#)  
[show pim join inet6 on page 3921](#)  
[show pim join inet6 star-g on page 3921](#)  
[show pim join instance <instance-name> on page 3921](#)  
[show pim join detail on page 3922](#)  
[show pim join extensive \(PIM Sparse Mode\) on page 3922](#)  
[show pim join extensive \(Bidirectional PIM\) on page 3923](#)  
[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 3924](#)  
[show pim join instance <instance-name> extensive on page 3925](#)  
[show pim join extensive \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 3925](#)  
[show pim join extensive \(Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 3926](#)  
[show pim join extensive \(Multipoint LDP with Multicast-Only Fast Reroute\) on page 3928](#)

**Output Fields**

[Table 414 on page 3917](#) describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

Table 414: show pim join Output Fields

| Field Name                               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output                            |
|------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------|
| <b>Instance</b>                          | Name of the routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>brief detail extensive summary none</b> |
| <b>Family</b>                            | Name of the address family: <b>inet</b> (IPv4) or <b>inet6</b> (IPv6).                                                                                                                                                                                                                                                                                                                                                                               | <b>brief detail extensive summary none</b> |
| <b>Route type</b>                        | Type of multicast route: (S,G) or (*G).                                                                                                                                                                                                                                                                                                                                                                                                              | <b>summary</b>                             |
| <b>Route count</b>                       | Number of (S,G) routes and number of (*G) routes.                                                                                                                                                                                                                                                                                                                                                                                                    | <b>summary</b>                             |
| <b>R</b>                                 | Rendezvous Point Tree.                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>brief detail extensive none</b>         |
| <b>S</b>                                 | Sparse.                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>brief detail extensive none</b>         |
| <b>W</b>                                 | Wildcard.                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>brief detail extensive none</b>         |
| <b>Group</b>                             | Group address.                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>brief detail extensive none</b>         |
| <b>Bidirectional group prefix length</b> | For bidirectional PIM, length of the IP prefix for RP group ranges.                                                                                                                                                                                                                                                                                                                                                                                  | All levels                                 |
| <b>Source</b>                            | Multicast source: <ul style="list-style-type: none"> <li>• * (wildcard value)</li> <li>• <i>ipv4-address</i></li> <li>• <i>ipv6-address</i></li> </ul>                                                                                                                                                                                                                                                                                               | <b>brief detail extensive none</b>         |
| <b>RP</b>                                | Rendezvous point for the PIM group.                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>brief detail extensive none</b>         |
| <b>Flags</b>                             | PIM flags: <ul style="list-style-type: none"> <li>• <b>bidirectional</b>—Bidirectional mode entry.</li> <li>• <b>dense</b>—Dense mode entry.</li> <li>• <b>rptree</b>—Entry is on the rendezvous point tree.</li> <li>• <b>sparse</b>—Sparse mode entry.</li> <li>• <b>spt</b>—Entry is on the shortest-path tree for the source.</li> <li>• <b>wildcard</b>—Entry is on the shared tree.</li> </ul>                                                 | <b>brief detail extensive none</b>         |
| <b>Upstream interface</b>                | <p>RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*G).</p> <p>For bidirectional PIM, <b>RP Link</b> means that the interface is directly connected to a subnet that contains a phantom RP address.</p> <p>A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.</p> | <b>brief detail extensive none</b>         |

Table 414: show pim join Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output  |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Upstream neighbor</b>        | <p>Information about the upstream neighbor: <b>Direct</b>, <b>Local</b>, <b>Unknown</b>, or a specific IP address.</p> <p>For bidirectional PIM, <b>Direct</b> means that the interface is directly connected to a subnet that contains a phantom RP address.</p> <p>The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>extensive</b> |
| Active upstream interface       | When multicast-only fast reroute (MoFRR) is configured in a PIM domain, the upstream interface for the active path. A PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. Preference is given to two paths that do not converge to the same immediate upstream router. PIM installs appropriate multicast routes with upstream neighbors as RPF next hops with two (primary and backup) interfaces.                                                                                                                                                                                                                                                                                                                                                                                        | <b>extensive</b> |
| Active upstream neighbor        | On the MoFRR primary path, the IP address of the neighbor that is directly connected to the active upstream interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b> |
| MoFRR Backup upstream interface | <p>The MoFRR upstream interface that is used when the primary path fails.</p> <p>When the primary path fails, the backup path is upgraded to primary, and traffic is forwarded accordingly. If there are alternate paths available, a new backup path is calculated and the appropriate multicast route is updated or installed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b> |
| <b>Upstream state</b>           | <p>Information about the upstream interface:</p> <ul style="list-style-type: none"> <li>• <b>Join to RP</b>—Sending a join to the rendezvous point.</li> <li>• <b>Join to Source</b>—Sending a join to the source.</li> <li>• <b>Local RP</b>—Sending neither join messages nor prune messages toward the RP, because this routing device is the rendezvous point.</li> <li>• <b>Local Source</b>—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device.</li> <li>• <b>Prune to RP</b>—Sending a prune to the rendezvous point.</li> <li>• <b>Prune to Source</b>—Sending a prune to the source.</li> </ul> <p><b>NOTE:</b> RP group range entries have <b>None</b> in the <b>Upstream state</b> field because RP group ranges do not trigger actual PIM join messages between routing devices.</p> | <b>extensive</b> |

Table 414: show pim join Output Fields (*continued*)

| Field Name                                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output  |
|-------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Downstream neighbors</b>               | <p>Information about downstream interfaces:</p> <ul style="list-style-type: none"> <li>• <b>Interface</b>—Interface name for the downstream neighbor.<br/>A pseudo PIM-SM interface appears for all IGMP-only interfaces.<br/>A pseudo multipoint LDP (M-LDP) interface appears on ingress root nodes in M-LDP point-to-multipoint LSPs with inband signaling.</li> <li>• <b>Interface address</b>—Address of the downstream neighbor.</li> <li>• <b>State</b>—Information about the downstream neighbor: <b>join</b> or <b>prune</b>.</li> <li>• <b>Flags</b>—PIM join flags: <b>R (RPtree)</b>, <b>S (Sparse)</b>, <b>W (Wildcard)</b>, or <b>zero</b>.</li> <li>• <b>Uptime</b>—Time since the downstream interface joined the group.</li> <li>• <b>Time since last Join</b>—Time since the last join message was received from the downstream interface.</li> <li>• <b>Time since last Prune</b>—Time since the last prune message was received from the downstream interface.</li> </ul> | <b>extensive</b> |
| <b>Number of downstream interfaces</b>    | Total number of outgoing interfaces for each (S,G) entry.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>extensive</b> |
| <b>Assert Timeout</b>                     | Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>extensive</b> |
| <b>Keepalive timeout</b>                  | Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, <b>Keepalive timeout</b> is <b>Infinity</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>extensive</b> |
| <b>Uptime</b>                             | Time since the creation of (S,G) or (*,G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*,G) state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>extensive</b> |
| <b>Bidirectional accepting interfaces</b> | <p>Interfaces on the routing device that forward bidirectional PIM traffic.</p> <p>The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (<b>DF Winner</b>), or the interface is the reverse path forwarding (RPF) interface toward the RP (<b>RPF</b>).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b> |

## Sample Output

### show pim join summary

```
user@host> show pim join summary
Instance: PIM.master Family: INET

Route type          Route count
(s,g)               2
(*,g)               1

Instance: PIM.master Family: INET6
```

### show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

### show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
Bidirectional group prefix length: 24
Source: *
```



```

RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

### show pim join inet6

```

user@host> show pim join inet6
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
Source: *
RP: ::46.0.0.13
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: ff04::e000:101
Source: ::1.1.1.1
Flags: sparse
Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
Source: ::1.1.1.1
Flags: sparse
Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
Source: ::1.1.1.2
Flags: sparse
Upstream interface: unknown (no neighbor)

```

### show pim join inet6 star-g

```

user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
Source: *
RP: ::46.0.0.13
Flags: sparse,rptree,wildcard
Upstream interface: Local

```

### show pim join instance <instance-name>

```

user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
Source: *
RP: 10.10.47.100

```

Flags: sparse,rptree,wildcard  
Upstream interface: Local

Group: 235.1.1.2  
Source: 192.168.195.74  
Flags: sparse,spt  
Upstream interface: at-0/3/1.0

Group: 235.1.1.2  
Source: 192.168.195.169  
Flags: sparse  
Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6  
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

#### show pim join detail

user@host> show pim join detail  
Instance: PIM.master Family: INET  
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1  
Source: \*  
RP: 10.255.14.144  
Flags: sparse,rptree,wildcard  
Upstream interface: Local

Group: 239.1.1.1  
Source: 10.255.14.144  
Flags: sparse,spt  
Upstream interface: Local

Group: 239.1.1.1  
Source: 10.255.70.15  
Flags: sparse,spt  
Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6  
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

#### show pim join extensive (PIM Sparse Mode)

user@host> show pim join extensive  
Instance: PIM.master Family: INET  
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1  
Source: \*  
RP: 10.255.14.144  
Flags: sparse,rptree,wildcard  
Upstream interface: Local  
Upstream neighbor: Local  
Upstream state: Local RP  
Uptime: 00:03:49  
Downstream neighbors:  
  Interface: so-1/0/0.0  
    10.111.10.2 State: Join Flags: SRW Timeout: 174  
    Uptime: 00:03:49 Time since last Join: 00:01:49  
  Interface: mt-1/1/0.32768  
    10.10.47.100 State: Join Flags: SRW Timeout: Infinity

```

        Uptime: 00:03:49 Time since last Join: 00:01:49
        Number of downstream interfaces: 2

Group: 239.1.1.1
  Source: 10.255.14.144
  Flags: sparse,spt
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local Source, Local RP
  Keepalive timeout: 344
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: so-1/0/0.0
      10.111.10.2 State: Join Flags: S Timeout: 174
      Uptime: 00:03:49 Time since last Prune: 00:01:49
    Interface: mt-1/1/0.32768
      10.10.47.100 State: Join Flags: S Timeout: Infinity
      Uptime: 00:03:49 Time since last Prune: 00:01:49
  Number of downstream interfaces: 2

Group: 239.1.1.1
  Source: 10.255.70.15
  Flags: sparse,spt
  Upstream interface: so-1/0/0.0
  Upstream neighbor: 10.111.10.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout: 344
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: Pseudo-GMP
      fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
    Interface: so-1/0/0.0 (pruned)
      10.111.10.2 State: Prune Flags: SR Timeout: 174
      Uptime: 00:03:49 Time since last Prune: 00:01:49
    Interface: mt-1/1/0.32768
      10.10.47.100 State: Join Flags: S Timeout: Infinity
      Uptime: 00:03:49 Time since last Prune: 00:01:49
  Number of downstream interfaces: 3

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

### show pim join extensive (Bidirectional PIM)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0 (RPF)
    Interface: lo0.0 (DF Winner)
  Number of downstream interfaces: 0

```

```
Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
  Downstream neighbors:
    Interface: lt-1/0/10.24
      10.0.24.4 State: Join   RW   Timeout: 185
    Interface: lt-1/0/10.23
      10.0.23.3 State: Join   RW   Timeout: 184
  Number of downstream interfaces: 2
```

```
Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
    Interface: xe-4/1/0.0      (DF Winner)
  Number of downstream interfaces: 0
```

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

#### show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

```
Group: 224.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
    Interface: xe-4/1/0.0      (DF Winner)
  Number of downstream interfaces: 0
```

**show pim join instance <instance-name> extensive**

```

user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
Source: *
RP: 10.10.47.100
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 00:03:49
Downstream neighbors:
  Interface: mt-1/1/0.32768
    10.10.47.101 State: Join Flags: SRW Timeout: 156
    Uptime: 00:03:49 Time since last Join: 00:01:49
Number of downstream interfaces: 1

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0
Upstream neighbor: 10.111.30.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0
Upstream neighbor: 10.111.20.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52

```

**show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)**

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:55
Downstream neighbors:
  Interface: Pseudo-MLDP
    Interface: lt-1/2/0.25
      1.2.5.2 State: Join Flags: S Timeout: Infinity
      Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
Source: 192.168.219.11
Flags: sparse,spt

```

```
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:41
Downstream neighbors:
  Interface: Pseudo-MLDP
```

```
Group: 232.1.1.3
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:41
Downstream neighbors:
  Interface: Pseudo-MLDP
```

```
Group: 232.2.2.2
Source: 1.2.7.7
Flags: sparse,spt
Upstream interface: lt-1/2/0.27
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:25
Downstream neighbors:
  Interface: Pseudo-MLDP
```

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

```
Group: ff3e::1:2
Source: abcd::1:2:7:7
Flags: sparse,spt
Upstream interface: lt-1/2/0.27
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:26
Downstream neighbors:
  Interface: Pseudo-MLDP
```

#### show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

```
Group: 227.1.1.1
Source: *
RP: 1.1.1.1
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 11:31:33
Downstream neighbors:
  Interface: fe-1/3/0.0
    192.168.209.9 State: Join Flags: SRW Timeout: Infinity
```

Uptime: 11:31:33 Time since last Join: 11:31:32

Group: 232.1.1.1

Source: 192.168.219.11  
 Flags: sparse,spt  
 Upstream protocol: MLDP  
 Upstream interface: Pseudo MLDP  
 Upstream neighbor: MLDP LSP root <1.1.1.2>  
 Upstream state: Join to Source  
 Keepalive timeout:  
 Uptime: 11:31:32  
 Downstream neighbors:  
   Interface: so-0/1/3.0  
     192.168.92.9 State: Join Flags: S   Timeout: Infinity  
     Uptime: 11:31:30 Time since last Join: 11:31:30  
 Downstream neighbors:  
   Interface: fe-1/3/0.0  
     192.168.209.9 State: Join Flags: S   Timeout: Infinity  
     Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.2

Source: 192.168.219.11  
 Flags: sparse,spt  
 Upstream protocol: MLDP  
 Upstream interface: Pseudo MLDP  
 Upstream neighbor: MLDP LSP root <1.1.1.2>  
 Upstream state: Join to Source  
 Keepalive timeout:  
 Uptime: 11:31:32  
 Downstream neighbors:  
   Interface: so-0/1/3.0  
     192.168.92.9 State: Join Flags: S   Timeout: Infinity  
     Uptime: 11:31:30 Time since last Join: 11:31:30  
 Downstream neighbors:  
   Interface: lt-1/2/0.14  
     1.1.4.4 State: Join Flags: S Timeout: 177  
     Uptime: 11:30:33 Time since last Join: 00:00:33  
 Downstream neighbors:  
   Interface: fe-1/3/0.0  
     192.168.209.9 State: Join Flags: S   Timeout: Infinity  
     Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.3

Source: 192.168.219.11  
 Flags: sparse,spt  
 Upstream protocol: MLDP  
 Upstream interface: Pseudo MLDP  
 Upstream neighbor: MLDP LSP root <1.1.1.2>  
 Upstream state: Join to Source  
 Keepalive timeout:  
 Uptime: 11:31:32  
 Downstream neighbors:  
   Interface: fe-1/3/0.0  
     192.168.209.9 State: Join Flags: S   Timeout: Infinity  
     Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.2.2.2

Source: 1.2.7.7  
 Flags: sparse,spt  
 Upstream protocol: MLDP  
 Upstream interface: Pseudo MLDP

```
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 11:31:30
Downstream neighbors:
  Interface: so-0/1/3.0
    192.168.92.9 State: Join Flags: S   Timeout: Infinity
    Uptime: 11:31:30 Time since last Join: 11:31:30

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
Source: abcd::1:2:7:7
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 11:31:32
Downstream neighbors:
  Interface: fe-1/3/0.0
    fe80::21f:12ff:fea5:c4db State: Join Flags: S   Timeout: Infinity
    Uptime: 11:31:32 Time since last Join: 11:31:32
```

#### show pim join extensive (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show pim join 225.1.1.1 extensive sg
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.1.1.1
Source: 10.0.0.1
Flags: sparse,spt
Active upstream interface: fe-1/2/13.0
Active upstream neighbor: 10.0.0.9
MoFRR Backup upstream interface: fe-1/2/14.0
MoFRR Backup upstream neighbor: 10.0.0.21
Upstream state: Join to Source, No Prune to RP
Keepalive timeout: 354
Uptime: 00:00:06
Downstream neighbors:
  Interface: fe-1/2/15.0
    10.0.0.13 State: Join Flags: S   Timeout: Infinity
    Uptime: 00:00:06 Time since last Join: 00:00:06
Number of downstream interfaces: 1
```



## show pim neighbors

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3929</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3929</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | <pre>show pim neighbors &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance (<i>instance-name</i>   all)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim neighbors &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance (<i>instance-name</i>   all)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Support for the <b>instance all</b> option added in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                     |
| <b>Description</b>                                  | Display information about Protocol Independent Multicast (PIM) neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                                      | <p><b>none</b>—(Same as <b>brief</b>) Display standard information about PIM neighbors for all supported family addresses for the main instance.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance (<i>instance-name</i>   all)</b>—(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                        | <a href="#">show pim neighbors on page 3931</a><br><a href="#">show pim neighbors brief on page 3931</a><br><a href="#">show pim neighbors instance on page 3931</a><br><a href="#">show pim neighbors detail on page 3931</a><br><a href="#">show pim neighbors detail (With BFD) on page 3932</a>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>                                | <p><a href="#">Table 415 on page 3930</a> describes the output fields for the <b>show pim neighbors</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 415: show pim neighbors Output Fields

| Field Name                                       | Field Description                                                                                                                                                                                                                                                                                              | Level of Output   |
|--------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| <b>Instance</b>                                  | Name of the routing instance.                                                                                                                                                                                                                                                                                  | All levels        |
| <b>Interface</b>                                 | Interface through which the neighbor is reachable.                                                                                                                                                                                                                                                             | All levels        |
| <b>Neighbor addr</b>                             | Address of the neighboring PIM routing device.                                                                                                                                                                                                                                                                 | All levels        |
| <b>IP</b>                                        | IP version: 4 or 6.                                                                                                                                                                                                                                                                                            | All levels        |
| <b>V</b>                                         | PIM version running on the neighbor: 1 or 2.                                                                                                                                                                                                                                                                   | All levels        |
| <b>Mode</b>                                      | PIM mode of the neighbor: <b>Sparse</b> , <b>Dense</b> , <b>SparseDense</b> , or <b>Unknown</b> . When the neighbor is running PIM version 2, this mode is always <b>Unknown</b> .                                                                                                                             | All levels        |
| <b>Option</b>                                    | Can be one or more of the following: <ul style="list-style-type: none"> <li>• <b>B</b>—Bidirectional Capable.</li> <li>• <b>H</b>—Hello Option Holdtime.</li> <li>• <b>G</b>—Generation Identifier.</li> <li>• <b>P</b>—Hello Option DR Priority.</li> <li>• <b>L</b>—Hello Option LAN Prune Delay.</li> </ul> | <b>brief</b> none |
| <b>Uptime</b>                                    | Time the neighbor has been operational since the PIM process was last initialized, in the format <b>dd:hh:mm:ss ago</b> for less than a week and <b>nwnd:hh:mm:ss ago</b> for more than a week.                                                                                                                | All levels        |
| <b>Address</b>                                   | Address of the neighboring PIM routing device.                                                                                                                                                                                                                                                                 | <b>detail</b>     |
| <b>BFD</b>                                       | Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: <b>Enabled</b> , <b>Operational state is up</b> , or <b>Disabled</b> .                                                                                                                                 | <b>detail</b>     |
| <b>Hello Option Holdtime</b>                     | Time for which the neighbor is available, in seconds. The range of values is 0 through 65,535.                                                                                                                                                                                                                 | <b>detail</b>     |
| <b>Hello Default Holdtime</b>                    | Default holdtime and the time remaining if the <b>holdtime</b> option is not in the received hello message.                                                                                                                                                                                                    | <b>detail</b>     |
| <b>Hello Option DR Priority</b>                  | Designated router election priority. The range of values is 0 through 255.                                                                                                                                                                                                                                     | <b>detail</b>     |
| <b>Hello Option Generation ID</b>                | 9-digit or 10-digit number used to tag hello messages.                                                                                                                                                                                                                                                         | <b>detail</b>     |
| <b>Hello Option Bi-Directional PIM supported</b> | Neighbor can process bidirectional PIM messages.                                                                                                                                                                                                                                                               | <b>detail</b>     |
| <b>Hello Option LAN Prune Delay</b>              | Time to wait before the neighbor receives prune messages, in the format <b>delay nnn ms override nnnn ms</b> .                                                                                                                                                                                                 | <b>detail</b>     |

Table 415: show pim neighbors Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                   | Level of Output |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Join Suppression supported | Neighbor is capable of join suppression.                                                                                                                                                                                                                                            | detail          |
| Rx Join                    | Information about joins received from the neighbor. <ul style="list-style-type: none"> <li><b>Group</b>—Group addresses in the join message.</li> <li><b>Source</b>—Address of the source in the join message.</li> <li><b>Timeout</b>—Time for which the join is valid.</li> </ul> | detail          |

## Sample Output

### show pim neighbors

```

user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface      IP V Mode      Option      Uptime Neighbor addr
so-1/0/0.0      4 2            HPLG        00:07:10 10.111.10.2

```

### show pim neighbors brief

The output for the **show pim neighbors brief** command is identical to that for the **show pim neighbors** command. For sample output, see [show pim neighbors on page 3931](#).

### show pim neighbors instance

```

user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface      IP V Mode      Option      Uptime Neighbor addr
at-0/3/1.0      4 2            HPLG        00:07:54 10.111.30.2
mt-1/1/0.32768  4 2            HPLG        00:07:22 10.10.47.101
so-1/0/1.0      4 2            HPLG        00:07:50 10.111.20.2

```

### show pim neighbors detail

```

user@host> show pim neighbors detail
Instance: PIM.master
Interface: ge-0/0/1.0

Address: 10.10.1.1, IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, tsf
Join Count: 2
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 2053759302
Hello Option Bi-Directional PIM supported
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

```

```
Address: 10.10.1.2, IPv4, PIM v2, sg Join Count: 0, tsg Join Count: 2
  BFD: Disabled
  Hello Option Holdtime: 105 seconds 93 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1734018161
  Hello Option Bi-Directional PIM supported
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Join Suppression supported
```

Interface: lo0.0

```
Address: 10.255.179.246, IPv4, PIM v2, Mode: SparseDense, sg Join Count:
0, tsg Join Count: 0
  Hello Option Holdtime: 65535 seconds
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1997462267
  Hello Option Bi-Directional PIM supported
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Join Suppression supported
```

#### show pim neighbors detail (With BFD)

```
user@host> show pim neighbors detail
```

Instance: PIM.master

Interface: fe-1/0/0.0

```
Address: 192.168.11.1, IPv4, PIM v2, Mode: Sparse
  Hello Option Holdtime: 65535 seconds
  Hello Option DR Priority: 1
  Hello Option Generation ID: 836607909
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Address: 192.168.11.2, IPv4, PIM v2
  BFD: Enabled, Operational state is up
  Hello Default Holdtime: 105 seconds 104 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1907549685
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Interface: fe-1/0/1.0

```
Address: 192.168.12.1, IPv4, PIM v2
  BFD: Disabled
  Hello Default Holdtime: 105 seconds 80 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1971554705
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

## show pim rps

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3933</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3933</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax</b>                                       | <pre>show pim rps &lt;brief   detail   extensive&gt; &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim rps &lt;brief   detail   extensive&gt; &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                                  | Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about PIM RPs for all groups and family addresses for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>group-address</b>—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.</p> <p><b>inet   inet6</b>—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance instance-name</b>—(Optional) Display information about RPs for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Bidirectional PIM</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                        | <a href="#">show pim rps on page 3936</a><br><a href="#">show pim rps brief on page 3936</a><br><a href="#">show pim rps &lt;group-address&gt; (Bidirectional PIM) on page 3936</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

[show pim rps <group-address> \(PIM Dense Mode\) on page 3936](#)  
[show pim rps <group-address> \(SSM Range Without asm-override-ssm Configured\) on page 3936](#)  
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Sparse-Mode RP\) on page 3937](#)  
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Bidirectional RP\) on page 3937](#)  
[show pim rps instance on page 3937](#)  
[show pim rps extensive \(PIM Sparse Mode\) on page 3937](#)  
[show pim rps extensive \(Bidirectional PIM\) on page 3938](#)  
[show pim rps extensive \(PIM Anycast RP in Use\) on page 3938](#)

**Output Fields** [Table 416 on page 3934](#) describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

**Table 416: show pim rps Output Fields**

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                          | Level of Output         |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Instance</b>                 | Name of the routing instance.                                                                                                                                                                                                                                                                                                                                                              | All levels              |
| <b>Family or Address family</b> | Name of the address family: <b>inet</b> (IPv4) or <b>inet6</b> (IPv6).                                                                                                                                                                                                                                                                                                                     | All levels              |
| <b>RP address</b>               | Address of the rendezvous point.                                                                                                                                                                                                                                                                                                                                                           | All levels              |
| <b>Type</b>                     | Type of RP: <ul style="list-style-type: none"> <li><b>auto-rp</b>—Address of the RP known through the Auto-RP protocol.</li> <li><b>bootstrap</b>—Address of the RP known through the bootstrap router protocol (BSR).</li> <li><b>embedded</b>—Address of the RP known through an embedded RP (IPv6).</li> <li><b>static</b>—Address of RP known through static configuration.</li> </ul> | <b>brief none</b>       |
| <b>Holdtime</b>                 | How long to keep the RP active, with time remaining, in seconds.                                                                                                                                                                                                                                                                                                                           | All levels              |
| <b>Timeout</b>                  | How long until the local routing device determines the RP to be unreachable, in seconds.                                                                                                                                                                                                                                                                                                   | All levels              |
| <b>Groups</b>                   | Number of groups currently using this RP.                                                                                                                                                                                                                                                                                                                                                  | All levels              |
| <b>Group prefixes</b>           | Addresses of groups that this RP can span.                                                                                                                                                                                                                                                                                                                                                 | <b>brief none</b>       |
| <b>Learned via</b>              | Address and method by which the RP was learned.                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Mode</b>                     | The PIM mode of the RP: bidirectional or sparse.<br><br>If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats.                                                                                                                                                                                                        | All levels              |
| <b>Time Active</b>              | How long the RP has been active, in the format <b>hh:mm:ss</b> .                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b> |

Table 416: show pim rps Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output                                     |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
| <b>Device Index</b>                   | Index value of the order in which Junos OS finds and initializes the interface.<br><br>For bidirectional RPs, the <b>Device Index</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>                             |
| <b>Subunit</b>                        | Logical unit number of the interface.<br><br>For bidirectional RPs, the <b>Subunit</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>                             |
| <b>Interface</b>                      | Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively.<br><br>For bidirectional RPs, the <b>Interface</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>                             |
| <b>Group Ranges</b>                   | Addresses of groups that this RP spans.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b><br><br><i>group-address</i> |
| <b>Active groups using RP</b>         | Number of groups currently using this RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>                             |
| <b>total</b>                          | Total number of active groups for this RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b>                             |
| <b>Register State for RP</b>          | Current register state for each group: <ul style="list-style-type: none"> <li>• <b>Group</b>—Multicast group address.</li> <li>• <b>Source</b>—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively:</li> <li>• <b>First Hop</b>—PIM-designated routing device that sent the Register message (the source address in the IP header).</li> <li>• <b>RP Address</b>—RP to which the Register message was sent (the destination address in the IP header).</li> <li>• <b>State</b>: <ul style="list-style-type: none"> <li>On the designated router: <ul style="list-style-type: none"> <li>• <b>Send</b>—Sending Register messages.</li> <li>• <b>Probe</b>—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages.</li> <li>• <b>Suppress</b>—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to <b>Probe</b> state.</li> </ul> </li> <li>On the RP: <ul style="list-style-type: none"> <li>• <b>Receive</b>—Receiving Register messages.</li> </ul> </li> </ul> </li> </ul> | <b>extensive</b>                                    |
| <b>Anycast-PIM rpset</b>              | If anycast RP is configured, the addresses of the RPs in the set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>extensive</b>                                    |
| <b>Anycast-PIM local address used</b> | If anycast RP is configured, the local address used by the RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>                                    |

Table 416: show pim rps Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output      |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| <b>Anycast-PIM Register State</b> | <p>If anycast RP is configured, the current register state for each group:</p> <ul style="list-style-type: none"> <li>• <b>Group</b>—Multicast group address.</li> <li>• <b>Source</b>—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively.</li> <li>• <b>Origin</b>—How the information was obtained: <ul style="list-style-type: none"> <li>• <b>DIRECT</b>—From a local attachment</li> <li>• <b>MSDP</b>—From the Multicast Source Discovery Protocol (MSDP)</li> <li>• <b>DR</b>—From the designated router</li> </ul> </li> </ul> | <b>extensive</b>     |
| <b>RP selected</b>                | For sparse mode and bidirectional mode, the identity of the RP for the specified group address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <i>group-address</i> |

## Sample Output

### show pim rps

```

user@host> show pim rps
Instance: PIM.master
Address family INET
RP address      Type      Mode   Holdtime Timeout Groups  Group prefixes
10.10.1.3       static   bidir   150     None     2  224.1.3.0/24
                225.1.3.0/24
10.10.13.2      static   bidir   150     None     2  224.1.1.0/24
                225.1.1.0/24

```

### show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see [show pim rps on page 3936](#).

### show pim rps <group-address> (Bidirectional PIM)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

224.1.0.0/16
  11.4.12.75 (Bidirectional)

RP selected: 11.4.12.75

```

### show pim rps <group-address> (PIM Dense Mode)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

Dense Mode active for group 224.1.1.1

```

### show pim rps <group-address> (SSM Range Without asm-override-ssm Configured)

```

user@host> show pim rps 224.1.1.1

```



Instance: PIM.master

Source-specific Mode (SSM) active for group 224.1.1.1

#### show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP)

user@host> show pim rps 224.1.1.1

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16  
11.4.12.75

RP selected: 11.4.12.75

#### show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP)

user@host> show pim rps 224.1.1.1

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16  
11.4.12.75 (Bidirectional)

RP selected: (null)

#### show pim rps instance

user@host> show pim rps instance VPN-A

Instance: PIM.VPN-A

Address family INET

| RP address   | Type   | Holdtime | Timeout | Groups | Group prefixes |
|--------------|--------|----------|---------|--------|----------------|
| 10.10.47.100 | static | 0        | None    | 1      | 224.0.0.0/4    |

Address family INET6

#### show pim rps extensive (PIM Sparse Mode)

user@host> show pim rps extensive

Instance: PIM.master

Family: INET

RP: 10.255.245.91

Learned via: static configuration

Time Active: 00:05:48

Holdtime: 45 with 36 remaining

Device Index: 122

Subunit: 32768

Interface: pd-6/0/0.32768

Group Ranges:

224.0.0.0/4, 36s remaining

Active groups using RP:

225.1.1.1

total 1 groups active

Register State for RP:

| Group     | Source         | FirstHop      | RP Address    | State   | Timeout |
|-----------|----------------|---------------|---------------|---------|---------|
| 225.1.1.1 | 192.168.195.78 | 10.255.14.132 | 10.255.245.91 | Receive | 0       |

**show pim rps extensive (Bidirectional PIM)**

```
user@host> show pim rps extensive
Instance: PIM.master
Address family INET

RP: 10.10.1.3
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
    224.1.3.0/24
    225.1.3.0/24

RP: 10.10.13.2
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
    224.1.1.0/24
    225.1.1.0/24
```

**show pim rps extensive (PIM Anycast RP in Use)**

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.10.10.2
Learned via: static configuration
Time Active: 00:54:52
Holdtime: 0
Device Index: 130
Subunit: 32769
Interface: pimd.32769
Group Ranges:
    224.0.0.0/4
Active groups using RP:
    224.10.10.10

    total 1 groups active

Anycast-PIM rpset:
    10.100.111.34
    10.100.111.17
    10.100.111.55

Anycast-PIM local address used: 10.100.111.1
Anycast-PIM Register State:

```

| Group        | Source     | Origin |
|--------------|------------|--------|
| 224.1.1.1    | 10.10.95.2 | DIRECT |
| 224.1.1.2    | 10.10.95.2 | DIRECT |
| 224.10.10.10 | 10.10.70.1 | MSDP   |
| 224.10.10.11 | 10.10.70.1 | MSDP   |
| 224.20.20.1  | 10.10.71.1 | DR     |

```
Address family INET6

Anycast-PIM rpset:
```

```
ab::1
ab::2
Anycast-PIM local address used: cd::1
```

Anycast-PIM Register State:

| Group         | Source       | Origin |
|---------------|--------------|--------|
| ::224.1.1.1   | ::10.10.95.2 | DIRECT |
| ::224.1.1.2   | ::10.10.95.2 | DIRECT |
| ::224.20.20.1 | ::10.10.71.1 | DR     |

## show pim source

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3940</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3940</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax</b>                                       | <pre>show pim source &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;source-prefix&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim source &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;source-prefix&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                  | Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about the PIM RPF state for all supported family addresses for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>source-prefix</b>—(Optional) Display the state for source RPF states in the given range.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>                        | <a href="#">show pim source on page 3941</a><br><a href="#">show pim source brief on page 3941</a><br><a href="#">show pim source detail on page 3941</a><br><a href="#">show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 3942</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>                                | <a href="#">Table 417 on page 3941</a> describes the output fields for the <b>show pim source</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

Table 417: show pim source Output Fields

| Field Name                | Field Description                                                                                                                                                                     |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b>           | Name of the routing instance.                                                                                                                                                         |
| <b>Source</b>             | Address of the source or reverse path.                                                                                                                                                |
| <b>Prefix/length</b>      | Prefix and prefix length for the route used to reach the RPF address.                                                                                                                 |
| <b>Upstream Protocol</b>  | Protocol toward the source address.                                                                                                                                                   |
| <b>Upstream interface</b> | RPF interface toward the source address.<br><br>A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.            |
| <b>Upstream Neighbor</b>  | Address of the RPF neighbor used to reach the source address.<br><br>The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling. |

## Sample Output

### show pim source

```

user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
  Prefix 10.255.14.144/32
  Upstream interface Local
  Upstream neighbor Local

Source 10.255.70.15
  Prefix 10.255.70.15/32
  Upstream interface so-1/0/0.0
  Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6

```

### show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see [show pim source on page 3941](#).

### show pim source detail

```

user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
  Prefix 10.255.14.144/32
  Upstream interface Local
  Upstream neighbor Local
  Active groups:228.0.0.0
  239.1.1.1

```

239.1.1.1

Source 10.255.70.15  
Prefix 10.255.70.15/32  
Upstream interface so-1/0/0.0  
Upstream neighbor 10.111.10.2  
Active groups:239.1.1.1

Instance: PIM.master Family: INET6

#### show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

user@host> show pim source

Instance: PIM.master Family: INET

Source 1.1.1.1  
Prefix 1.1.1.1/32  
Upstream interface Local  
Upstream neighbor Local

Source 1.2.7.7  
Prefix 1.2.7.0/24  
Upstream protocol MLDP  
Upstream interface Pseudo MLDP  
Upstream neighbor MLDP LSP root <1.1.1.2>

Source 192.168.219.11  
Prefix 192.168.219.0/28  
Upstream protocol MLDP  
Upstream interface Pseudo MLDP  
Upstream neighbor MLDP LSP root <1.1.1.2>

Instance: PIM.master Family: INET6

Source abcd::1:2:7:7  
Prefix abcd::1:2:7:0/120  
Upstream protocol MLDP  
Upstream interface Pseudo MLDP  
Upstream neighbor MLDP LSP root <1.1.1.2>

## show pim statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 3943</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 3943</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax</b>                                       | <pre>show pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>                                                                                                                                                                              |
| <b>Description</b>                                  | Display Protocol Independent Multicast (PIM) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display PIM statistics.</p> <p><b>inet   inet6</b>—(Optional) Display IPv4 or IPv6 PIM statistics, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM).</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display statistics about the specified interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">clear pim statistics on page 3846</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                        | <a href="#">show pim statistics on page 3950</a><br><a href="#">show pim statistics inet interface &lt;interface-name&gt; on page 3952</a><br><a href="#">show pim statistics inet6 interface &lt;interface-name&gt; on page 3952</a><br><a href="#">show pim statistics instance &lt;instance-name&gt; on page 3953</a><br><a href="#">show pim statistics interface &lt;interface-name&gt; on page 3954</a>                                                                                                                                                         |
| <b>Output Fields</b>                                | <p><a href="#">Table 418 on page 3944</a> describes the output fields for the <b>show pim statistics</b> command.</p> <p>Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                  |

Table 418: show pim statistics Output Fields

| Field Name       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Instance         | <p>Name of the routing instance.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> <li>• <b>inet interface <i>interface-name</i></b></li> <li>• <b>inet6 interface <i>interface-name</i></b></li> <li>• <b>interface <i>interface-name</i></b></li> </ul>                                                                                                 |
| Family           | <p>Output is for IPv4 or IPv6 PIM statistics. <b>INET</b> indicates IPv4 statistics, and <b>INET6</b> indicates IPv6 statistics.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> <li>• <b>inet interface <i>interface-name</i></b></li> <li>• <b>inet6 interface <i>interface-name</i></b></li> <li>• <b>interface <i>interface-name</i></b></li> </ul> |
| PIM statistics   | PIM statistics for all interfaces or for the specified interface.                                                                                                                                                                                                                                                                                                                                                       |
| PIM message type | Message type for which statistics are displayed.                                                                                                                                                                                                                                                                                                                                                                        |
| Received         | Number of received statistics.                                                                                                                                                                                                                                                                                                                                                                                          |
| Sent             | Number of messages sent of a certain type.                                                                                                                                                                                                                                                                                                                                                                              |
| Rx errors        | Number of received packets that contained errors.                                                                                                                                                                                                                                                                                                                                                                       |
| V2 Hello         | PIM version 2 hello packets.                                                                                                                                                                                                                                                                                                                                                                                            |
| V2 Register      | PIM version 2 register packets.                                                                                                                                                                                                                                                                                                                                                                                         |
| V2 Register Stop | PIM version 2 register stop packets.                                                                                                                                                                                                                                                                                                                                                                                    |
| V2 Join Prune    | PIM version 2 join and prune packets.                                                                                                                                                                                                                                                                                                                                                                                   |
| V2 Bootstrap     | PIM version 2 bootstrap packets.                                                                                                                                                                                                                                                                                                                                                                                        |
| V2 Assert        | PIM version 2 assert packets.                                                                                                                                                                                                                                                                                                                                                                                           |
| V2 Graft         | PIM version 2 graft packets.                                                                                                                                                                                                                                                                                                                                                                                            |
| V2 Graft Ack     | PIM version 2 graft acknowledgment packets.                                                                                                                                                                                                                                                                                                                                                                             |
| V2 Candidate RP  | PIM version 2 candidate RP packets.                                                                                                                                                                                                                                                                                                                                                                                     |
| V2 State Refresh | <p>PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh.</p> <p>State refresh is an extension to PIM-DM. It not supported in Junos OS.</p>                                                                                                                                                                                                                                                   |



Table 418: show pim statistics Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                   |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------|
| V2 DF Election                   | PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election.            |
| V1 Query                         | PIM version 1 query packets.                                                                                        |
| V1 Register                      | PIM version 1 register packets.                                                                                     |
| V1 Register Stop                 | PIM version 1 register stop packets.                                                                                |
| V1 Join Prune                    | PIM version 1 join and prune packets.                                                                               |
| V1 RP Reachability               | PIM version 1 RP reachability packets.                                                                              |
| V1 Assert                        | PIM version 1 assert packets.                                                                                       |
| V1 Graft                         | PIM version 1 graft packets.                                                                                        |
| V1 Graft Ack                     | PIM version 1 graft acknowledgment packets.                                                                         |
| AutoRP Announce                  | Auto-RP announce packets.                                                                                           |
| AutoRP Mapping                   | Auto-RP mapping packets.                                                                                            |
| AutoRP Unknown type              | Auto-RP packets with an unknown type.                                                                               |
| Anycast Register                 | Auto-RP announce packets.                                                                                           |
| Anycast Register Stop            | Auto-RP announce packets.                                                                                           |
| Global Statistics                | Summary of PIM statistics for all interfaces.                                                                       |
| Hello dropped on neighbor policy | Number of hello packets dropped because of a configured neighbor policy.                                            |
| Unknown type                     | Number of PIM control packets received with an unknown type.                                                        |
| V1 Unknown type                  | Number of PIM version 1 control packets received with an unknown type.                                              |
| Unknown Version                  | Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.          |
| Neighbor unknown                 | Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.              |
| Bad Length                       | Number of PIM control packets received for which the packet size does not match the PIM length field in the packet. |

Table 418: show pim statistics Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| Bad Checksum                    | Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet. |
| Bad Receive If                  | Number of PIM control packets received on an interface that does not have PIM configured.                                 |
| Rx Bad Data                     | Number of PIM control packets received that contain data for TCP Bad register packets.                                    |
| Rx Intf disabled                | Number of PIM control packets received on an interface that has PIM disabled.                                             |
| Rx V1 Require V2                | Number of PIM version 1 control packets received on an interface configured for PIM version 2.                            |
| Rx V2 Require V1                | Number of PIM version 2 control packets received on an interface configured for PIM version 1.                            |
| Rx Register not RP              | Number of PIM register packets received when the routing device is not the RP for the group.                              |
| Rx Register no route            | Number of PIM register packets received when the RP does not have a unicast route back to the source.                     |
| Rx Register no decap if         | Number of PIM register packets received when the RP does not have a de-encapsulation interface.                           |
| Null Register Timeout           | Number of NULL register timeout packets.                                                                                  |
| RP Filtered Source              | Number of PIM packets received when the routing device has a source address filter configured for the RP.                 |
| Rx Unknown Reg Stop             | Number of register stop messages received with an unknown type.                                                           |
| Rx Join/Prune no state          | Number of join and prune messages received for which the routing device has no state.                                     |
| Rx Join/Prune on upstream if    | Number of join and prune messages received on the interface used to reach the upstream routing device, toward the RP.     |
| Rx Join/Prune for invalid group | Number of join or prune messages received for invalid multicast group addresses.                                          |
| Rx Join/Prune messages dropped  | Number of join and prune messages received and dropped.                                                                   |
| Rx sparse join for dense group  | Number of PIM sparse mode join messages received for a group that is configured for dense mode.                           |

Table 418: show pim statistics Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                           |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rx Graft/Graft Ack no state | Number of graft and graft acknowledgment messages received for which the router or switch has no state.                                                                     |
| Rx Graft on upstream if     | Number of graft messages received on the interface used to reach the upstream routing device, toward the RP.                                                                |
| Rx CRP not BSR              | Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap.                                                                 |
| Rx BSR when BSR             | Number of BSR messages received in which the PIM message type is Bootstrap.                                                                                                 |
| Rx BSR not RPF if           | Number of BSR messages received on an interface that is not the RPF interface.                                                                                              |
| Rx unknown hello opt        | Number of PIM hello packets received with options that Junos OS does not support.                                                                                           |
| Rx data no state            | Number of PIM control packets received for which the routing device has no state for the data type.                                                                         |
| Rx RP no state              | Number of PIM control packets received for which the routing device has no state for the RP.                                                                                |
| Rx aggregate                | Number of PIM aggregate MDT packets received.                                                                                                                               |
| Rx malformed packet         | Number of PIM control packets received with a malformed IP unicast or multicast address family.                                                                             |
| No RP                       | Number of PIM control packets received with no RP address.                                                                                                                  |
| No register encaps if       | Number of PIM register packets received when the first-hop routing device does not have an encapsulation interface.                                                         |
| No route upstream           | Number of PIM control packets received when the routing device does not have a unicast route to the the interface used to reach the upstream routing device, toward the RP. |
| Nexthop Unusable            | Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down.                                                  |
| RP mismatch                 | Number of PIM control packets received for which the routing device has an RP mismatch.                                                                                     |
| RP mode mismatch            | RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages.                                                                           |
| RPF neighbor unknown        | Number of PIM control packets received for which the routing device has an unknown RPF neighbor for the source.                                                             |

Table 418: show pim statistics Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rx Joins/Prunes filtered            | The number of join and prune messages filtered because of configured route filters and source address filters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Tx Joins/Prunes filtered            | The number of join and prune messages filtered because of configured route filters and source address filters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Embedded-RP invalid addr            | Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Embedded-RP limit exceed            | Number of times the limit configured with the <b>maximum-rps</b> statement is exceeded. The <b>maximum-rps</b> statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Embedded-RP added                   | <p>Number of packets in which the embedded RP for IPv6 is added.</p> <p>The following receive events trigger extraction of an IPv6 embedded RP address on the routing device:</p> <ul style="list-style-type: none"> <li>• Multicast Listener Discovery (MLD) report for an embedded RP multicast group address</li> <li>• PIM join message with an embedded RP multicast group address</li> <li>• Static embedded RP multicast group address associated with an interface</li> <li>• Packets sent to an embedded RP multicast group address received on the DR</li> </ul> <p>An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.</p> |
| Embedded-RP removed                 | Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Rx Register msgs filtering drop     | Number of received register messages dropped because of a filter configured for PIM register messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Tx Register msgs filtering drop     | Number of register messages dropped because of a filter configured for PIM register messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Rx Bidir Join/Prune on non-Bidir if | Error counter for join and prune messages received on non-bidirectional PIM interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Rx Bidir Join/Prune on non-DF if    | Error counter for join and prune messages received on non-designated forwarder interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| V4 (S,G) Maximum                    | Maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

Table 418: show pim statistics Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                   |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| V4 (S,G) Accepted                | Number of accepted (S,G) IPv4 multicast routes.                                                                                                                                     |
| V4 (S,G) Threshold               | Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device).                                            |
| V4 (S,G) Log Interval            | Time (in seconds) between consecutive log messages.                                                                                                                                 |
| V6 (S,G) Maximum                 | Maximum number of (S,G) IPv6 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted. |
| V6 (S,G) Accepted                | Number of accepted (S,G) IPv6 multicast routes.                                                                                                                                     |
| V6 (S,G) Threshold               | Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv6 multicast routes accepted by the device).                                            |
| V6 (S,G) Log Interval            | Time (in seconds) between consecutive log messages.                                                                                                                                 |
| V4 (grp-prefix, RP) Maximum      | Maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.        |
| V4 (grp-prefix, RP) Accepted     | Number of accepted group-to-RP IPv4 multicast mappings.                                                                                                                             |
| V4 (grp-prefix, RP) Threshold    | Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device).                                    |
| V4 (grp-prefix, RP) Log Interval | Time (in seconds) between consecutive log messages.                                                                                                                                 |
| V6 (grp-prefix, RP) Maximum      | Maximum number of group-to RP IPv6 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.                           |
| V6 (grp-prefix, RP) Accepted     | Number of accepted group-to-RP IPv6 multicast mappings.                                                                                                                             |
| V6 (grp-prefix, RP) Threshold    | Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv6 multicast mappings accepted by the device).                                    |
| V6 (grp-prefix, RP) Log Interval | Time (in seconds) between consecutive log messages.                                                                                                                                 |

Table 418: show pim statistics Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                 |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| V4 Register Maximum                   | Maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted.<br><br>You configure the register limits on the RP.                                                 |
| V4 Register Accepted                  | Number of accepted IPv4 PIM registers.                                                                                                                                                                                                            |
| V4 Register Threshold                 | Threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device).                                                                                                                   |
| V4 Register Log Interval              | Time (in seconds) between consecutive log messages.                                                                                                                                                                                               |
| V6 Register Maximum                   | Maximum number of IPv6 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted.<br><br>You configure the register limits on the RP.                                                 |
| V6 Register Accepted                  | Number of accepted IPv6 PIM registers.                                                                                                                                                                                                            |
| V6 Register Threshold                 | Threshold at which a warning message is logged (percentage of the maximum number of IPv6 PIM registers accepted by the device).                                                                                                                   |
| V6 Register Log Interval              | Time (in seconds) between consecutive log messages.                                                                                                                                                                                               |
| (*G) Join drop due to SSM range check | PIM join messages that are dropped because the multicast addresses are outside of the SSM address range of 232.0.0.0 through 232.255.255.255. You can extend the accepted SSM address range by configuring the <code>ssm-groups</code> statement. |

## Sample Output

### show pim statistics

```

user@host> show pim statistics
PIM Message type    Received    Sent    Rx errors
V2 Hello            15          32         0
V2 Register          0          362        0
V2 Register Stop     483         0         0
V2 Join Prune        18          518        0
V2 Bootstrap         0           0         0
V2 Assert            0           0         0
V2 Graft             0           0         0
V2 Graft Ack         0           0         0
V2 Candidate RP      0           0         0
V2 State Refresh     0           0         0
V2 DF Election       0           0         0
V1 Query             0           0         0
V1 Register          0           0         0
V1 Register Stop     0           0         0
V1 Join Prune        0           0         0

```

|                       |   |   |   |
|-----------------------|---|---|---|
| V1 RP Reachability    | 0 | 0 | 0 |
| V1 Assert             | 0 | 0 | 0 |
| V1 Graft              | 0 | 0 | 0 |
| V1 Graft Ack          | 0 | 0 | 0 |
| AutoRP Announce       | 0 | 0 | 0 |
| AutoRP Mapping        | 0 | 0 | 0 |
| AutoRP Unknown type   | 0 |   |   |
| Anycast Register      | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

## Global Statistics

|                                         |   |
|-----------------------------------------|---|
| Hello dropped on neighbor policy        | 0 |
| Unknown type                            | 0 |
| V1 Unknown type                         | 0 |
| Unknown Version                         | 0 |
| ipv4 BSR pkt drop due to excessive rate | 0 |
| ipv6 BSR pkt drop due to excessive rate | 0 |
| Neighbor unknown                        | 0 |
| Bad Length                              | 0 |
| Bad Checksum                            | 0 |
| Bad Receive If                          | 0 |
| Rx Bad Data                             | 0 |
| Rx Intf disabled                        | 0 |
| Rx V1 Require V2                        | 0 |
| Rx V2 Require V1                        | 0 |
| Rx Register not RP                      | 0 |
| Rx Register no route                    | 0 |
| Rx Register no decap if                 | 0 |
| Null Register Timeout                   | 0 |
| RP Filtered Source                      | 0 |
| Rx Unknown Reg Stop                     | 0 |
| Rx Join/Prune no state                  | 0 |
| Rx Join/Prune on upstream if            | 0 |
| Rx Join/Prune for invalid group         | 5 |
| Rx Join/Prune messages dropped          | 0 |
| Rx sparse join for dense group          | 0 |
| Rx Graft/Graft Ack no state             | 0 |
| Rx Graft on upstream if                 | 0 |
| Rx CRP not BSR                          | 0 |
| Rx BSR when BSR                         | 0 |
| Rx BSR not RPF if                       | 0 |
| Rx unknown hello opt                    | 0 |
| Rx data no state                        | 0 |
| Rx RP no state                          | 0 |
| Rx aggregate                            | 0 |
| Rx malformed packet                     | 0 |
| Rx illegal TTL                          | 0 |
| Rx illegal destination address          | 0 |
| No RP                                   | 0 |
| No register encap if                    | 0 |
| No route upstream                       | 0 |
| Nexthop Unusable                        | 0 |
| RP mismatch                             | 0 |
| RP mode mismatch                        | 0 |
| RPF neighbor unknown                    | 0 |
| Rx Joins/Prunes filtered                | 0 |
| Tx Joins/Prunes filtered                | 0 |
| Embedded-RP invalid addr                | 0 |
| Embedded-RP limit exceed                | 0 |
| Embedded-RP added                       | 0 |

```
Embedded-RP removed                0
Rx Register msgs filtering drop      0
Tx Register msgs filtering drop      0
Rx Bidir Join/Prune on non-Bidir if  0
Rx Bidir Join/Prune on non-DF if     0
(*,G) Join drop due to SSM range check 0
```

## Sample Output

**show pim statistics inet interface <interface-name>**

```
user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 4    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| V1 Query              | 0        | 0    | 0         |
| V1 Register           | 0        | 0    | 0         |
| V1 Register Stop      | 0        | 0    | 0         |
| V1 Join Prune         | 0        | 0    | 0         |
| V1 RP Reachability    | 0        | 0    | 0         |
| V1 Assert             | 0        | 0    | 0         |
| V1 Graft              | 0        | 0    | 0         |
| V1 Graft Ack          | 0        | 0    | 0         |
| AutoRP Announce       | 0        | 0    | 0         |
| AutoRP Mapping        | 0        | 0    | 0         |
| AutoRP Unknown type   | 0        |      |           |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |

## Sample Output

**show pim statistics inet6 interface <interface-name>**

```
user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 4    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |



**show pim statistics instance <instance-name>**

```

user@host> show pim statistics instance VPN-A
PIM Message type      Received      Sent  Rx errors
V2 Hello              31           37      0
V2 Register           0            0      0
V2 Register Stop      0            0      0
V2 Join Prune         0           16      0
V2 Bootstrap          0            0      0
V2 Assert             0            0      0
V2 Graft              0            0      0
V2 Graft Ack          0            0      0
V2 Candidate RP       0            0      0
V2 State Refresh      0            0      0
V2 DF Election        0            0      0
V1 Query              0            0      0
V1 Register           0            0      0
V1 Register Stop      0            0      0
V1 Join Prune         0            0      0
V1 RP Reachability    0            0      0
V1 Assert             0            0      0
V1 Graft              0            0      0
V1 Graft Ack          0            0      0
AutoRP Announce       0            0      0
AutoRP Mapping         0            0      0
AutoRP Unknown type   0            0      0
Anycast Register       0            0      0
Anycast Register Stop 0            0      0

```

**Global Statistics**

```

Hello dropped on neighbor policy      0
Unknown type                          0
V1 Unknown type                       0
Unknown Version                       0
Neighbor unknown                      0
Bad Length                            0
Bad Checksum                          0
Bad Receive If                        0
Rx Bad Data                           0
Rx Intf disabled                      0
Rx V1 Require V2                      0
Rx V2 Require V1                      0
Rx Register not RP                    0
Rx Register no route                  0
Rx Register no decap if                0
Null Register Timeout                 0
RP Filtered Source                    0
Rx Unknown Reg Stop                   0
Rx Join/Prune no state                0
Rx Join/Prune on upstream if          0
Rx Join/Prune for invalid group        0
Rx Join/Prune messages dropped         0
Rx sparse join for dense group         0
Rx Graft/Graft Ack no state           0
Rx Graft on upstream if               0
Rx CRP not BSR                        0
Rx BSR when BSR                       0
Rx BSR not RPF if                     0
Rx unknown hello opt                  0
Rx data no state                      0

```

|                                        |     |
|----------------------------------------|-----|
| Rx RP no state                         | 0   |
| Rx aggregate                           | 0   |
| Rx malformed packet                    | 0   |
| Rx illegal TTL                         | 0   |
| Rx illegal destination address         | 0   |
| No RP                                  | 0   |
| No register encap if                   | 0   |
| No route upstream                      | 28  |
| Nexthop Unusable                       | 0   |
| RP mismatch                            | 0   |
| RP mode mismatch                       | 0   |
| RPF neighbor unknown                   | 0   |
| Rx Joins/Prunes filtered               | 0   |
| Tx Joins/Prunes filtered               | 0   |
| Embedded-RP invalid addr               | 0   |
| Embedded-RP limit exceed               | 0   |
| Embedded-RP added                      | 0   |
| Embedded-RP removed                    | 0   |
| Rx Register msgs filtering drop        | 0   |
| Tx Register msgs filtering drop        | 0   |
| Rx Bidir Join/Prune on non-Bidir if    | 0   |
| Rx Bidir Join/Prune on non-DF if       | 0   |
| V4 (S,G) Maximum                       | 10  |
| V4 (S,G) Accepted                      | 9   |
| V4 (S,G) Threshold                     | 80  |
| V4 (S,G) Log Interval                  | 80  |
| V6 (S,G) Maximum                       | 8   |
| V6 (S,G) Accepted                      | 8   |
| V6 (S,G) Threshold                     | 50  |
| V6 (S,G) Log Interval                  | 100 |
| V4 (grp-prefix, RP) Maximum            | 100 |
| V4 (grp-prefix, RP) Accepted           | 5   |
| V4 (grp-prefix, RP) Threshold          | 80  |
| V4 (grp-prefix, RP) Log Interval       | 10  |
| V6 (grp-prefix, RP) Maximum            | 20  |
| V6 (grp-prefix, RP) Accepted           | 0   |
| V6 (grp-prefix, RP) Threshold          | 90  |
| V6 (grp-prefix, RP) Log Interval       | 20  |
| V4 Register Maximum                    | 100 |
| V4 Register Accepted                   | 10  |
| V4 Register Threshold                  | 80  |
| V4 Register Log Interval               | 10  |
| V6 Register Maximum                    | 20  |
| V6 Register Accepted                   | 0   |
| V6 Register Threshold                  | 90  |
| V6 Register Log Interval               | 20  |
| (*,G) Join drop due to SSM range check | 0   |

## Sample Output

show pim statistics interface <interface-name>

```

user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type      Received      Sent  Rx errors
V2 Hello                0             3         0
V2 Register             0             0         0
V2 Register Stop        0             0         0

```

|                       |   |   |   |
|-----------------------|---|---|---|
| V2 Join Prune         | 0 | 0 | 0 |
| V2 Bootstrap          | 0 | 0 | 0 |
| V2 Assert             | 0 | 0 | 0 |
| V2 Graft              | 0 | 0 | 0 |
| V2 Graft Ack          | 0 | 0 | 0 |
| V2 Candidate RP       | 0 | 0 | 0 |
| V1 Query              | 0 | 0 | 0 |
| V1 Register           | 0 | 0 | 0 |
| V1 Register Stop      | 0 | 0 | 0 |
| V1 Join Prune         | 0 | 0 | 0 |
| V1 RP Reachability    | 0 | 0 | 0 |
| V1 Assert             | 0 | 0 | 0 |
| V1 Graft              | 0 | 0 | 0 |
| V1 Graft Ack          | 0 | 0 | 0 |
| AutoRP Announce       | 0 | 0 | 0 |
| AutoRP Mapping        | 0 | 0 | 0 |
| AutoRP Unknown type   | 0 |   |   |
| Anycast Register      | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 3    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |



## PART 21

# Network Managment and Monitoring

- [Overview on page 3959](#)
- [Configuration on page 3995](#)
- [Administration on page 4263](#)
- [Troubleshooting on page 4413](#)



## CHAPTER 66

# Overview

- [Real-Time Performance Monitoring on page 3959](#)
- [Analyzers and Port Mirroring on page 3963](#)
- [sFlow Monitoring Technology on page 3969](#)
- [Ethernet OAM Link Fault Management and Connectivity Fault Management on page 3971](#)
- [Uplink Failure Detection on page 3976](#)
- [Network Analytics on page 3978](#)

### **[Real-Time Performance Monitoring](#)**

---

- [Understanding Real-Time Performance Monitoring on EX Series Switches on page 3960](#)

## Understanding Real-Time Performance Monitoring on EX Series Switches

Real-time performance monitoring (RPM) enables you to configure active probes to track and monitor traffic across the network and to investigate network problems. You can use RPM with Juniper Networks EX Series Ethernet Switches.

The ways in which you can use RPM include:

- Monitor time delays between devices.
- Monitor time delays at the protocol level.
- Set thresholds to trigger SNMP traps when values are exceeded.

You can configure thresholds for round-trip time, ingress or egress delay, standard deviation, jitter, successive lost probes, and total lost probes per test. (SNMP trap results are stored in **pingResultsTable**, **jnxPingResultsTable**, **jnxPingProbeHistoryTable**, and **pingProbeHistoryTable**.)

- Determine automatically whether a path exists between a host router or switch and its configured BGP neighbors. You can view the results of the discovery using an SNMP client.
- Use the history of the most recent 50 probes to analyze trends in your network and predict future needs.

RPM provides MIB support with extensions for RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*.

This topic includes:

- [RPM Packet Collection on page 3960](#)
- [Tests and Probe Types on page 3960](#)
- [Hardware Timestamps on page 3961](#)
- [Limitations of RPM on EX Series Switches on page 3963](#)

---

### RPM Packet Collection

Probes collect packets per destination and per application, including ping Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets.

---

### Tests and Probe Types

A test can contain multiple probes. The probe type specifies the packet and protocol contents of the probe.

EX Series switches support the following tests and probe types:

- Ping tests:
  - ICMP echo probe



- ICMP timestamp probe
- HTTP tests:
  - HTTP get probe (not available for BGP RPM services)
  - HTTP get metadata probe
- UDP and TCP tests with user-configured ports:
  - UDP echo probe
  - TCP connection probe
  - UDP timestamp probe

### Hardware Timestamps

To account for latency or jitter in the communication of probe messages, you can enable timestamping of the probe packets (hardware timestamps). If hardware timestamps are not configured, then timers are generated at the software level and are less accurate than they would have been with hardware timestamps.



**NOTE:** EX Series switches support hardware timestamps for UDP and ICMP probes. EX Series switches do not support hardware timestamps for HTTP or TCP probes.

You can timestamp the following RPM probes to improve the measurement of latency or jitter:

- ICMP ping
- ICMP ping timestamp
- UDP ping
- UDP ping timestamp


You should configure the requester (the RPM client) with hardware timestamps (see [Figure 51 on page 3962](#)) to get more meaningful results than you would get without the timestamps. The responder (the RPM server) does not need to be configured to support hardware timestamps. If the responder supports hardware timestamps, it timestamps the RPM probes. If the responder does not support hardware timestamps, RPM can only report round-trip measurements that include the processing time on the responder.



**NOTE:** Hardware timestamps are supported on all EX Series switches.

[Figure 51 on page 3962](#) shows the timestamps:

Figure 51: RPM Timestamps



- T1 is the time the packet leaves the requester port.
- T2 is the time the responder receives the packet.
- T3 is the time the responder sends the response.
- T4 is the time the requester receives the response.

The round-trip time is  $(T2 - T1) + (T4 - T3)$ . If the responder does not support hardware timestamps, then the round-trip time is  $(T4 - T1) / 2$ , and thus includes the processing time of the responder.

You can use RPM probes to find the following time measurements:

- Minimum round-trip time
- Maximum round-trip time
- Average round-trip time
- Standard deviation of the round-trip time
- Jitter of the round-trip time—Difference between the minimum and maximum round-trip time



**NOTE:** See “[Configuring the Interface for RPM Timestamping for Client/Server on an EX Series Switch \(CLI Procedure\)](#)” on [page 4057](#) for information on how to configure hardware timestamps on the requester.

The RPM feature provides a configuration option to set one-way hardware timestamps. Use one-way timestamps when you want information about one-way time, rather than round-trip times, for packets to traverse the network between the requester and the responder. As shown in [Figure 51 on page 3962](#), one-way timestamps represent the time  $T2 - T1$  and the time from  $T4 - T3$ . Use one-way timestamps when you want to gather information about delay in each direction and to find egress and ingress jitter values.



**NOTE:** For correct one-way measurement, the clocks of the requester and responder must be synchronized. If the clocks are not synchronized, one-way jitter measurements and calculations can include significant variations, in some cases orders of magnitude greater than the round-trip times.

When you enable one-way timestamps in a probe, the following one-way measurements are reported:

- Minimum, maximum, standard deviation, and jitter measurements for egress and ingress times
- Number of probes sent
- Number of probe responses received
- Percentage of lost probes

---

### Limitations of RPM on EX Series Switches

- Two-Way Active Measurement Protocol (TWAMP) is not supported on EX Series switches.
- EX Series switches do not support user-configured class-of-service (CoS) classifiers or prioritization of RPM packets over regular data packets received on an input interface.
- Timestamps:
  - If the responder does not support hardware timestamps, RPM can only report the round-trip measurements and cannot calculate round-trip jitter.
  - EX Series switches do not support hardware timestamps for HTTP and TCP probes.
  - Timestamps apply only to IPv4 traffic.

#### Related Documentation

- [For further details about RPM, see \*Junos OS Services Interfaces Configuration Guide\*](#)
- [Configuring the Interface for RPM Timestamping for Client/Server on an EX Series Switch \(CLI Procedure\) on page 4057](#)
- [Configuring Real-Time Performance Monitoring \(J-Web Procedure\) on page 4050](#)
- [Configuring SNMP \(J-Web Procedure\) on page 4040](#)
- [Monitoring Network Traffic Using Traceroute on page 4265](#)

---

## Analyzers and Port Mirroring

- [Understanding Port Mirroring and Analyzers on EX4300 Switches on page 3964](#)

## Understanding Port Mirroring and Analyzers on EX4300 Switches



**NOTE:** This concept uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Understanding Port Mirroring on EX Series Switches*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

Mirroring might be needed for traffic analysis on a switch because a switch, unlike a hub, does not broadcast packets to every port on the destination device. The switch sends packets only to the port to which the destination device is connected.

Juniper Networks EX4300 Ethernet Switches support the following mirroring methods: port mirroring and analyzers. You can use port mirroring or analyzers to facilitate analyzing traffic on EX4300 switches at the packet level. You might use analyzers as part of monitoring switch traffic for such purposes as enforcing policies concerning network usage and file sharing and for identifying sources of problems on your network by locating abnormal or heavy bandwidth usage by particular stations or applications.

Mirrored packets can be copied either to a local interface for local monitoring or to a VLAN for remote monitoring. The following packets can be copied:

- **Packets entering or exiting a port**—You can mirror the packets in any combination of packets entering or exiting ports on up to 256 ports. For example, you can send copies of the packets entering some ports and the packets exiting other ports to the same local analyzer port or analyzer VLAN.
- **Packets entering a VLAN**—You can mirror the packets entering a VLAN to either a local analyzer port or to an analyzer VLAN. You can configure multiple VLANs (up to 256 VLANs), including a VLAN range and PVLANS, as ingress input to an analyzer.
- **Policy-based sample packets**—You can mirror a policy-based sample of packets that are entering a port or a VLAN. You configure a firewall filter to establish a policy to select the packets to be mirrored. You can send the sample to a port-mirroring instance or to an analyzer VLAN.

This topic describes:

- [Port Mirroring Overview on page 3964](#)
- [Analyzer Overview on page 3965](#)
- [Port Mirroring and Analyzer Terminologies on page 3965](#)
- [Configuration Guidelines for Port Mirroring and Analyzers on EX4300 Switches on page 3966](#)

---

### Port Mirroring Overview

You configure port mirroring on an EX4300 switch to send copies of unicast traffic to an output destination such as an interface, a routing-instance, or a VLAN. Then, you can analyze the mirrored traffic by using a protocol analyzer application. The protocol analyzer

application can run either on a computer connected to the analyzer output interface or on a remote monitoring station. For the input traffic, you can configure a firewall filter term to specify whether port mirroring must be applied to all packets at the interface to which the firewall filter is applied. You can apply a firewall filter configured with the action **port-mirror** or **port-mirror-instance *name*** to the input or output logical interfaces (including aggregated Ethernet logical interfaces), to traffic forwarded or flooded to a VLAN, or traffic forwarded or flooded to a VPLS routing instance. EX4300 switches support port mirroring of VPLS (**family ethernet-switching** or **family vpls**) traffic and VPN traffic with **family ccc** in a Layer 2 environment. Within a firewall filter term, you can specify the port-mirroring properties under the **then** statement in either of the following ways:

- Implicitly reference the port-mirroring properties in effect on the port.
- Explicitly reference a particular named instance of port mirroring.

You can configure port mirroring at the **[edit forwarding-options port-mirroring]** hierarchy level.

### Analyzer Overview

You can configure an analyzer to define both the input traffic and output traffic in the same analyzer configuration. The input traffic to be analyzed can be traffic that enters or exits an interface, or traffic that enters a VLAN. The analyzer configuration enables you to send this traffic to an output interface, instance, or VLAN. You can configure an analyzer at the **[edit forwarding-options analyzer]** hierarchy.

### Port Mirroring and Analyzer Terminologies

Table 419 on page 3965 lists some port mirroring terms and their descriptions.

**Table 419: Mirroring Terminologies**

| Term                                                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Analyzer                                                  | <p>In a mirroring configuration (analyzer) on an EX4300 switch, the analyzer includes:</p> <ul style="list-style-type: none"> <li>• The name of the analyzer</li> <li>• Source (input) ports or VLAN</li> <li>• A destination for mirrored packets (either a monitor port or a monitor VLAN)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                         |
| Analyzer output interface<br>(Also known as monitor port) | <p>Interface to which mirrored traffic is sent and to which a protocol analyzer application is connected.</p> <p><b>NOTE:</b> Interfaces used as output for an analyzer must be configured under the <b>ethernet-switching</b> family.</p> <p>Analyzer output interfaces have the following limitations:</p> <ul style="list-style-type: none"> <li>• Cannot also be a source port.</li> <li>• Do not participate in Layer 2 protocols, such as Spanning Tree Protocol (STP), when part of a port-mirroring configuration.</li> <li>• If the bandwidth of the analyzer output interface is not sufficient to handle the traffic from the source ports, overflow packets are dropped.</li> </ul> |

Table 419: Mirroring Terminologies (*continued*)

| Term                                                                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Analyzer VLAN<br>(Also known as monitor VLAN)                             | VLAN to which mirrored traffic is sent. The mirrored traffic can be used by a protocol analyzer application. The member interfaces in the monitor VLAN are spread across the switches in your network.                                                                                                                                                                                                                                                                                                      |
| Port mirroring                                                            | A port-mirroring configuration that does not specify an input source; it specifies only an output destination. A firewall filter configuration must be defined for the input source. A firewall filter configuration must be defined to mirror packets that match the match conditions defined in the firewall filter term. The action item <b>port-mirror-instance instance-name</b> in the firewall filter configuration is used to send packets to the analyzer and these packets form the input source. |
| Global port mirror                                                        | A port mirroring configuration that does not have an instance name. The firewall filter action <b>port-mirror</b> will be the action for the firewall filter configuration.                                                                                                                                                                                                                                                                                                                                 |
| Input interface<br>(Also known as mirrored ports or monitored interfaces) | An interface on the switch that is being mirrored. Traffic that is either entering or exiting this interface is mirrored.                                                                                                                                                                                                                                                                                                                                                                                   |
| LAG-based analyzer                                                        | An analyzer that has a link aggregation group (LAG) specified as the input (ingress) interface in the analyzer configuration.                                                                                                                                                                                                                                                                                                                                                                               |
| Local mirroring                                                           | An analyzer configuration in which packets are mirrored to a local analyzer port.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Monitoring station                                                        | A computer running a protocol analyzer application.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Native analyzer session                                                   | An analyzer session that has both <b>input</b> and <b>output</b> definitions in its analyzer configuration.                                                                                                                                                                                                                                                                                                                                                                                                 |
| Policy-based mirroring<br>(Also known as port mirroring)                  | Mirroring of packets that match the match items in the defined firewall filter term. The action item <b>port-mirror-instance instance-name</b> is used in the firewall filter to send the packets to the monitor port.                                                                                                                                                                                                                                                                                      |
| Port-based analyzer                                                       | An analyzer session whose configuration defines interfaces for both input and output.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Protocol analyzer application                                             | An application used to examine packets transmitted across a network segment. Also commonly called network analyzer, packet sniffer, or probe.                                                                                                                                                                                                                                                                                                                                                               |
| Remote port mirroring                                                     | Functions the same way as local port mirroring, except that the mirrored traffic is not copied to a local analyzer port but is flooded to an analyzer VLAN that you create specifically for the purpose of receiving mirrored traffic.                                                                                                                                                                                                                                                                      |
| VLAN-based analyzer                                                       | An analyzer session whose configuration uses VLANs for both input and output or for either input or output.                                                                                                                                                                                                                                                                                                                                                                                                 |

### Configuration Guidelines for Port Mirroring and Analyzers on EX4300 Switches

When you configure port mirroring or analyzers on EX4300 switches, we recommend that you follow certain guidelines to ensure that you obtain optimum benefit from mirroring. Additionally, we recommend that you disable mirroring when you are not using it and that you select specific interfaces for which packets must be mirrored (that is,

select specific interfaces as input to the analyzer) in preference to using the **all** keyword option, which will enable mirroring on all interfaces. Mirroring only the necessary packets reduces any potential performance impact.

With local mirroring, traffic from multiple ports is replicated to the analyzer output interface. If the output interface for an analyzer reaches capacity, packets are dropped. Thus, while configuring an analyzer, you must consider whether the traffic being mirrored exceeds the capacity of the analyzer output interface.

[Table 420 on page 3967](#) summarizes further configuration guidelines for mirroring on EX4300 switches.

**Table 420: Configuration Guidelines for Port Mirroring and Analyzers on EX4300 Switches**

| Guideline                                                                                    | Value or Support Information                                                                                                                                                                                                                                                                  | Comment                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Number of VLANs that you can use as ingress input to an analyzer.                            | 256                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Number of port-mirroring sessions and analyzers that you can enable concurrently.            | 4                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li>You can configure a total of four sessions and you can enable only one of the following at any point in time: <ul style="list-style-type: none"> <li>A maximum of four port-mirroring sessions (including the global port-mirroring session). See <a href="#">Table 419 on page 3965</a> for a description of global port mirror.</li> <li>A maximum of four analyzer sessions.</li> <li>A combination of port-mirroring and analyzer sessions, and the total of this combination must be four.</li> </ul> </li> <li>You can configure more than the specified number of port-mirroring instances or analyzers on the switch, but you can enable only the specified number for a session. Use <b>disable forwarding-options analyzer name</b> to disable an analyzer and use <b>disable forwarding-options port-mirroring instance name</b> to disable a port-mirroring instance.</li> </ul> |
| Types of ports on which you cannot mirror traffic.                                           | <ul style="list-style-type: none"> <li>Virtual Chassis ports (VCPs)</li> <li>Management Ethernet ports (<b>me0</b> or <b>vme0</b>)</li> <li>Integrated routing and bridging (IRB) interfaces; also known as routed VLAN interfaces (RVIs).</li> <li>VLAN-tagged Layer 3 interfaces</li> </ul> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Protocol families that you can include in a port-mirroring configuration for remote traffic. | any                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Table 420: Configuration Guidelines for Port Mirroring and Analyzers on EX4300 Switches** (*continued*)

| Guideline                                                                                                     | Value or Support Information | Comment                                                                           |
|---------------------------------------------------------------------------------------------------------------|------------------------------|-----------------------------------------------------------------------------------|
| Traffic directions that you can configure for mirroring on ports in firewall-filter-based configurations.     | Ingress only                 |                                                                                   |
| Mirrored packets exiting an interface reflect rewritten class-of-service (CoS) DSCP or 802.1p bits.           | Applicable                   |                                                                                   |
| Packets with physical layer errors are not sent to the local or remote analyzer.                              | Applicable                   | Packets with these errors are filtered out and thus are not sent to the analyzer. |
| Port mirroring does not support line-rate traffic.                                                            | Applicable                   | Port mirroring for line-rate traffic is done on a best-effort basis.              |
| Mirroring of packets egressing a VLAN.                                                                        | Not supported                |                                                                                   |
| Port-mirroring or analyzer output on a LAG interface.                                                         | Supported                    |                                                                                   |
| Maximum number of child members on a port-mirroring or analyzer output LAG interface.                         | 8                            |                                                                                   |
| Maximum number of interfaces in a remote port-mirroring or analyzer VLAN.                                     | 1                            |                                                                                   |
| Egress mirroring of host-generated control packets.                                                           | Not Supported                |                                                                                   |
| Configuring Layer 3 logical interfaces in the <b>input</b> stanza of an analyzer.                             | Not supported                | This functionality can be achieved by configuring port mirroring.                 |
| The analyzer input and output stanzas containing members of the same VLAN or the VLAN itself must be avoided. | Applicable                   |                                                                                   |

**Related Documentation**

- [Example: Configuring Mirroring for Local Monitoring of Employee Resource Use on EX4300 Switches on page 4000](#)
- [Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches on page 4006](#)
- [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 4047](#)
- [Configuring Mirroring on EX4300 Switches to Analyze Traffic \(CLI Procedure\) on page 4043](#)
- [Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704](#)



## sFlow Monitoring Technology

- [Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch on page 3969](#)

### Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology randomly samples network packets and sends the samples to a monitoring station. You can configure sFlow technology on a Juniper Networks EX Series Ethernet Switch to continuously monitor traffic at wire speed on all interfaces simultaneously.

This topic describes:

- [Sampling Mechanism and Architecture of sFlow Technology on EX Series Switches on page 3969](#)
- [Adaptive Sampling on page 3970](#)
- [sFlow Agent Address Assignment on page 3971](#)

### Sampling Mechanism and Architecture of sFlow Technology on EX Series Switches

sFlow technology uses the following two sampling mechanisms:

- **Packet-based sampling:** Samples one packet out of a specified number of packets from an interface enabled for sFlow technology.
- **Time-based sampling:** Samples interface statistics at a specified interval from an interface enabled for sFlow technology.

The sampling information is used to create a network traffic visibility picture. The Juniper Networks Junos operating system (Junos OS) fully supports the sFlow standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.



**NOTE:** sFlow technology on the switches samples only raw packet headers. A raw Ethernet packet is the complete Layer 2 network frame.

An sFlow monitoring system consists of an sFlow agent embedded in the switch and a centralized collector. The sFlow agent's two main activities are random sampling and statistics gathering. The sFlow agent combines interface counters and flow samples and sends them across the network to the sFlow collector in UDP datagrams, directing those datagrams to the IP address and UDP destination port of the collector. Each datagram contains the following information:

- The IP address of the sFlow agent
- The number of samples
- The interface through which the packets entered the agent

- The interface through which the packets exited the agent
- The source and destination interface for the packets
- The source and destination VLAN for the packets

EX Series switches adopt the distributed sFlow architecture. The sFlow agent has two separate sampling entities that are associated with each Packet Forwarding Engine. These sampling entities are known as subagents. Each subagent has a unique ID that is used by the collector to identify the data source. A subagent has its own independent state and forwards its own sample packets to the sFlow agent. The sFlow agent is responsible for packaging the samples into datagrams and sending them to the sFlow collector. Because sampling is distributed across subagents, the protocol overhead associated with sFlow technology is significantly reduced at the collector.



**NOTE:** You cannot configure sFlow monitoring on a link aggregation group (LAG), but you can configure it individually on a LAG member interface.



**NOTE:** If the mastership assignment changes in a Virtual Chassis setup, sFlow technology continues to function.

---

### Adaptive Sampling

The switches use adaptive sampling to ensure both sampling accuracy and efficiency. Adaptive sampling is a process of monitoring the overall incoming traffic rate on the network device and providing intelligent feedback to interfaces to dynamically adapt the sampling rates on interfaces on the basis of traffic conditions. Interfaces on which incoming traffic exceeds the system threshold are checked so that all violations can be regulated without affecting the traffic on other interfaces. Every 12 seconds, the agent checks interfaces to get the number of samples, and interfaces are grouped on the basis of the slot that they belong to. The top five interfaces that produce the highest number of samples are selected. Using the binary backoff algorithm, the sampling load on these interfaces is reduced by half and allotted to interfaces that have a lower sampling rate. Therefore, when the processor's sampling limit is reached, the sampling rate is adapted such that it does not load the processor any further. If the switch is rebooted, the adaptive sampling rate is reset to the user-configured sampling rate. Also, if you modify the sampling rate, the adaptive sampling rate changes.

The advantage of adaptive sampling is that the switch continues to operate at its optimum level even when there is a change in the traffic patterns in the interfaces. You do not need to make any changes. Because the sampling rate adapts dynamically to changing network conditions, the resources are utilized optimally resulting in a high-performance network.

Infrequent sampling flows might not be reported in the sFlow information, but over time, the majority of flows are reported. On the basis of the configured sampling rate  $N$ , 1 out of  $N$  packets is captured and sent to the collector. This type of sampling does not provide a result that is 100 percent accurate in the analysis, but it does provide a result of quantifiable accuracy. A user-configured polling interval defines how often the sFlow

data for a specific interface are sent to the collector, but an sFlow agent can also schedule polling.



**NOTE:** sFlow technology on EX Series switches does not support graceful restart. When a graceful restart occurs, the adaptive sampling rate is set to the user-configured sampling rate.

### sFlow Agent Address Assignment

The sFlow collector uses the sFlow agent's IP address to determine the source of the sFlow data. You can configure the IP address of the sFlow agent to ensure that the agent ID of the sFlow agent remains constant. If you do not configure the IP address of the sFlow agent, an IP address is automatically assigned to the agent. This is the IP address of one of the following interfaces configured on the switch taken in the given order of priority:

1. Virtual management Ethernet (VME) interface
2. Management Ethernet interface

If neither of the preceding interfaces has been configured, the IP address of any Layer 3 interface or the routed VLAN interface (RVI) is assigned to the agent. At least one interface must be configured on the switch for an IP address to be automatically assigned to the agent. When the agent's IP address is assigned automatically, the IP address is dynamic and changes when the switch reboots.

sFlow data can be used to provide network traffic visibility information. You can explicitly configure the IP address to be assigned to source data (sFlow datagrams). If you do not explicitly configure that address, the IP address of the configured Gigabit Ethernet interface, 10-Gigabit Ethernet interface, or the RVI is used as the source IP address.

#### Related Documentation

- [Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995](#)
- [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\) on page 4049](#)
- [Monitoring Interface Status and Traffic on page 2835](#)

## Ethernet OAM Link Fault Management and Connectivity Fault Management

- [Understanding Ethernet OAM Link Fault Management for an EX Series Switch on page 3972](#)
- [Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch on page 3973](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 3974](#)

## Understanding Ethernet OAM Link Fault Management for an EX Series Switch

Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities even as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning. Ethernet OAM should:

- Rely only on the media access control (MAC) address or virtual LAN identifier for troubleshooting.
- Work independently of the actual Ethernet transport and function over physical Ethernet ports or a virtual service such as pseudowire.
- Isolate faults over a flat (or single operator) network architecture or nested or hierarchical (or multiprovider) networks.

The following OAM LFM features are supported on EX Series switches:

- Discovery and Link Monitoring

The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard. You can specify the discovery mode used for IEEE 802.3ah OAM support. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery. The switch performs link monitoring by sending periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

- Remote Fault Detection

Remote fault detection uses flags and events. Flags are used to convey the following: Link Fault means a loss of signal, Dying Gasp means an unrecoverable condition such as a power failure, and Critical Event means an unspecified vendor-specific critical event. You can specify the periodic OAM PDU sending interval for fault detection. The EX Series switch uses the Event Notification OAM PDU to notify the remote OAM device when a problem is detected. You can specify the action to be taken by the system when the configured link-fault event occurs.

- Remote Loopback Mode

Remote loopback mode ensures link quality between the switch and a remote peer during installation or troubleshooting. In this mode, when the interface receives a frame that is not an OAM PDU or a pause frame, it sends it back on the same interface on which it was received. The link appears to be in the active state. You can use the returned loopback acknowledgement to test delay, jitter, and throughput.

Junos OS can place a remote DTE into loopback mode (if remote loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent and processed.

**Related Documentation**

- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 4062](#)
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027](#)

## Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch

Ethernet interfaces on Juniper Networks EX Series Ethernet Switches and Juniper Networks Junos operating system (Junos OS) for EX Series switches support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM). CFM monitors Ethernet networks that might comprise one or more service instances for network-compromising connectivity faults.

The major features of CFM are:

- Fault monitoring using the continuity check protocol. This is a neighbor discovery and health check protocol that discovers and maintains adjacencies at the VLAN or link level.
- Path discovery and fault verification using the linktrace protocol.
- Fault isolation using the loopback protocol.

CFM partitions the service network into various administrative domains. For example, operators, providers, and customers might be part of different administrative domains. Each administrative domain is mapped into one maintenance domain providing enough information to perform its own management, thus avoiding security breaches and making end-to-end monitoring possible.

In a CFM maintenance domain, each service instance is called a maintenance association. A maintenance association can be thought of as a full mesh of maintenance association endpoints (MEPs) having similar characteristics. MEPs are active CFM entities generating and responding to CFM protocol messages. There is also a maintenance intermediate point (MIP), which is a CFM entity similar to the MEP, but more passive (MIPs only respond to CFM messages).

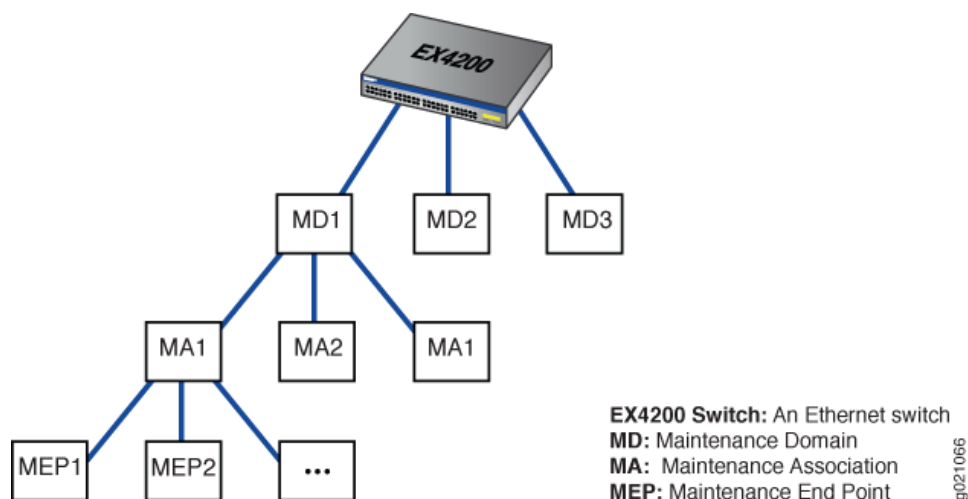
Each maintenance domain is associated with a maintenance domain level from 0 through 7. Level allocation is based on the network hierarchy, where outer domains are assigned

a higher level than the inner domains. Configure customer end points to have the highest maintenance domain level. The maintenance domain level is a mandatory parameter that indicates the nesting relationships between various maintenance domains. The level is embedded in each CFM frame. CFM messages within a given level are processed by MEPs at that same level.

To enable CFM on an Ethernet interface, you must configure maintenance domains, maintenance associations, and maintenance association end points (MEPs).

Figure 52 on page 3974 shows the relationships among maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs) configured on a switch.

**Figure 52: Relationship Among MEPs, MIPs, and Maintenance Domain Levels**



**Related Documentation**

- [Configuring Ethernet OAM Connectivity Fault Management \(CLI Procedure\) on page 4058](#)
- [Junos OS Network Interfaces Configuration Guide](#)

## Understanding Ethernet Frame Delay Measurements on Switches

Performance management depends on the accurate measurement of service-level agreement (SLA) objective parameters, which can include bandwidth and reliability. In many cases, a service provider could be subject to penalties imposed by regulation, statute, or contract if network performance is not within the bounds established for the service. One key performance objective is delay, along with its close relative, delay variation (often called jitter). Some applications (such as bulk file transfer) will function just as well with high delays across the network and high delay variations, while other applications (such as voice) can function only with low and stable delays. Many networks invoke protocols or features available at Layer 3 (the packet layer) or higher to measure network delays and jitter link by link. However, when the network consists of many Ethernet links, there are few protocols and features available at Layer 2 (the frame layer) that allow routers and switches to measure frame delay and jitter. This is where the ability to configure and monitor Ethernet frame delay is helpful.

This topic includes:

- [Ethernet Frame Delay Measurements on page 3975](#)
- [Types of Ethernet Frame Delay Measurements on page 3975](#)
- [Limitations on page 3976](#)

## Ethernet Frame Delay Measurements

---

You can perform Ethernet frame delay measurements (referred to as ETH-DM in Ethernet specifications) on Juniper Networks EX Series Ethernet Switches. This feature allows you to configure on-demand Operation, Administration, and Maintenance (OAM) statements for the measurement of frame delay and frame delay variation (jitter). You can configure Ethernet frame delay measurement in either one-way or two-way (round-trip) mode to gather frame delay statistics simultaneously from multiple sessions. Ethernet frame delay measurement provides fine control to operators for triggering delay measurement on a given service and can be used to monitor SLAs.

Ethernet frame delay measurement also collects other useful information, such as worst and best case delays, average delay, and average delay variation. It supports software-assisted timestamping in the receive direction for delay measurements. It also provides runtime display of delay statistics when two-way delay measurement is triggered. Ethernet frame delay measurement records the last 100 samples collected per remote maintenance association end point (MEP) or per connectivity fault management (CFM) session. You can retrieve the history at any time using simple commands. You can clear all Ethernet frame delay measurement statistics and PDU counters. Ethernet frame delay measurement is fully compliant with the ITU-T Y.1731 (*OAM Functions and Mechanisms for Ethernet-based Networks*) specification.

Ethernet frame delay measurement uses the IEEE 802.1ag CFM infrastructure.

Generally, Ethernet frame delay measurements are made in a peer fashion from one MEP or CFM session to another. However, these measurements are not made to maintenance association intermediate points (MIPs).

For a complete description of Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.

## Types of Ethernet Frame Delay Measurements

---

There are two types of Ethernet frame delay measurements:

- One-way
- Two-way (round-trip)

For one-way Ethernet frame delay measurement, either MEP can send a request to begin a one-way delay measurement to its peer MEP. However, the statistics are collected only at the receiver MEP. This feature requires the clocks at the transmitting and receiving MEPs to be synchronized. If these clocks fall out of synchronization, only one-way delay variation and average delay variation values are computed correctly (and will, therefore, be valid). Use the **show** commands at the receiver MEP to display one-way delay statistics.

For two-way (round-trip) Ethernet frame delay measurement, either MEP can send a request to begin a two-way delay measurement to its peer MEP, which responds with timestamp information. Run-time statistics are collected and displayed at the initiator MEP. The clocks do not need to be synchronized at the transmitting and receiving MEPs. Junos OS supports timestamps in delay measurement reply (DMR) frames to increase the accuracy of delay calculations.

Use the **show** commands at the initiator MEP to display two-way delay statistics, and at the receiver MEP to display one-way delay statistics.

You can create an iterator profile to periodically transmit SLA measurement packets in the form of ITU-Y.1731-compliant frames for delay measurement or loss measurement.

---

### Limitations

The following are some limitations with regard to using Ethernet frame delay measurement:

- Ethernet frame delay measurements are available only when distributed periodic packet management (PPM) is enabled.
- The statistics collected are lost after a graceful Routing Engine switchover (GRES).
- You can monitor only one session to the same remote MEP or MAC address.
- Accuracy is compromised when the system configuration changes (such as from reconfiguration). We recommend performing Ethernet frame delay measurements on a stable system.

#### Related Documentation

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 4065](#)
- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4066](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4069](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 4068](#)

---

## Uplink Failure Detection

- [Understanding Uplink Failure Detection on page 3977](#)



## Understanding Uplink Failure Detection

Uplink failure detection allows Juniper Networks EX Series Ethernet Switches to detect link failure on uplink interfaces and to propagate the failure to the downlink interfaces so that servers connected to those downlink interfaces can switch over to secondary interfaces.

Uplink failure detection supports network adapter teaming and provides network redundancy. In network adapter teaming, all the network interface cards (NICs) on a server are configured in a primary or secondary relationship and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link. With uplink failure detection, the switch monitors uplink interfaces for link failures. When it detects a failure, it disables the downlink interfaces. When the server detects disabled downlink interfaces, it switches over to the secondary link to help ensure balanced traffic flow on switches.

This topic describes:


- [Uplink Failure Detection Overview on page 3977](#)
- [Failure Detection Pair on page 3978](#)

### Uplink Failure Detection Overview

Uplink failure detection allows switches to monitor uplink interfaces to spot link failures. When a switch detects a link failure, it automatically disables the downlink interfaces in that group. The server that is connected to the disabled downlink interfaces triggers a network-adapter failover to a secondary link to avoid any information drop.

[Figure 53 on page 3977](#) illustrates a typical setup for uplink failure detection.

**Figure 53: Uplink Failure Detection Configuration on Switches**



g040577

For uplink failure detection, you specify a group of uplink interfaces to be monitored and downlink interfaces to be brought down when an uplink fails. The downlink interfaces are bound to the uplink interfaces within the group. If all uplink interfaces in a group go down, then the switch brings down all downlink interfaces within that group. If any uplink interface returns to service, then the switch brings all downlink interfaces in that group back to service.



**NOTE:** Routed VLAN interfaces (RVIs) cannot be configured as uplink interfaces to be monitored.

The switch can monitor both physical-interface links and logical-interface links for uplink failures, but you must put the two types of interfaces in separate groups.



**NOTE:** To detect failure of logical interfaces, the server must run some high level protocol such as keepalives between the switch and the server.

---

### Failure Detection Pair

Uplink failure detection requires that you create groups that contain uplink interfaces and downlink interfaces. Each group includes one of each of the following:

- A link-to-monitor interface—The link-to-monitor interfaces specify the uplink interfaces the switch monitors. You can configure a maximum of 48 uplink interfaces as link-to-monitor in a group.
- A link-to-disable interface—The link-to-disable interfaces specify the downlink interfaces the switch disables when the switch detects an uplink failure. You can configure a maximum of 48 downlink interfaces as link-to-disable in a group.

The link-to-disable interfaces are bound to the link-to-monitor interfaces within the group. When a link-to-monitor interface returns to service, the switch automatically enables all link-to-disable interfaces in the group.

#### Related Documentation

- [Configuring Interfaces for Uplink Failure Detection \(CLI Procedure\) on page 4064](#)

---

## Network Analytics

- [Network Analytics Overview on page 3978](#)
- [Understanding Enhanced Network Analytics Streaming Data on page 3985](#)
- [Understanding Enhanced Analytics Local File Output on page 3991](#)
- [Prototype File for the Google Protocol Buffer Stream Format on page 3993](#)

### Network Analytics Overview

The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the

data using sophisticated algorithms, and captures the results in reports. Network administrators can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed. The analytics manager (analyticsm) in the Packet Forwarding Engine collects traffic and queue statistics, and the analytics daemon (analyticsd) in the Routing Engine analyzes the data and generates reports. You can enable network analytics by configuring microburst monitoring and high-frequency traffic statistics monitoring.



**NOTE:** In Junos OS Release 13.2X51-D15, the network analytics feature was enhanced, and extensive changes were made to the CLI statements and hierarchies. If you upgrade to Junos OS Release 13.2X51-D15 or later from a release prior to 13.2X51-D15, network analytics configurations committed in previous releases will appear on your device, but the feature is disabled. To enable this feature, you must reconfigure it using the new CLI statements and hierarchies.

For more information, see:

- [Analytics Feature Overview on page 3979](#)
- [Network Analytics Enhancements Overview on page 3980](#)
- [Summary of CLI Changes on page 3981](#)

### Analytics Feature Overview

You enable network analytics by configuring queue (microburst) monitoring and high-frequency traffic statistics monitoring. You use microburst monitoring to look at traffic queue conditions in the network. A microburst occurrence indicates to the Packet Forwarding Engine that a user-specified queue depth or latency threshold is reached. The queue depth is the buffer (in bytes) containing the data, and latency is the time (in nanoseconds or microseconds) the data stays in the queue.

You can configure queue monitoring based on either queue depth or latency (but not both), and configure the frequency (polling interval) at which the Packet Forwarding Engine checks for microbursts and sends the data to the Routing Engine for processing. You may configure queue monitoring globally for all physical interfaces on the system, or for a specific interface on the switch. However, the specified queue monitoring interval applies either to all interfaces, or none; you cannot configure the interval for each interface.

You use high-frequency traffic statistics monitoring to collect traffic statistics at specified polling intervals. Similar to the queue monitoring interval, the traffic monitoring interval applies either to all interfaces, or none; you cannot configure the interval for each interface.

Both traffic and queue monitoring are disabled by default. You must configure each type of monitoring using the CLI. In each case, the configuration for an interface always takes precedence over the global configuration.



**NOTE:** You can configure traffic and queue monitoring for physical interfaces only; logical interfaces and Virtual Chassis port (VCP) interfaces are not supported.

The analyticsd daemon in the Routing Engine generates local log files containing queue and traffic statistics records. You can specify the log filename and size, and the number of log files. If you do not configure a filename, the data is not saved.

You can display the local log file or specify a server to receive the streaming data containing the queue and traffic statistics.

For each port, information for the last 10 records of traffic statistics and 100 records of queue statistics is cached. You may view this information by using the **show analytics** commands.

To store traceoptions data, you configure the **traceoptions** statement at the **[edit services analytics]** hierarchy level.

### Network Analytics Enhancements Overview

---

Beginning in Junos OS Release 13.2X51-D15, the network analytics feature provides the following enhancements:

- **Resources**—Consist of interfaces and system. The interfaces resource allows you to configure an interface name and an associated resource profile name for each interface. With the system resource, you can configure the polling intervals for queue monitoring and traffic monitoring, and an associated resource profile for the system.
- **Resource profile**—A template that contains the configurations for queue and traffic monitoring, such as depth threshold and latency threshold values, and whether each type of monitoring is enabled or disabled. Once a resource profile is configured, you apply it to a system or interfaces resource.
- **Collector**—A server for collecting queue and traffic monitoring statistics, and can be a local or remote server. You can configure a local server to store monitoring statistics in a log file, or a remote server to receive streamed statistics data.
- **Export profile**—You must configure an export profile if you wish to send streaming data to a remote collector. In the export profile, you define the category of streamed data (system-wide or interface-specific) to determine stream type the collector will receive. You can specify both system and interface stream categories. System data includes system information and status of queue and traffic monitoring. Interface-specific data includes interface information, queue and traffic statistics, and link, queue, and traffic status.
- **Google Protocol Buffer (GBP) stream format**—A new streaming format for monitoring statistics data that is sent to a remote collector in a single AnRecord message. This stream format provides nine types of information, including:
  - **System information**—General system information, including boot time, model information, serial number, number of ports, and so on.
  - **System queue status**—Queue status for the system in general.
  - **System traffic status**—Traffic status for the system in general.
  - **Interface information**—Includes SNMP index, slot, port, and other information.
  - **Queue statistics for interfaces**—Queue statistics for specific interfaces.

- Traffic statistics for interfaces—Traffic statistics for specific interfaces.
- Link status for interfaces—Includes link speed, state, and so on.
- Queue status for interfaces—Queue status for specific interfaces.
- Traffic status for interfaces—Traffic status for specific interfaces.
- The **analytics.proto** file—Provides a template for the GBP stream format. This file can be used for writing your analytics server application. To download the file, go to:  
[http://www.juniper.net/techpubs/en\\_US/junos13.2/topics/reference/proto-files/analytics-proto.txt](http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/proto-files/analytics-proto.txt)
- Use of threshold values—The Analytics Manager (analyticsm) will generate a queue statistics record when the lower queue depth or latency threshold value is exceeded.
- User Datagram Protocol (UDP)—Additional transport protocol you can configure, in addition to Transmission Control Protocol (TCP), for the remote streaming server port.
- Single file for local logging—Replaces the separate log files for queue and traffic statistics.
- Change in latency measurement—Configuration and reporting of latency values have changed from microseconds to nanoseconds.
- Change in reporting of the collection time in UTC format—Statistics collection time is reported in microseconds instead of milliseconds.
- New operational mode command **show analytics collector**—Replaces the **show analytics streaming-server** command.
- Changes in command output format—Include the following changes:
  - Addition of unicast, multicast, and broadcast packet counters in queue and traffic statistics.
  - Reversal of the sequence of statistics information in the output. The most recent record is displayed at the beginning, and the oldest record at the end of the output.
  - Removal of traffic or queue monitoring status information from the global portion of the **show analytics configuration** and **show analytics status** command output if there is no global configuration.
  - Addition of **n/a** to the interface-specific portion of the **show analytics configuration** and **show analytics status** command output if a parameter is not configured (for example, depth threshold or latency threshold).

### Summary of CLI Changes

Beginning in Junos OS Release 13.2X51-D15, enhancements to the network analytics feature result in changes in the CLI when you configure the feature. See [Table 421 on page 3982](#) for a summary of CLI changes.

Table 421: Network Analytics CLI Changes

| Task                                                                                                               | CLI for Junos OS Release 13.2X50-D15 and 13.2X51-D10                                                                                                                                                                    | CLI for Junos OS Release 13.2X51-D15 and later                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configuring global queue and traffic monitoring polling interval                                                   | <pre>[edit services analytics]  traffic-statistics {   interval <i>interval</i>; } queue-statistics {   interval <i>interval</i>; }</pre>                                                                               | <pre>[edit services analytics]  resource {   system {     polling-interval {       queue-monitoring <i>interval</i>;       traffic-monitoring <i>interval</i>;     }   } }</pre>                                                                                                                                                                                                                                                                                                                                                         |
| Configuring local files for traffic and queue statistics reporting                                                 | <pre>[edit services analytics]  traffic-statistics {   file <i>filename</i>;   size <i>size</i>;   files <i>number</i>; } queue-statistics {   file <i>filename</i>;   size <i>size</i>;   files <i>number</i>; }</pre> | <pre>[edit services analytics]  collector {   local {     file <i>filename</i> {       files <i>number</i>;       size <i>size</i>;     }   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                    |
| Enabling queue statistics and traffic monitoring, and specifying the depth threshold for all interfaces (globally) | <pre>[edit services analytics]  interfaces {   all {     queue-statistics;     traffic-statistics;     depth-threshold {       high <i>number</i>;       low <i>number</i>;     }   } }</pre>                           | <p>Requires defining a resource profile and applying it to the system:</p> <ol style="list-style-type: none"> <li>To define a resource profile: <pre>[edit services analytics]  resource-profiles {   <i>profile-name</i> {     queue-monitoring;     traffic-monitoring;     depth-threshold {       high <i>number</i>;       low <i>number</i>;     }   } }</pre> </li> <li>To apply a profile to the system: <pre>[edit services analytics]  resource {   system {     resource-profile <i>profile-name</i>;   } }</pre> </li> </ol> |

Table 421: Network Analytics CLI Changes (*continued*)

| Task                                                                                                                                                                                                                                            | CLI for Junos OS Release 13.2X50-D15 and 13.2X51-D10                                                                                                                                         | CLI for Junos OS Release 13.2X51-D15 and later                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enabling queue statistics and traffic monitoring, and specifying the latency threshold for one interface                                                                                                                                        | <pre>[edit services analytics] interfaces {   interface {     queue-statistics;     traffic-statistics;     latency-threshold       high <i>number</i>;       low <i>number</i>;   } }</pre> | <p>Requires defining a resource profile and applying it to the interface:</p> <ol style="list-style-type: none"> <li>To define a resource profile: <pre>[edit services analytics] resource-profiles {   profile-name {     queue-monitoring;     traffic-monitoring;     latency-threshold {       high <i>number</i>;       low <i>number</i>;     }   } }</pre> </li> <li>To apply a profile to the interface: <pre>[edit services analytics] resource {   interfaces {     interface-name {       resource-profile <i>profile-name</i>;     }   } }</pre> </li> </ol>      |
| <p>Configuring the streaming data format (JSON, CSV, or TSV) to send to a remote server</p> <p><b>NOTE:</b> Junos OS Release 13.2X51-D15 added support for the GPB stream format and configuration of the transport protocols (TCP or UDP).</p> | <pre>[edit services analytics] streaming-servers {   address <i>ip-address</i> {     port <i>number</i> {       stream-format <i>format</i>;     }   } }</pre>                               | <p>Requires defining the stream format in an export profile and applying the profile to the collector.</p> <ol style="list-style-type: none"> <li>To configure the stream format: <pre>[edit services analytics] export-profiles {   profile-name {     stream-format <i>format</i>;   } }</pre> </li> <li>To apply an export profile to the collector: <pre>[edit services analytics] collector {   address <i>ip-address</i> {     port <i>number</i> {       transport <i>protocol</i> {         export-profile <i>profile-name</i>;       }     }   } }</pre> </li> </ol> |

Table 421: Network Analytics CLI Changes (*continued*)

| Task                                                                                             | CLI for Junos OS Release 13.2X50-D15 and 13.2X51-D10                                                                                                                                        | CLI for Junos OS Release 13.2X51-D15 and later                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configuring the streaming message types (queue or traffic statistics) to send to a remote server | <pre> [edit services analytics] streaming-servers {   address <i>ip-address</i> {     port <i>number</i> {       stream-type <i>type</i>;       stream-type <i>type</i>;     }   } } </pre> | <p>Requires defining an export profile and applying it to the collector:</p> <ol style="list-style-type: none"> <li>To define an export profile: <pre> [edit services analytics] export-profiles {   <i>profile-name</i> {     interface {       information;       statistics {         queue;         traffic;       }       status {         link;         queue;         traffic;       }     }     system {       information;       status {         queue;         traffic;       }     }   } } </pre> </li> <li>To apply an export profile to the collector: <pre> [edit services analytics] collector {   address <i>ip-address</i> {     port <i>number</i> {       export-profile <i>profile-name</i>;     }   } } </pre> </li> </ol> |



Table 421: Network Analytics CLI Changes (*continued*)

| Task                                                                                | CLI for Junos OS Release 13.2X50-D15 and 13.2X51-D10               | CLI for Junos OS Release 13.2X51-D15 and later                                                                                                                                                                                                                                                                                |
|-------------------------------------------------------------------------------------|--------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configuring the transport protocol for sending streaming data to an external server | No configuration is available. Only the TCP protocol is supported. | Configuration is available. Both TCP and UDP protocols are supported, and can be configured for the same port.<br><br>[edit services analytics]<br><br>collector {<br>address <i>ip-address</i> {<br>port <i>number1</i> {<br>transport tcp;<br>transport udp;<br>}<br>port <i>number2</i> {<br>transport udp;<br>}<br>}<br>} |
| Show information about remote streaming server or collector                         | Issue the <b>show analytics streaming-sever</b> command.           | Issue the <b>show analytics collector</b> command.                                                                                                                                                                                                                                                                            |

**Related Documentation**

- [analytics on page 4167](#)

## Understanding Enhanced Network Analytics Streaming Data

Network analytics monitoring data can be streamed to remote servers called collectors. You can configure one or more collectors to receive streamed data containing queue and traffic statistics. This topic describes the streamed data output.



**NOTE:** This topic applies to Junos OS Release 13.2X51-D15 or later.

Starting in Junos OS Release 13.2X51-D15, network analytics supports the following streaming data formats and output:

- [Google Protocol Buffer \(GPB\) on page 3985](#)
- [JavaScript Object Notation \(JSON\) on page 3988](#)
- [Comma-separated Values \(CSV\) on page 3988](#)
- [Tab-separated Values \(TSV\) on page 3988](#)
- [Queue Statistics Output for JSON, CSV, and TSV on page 3989](#)
- [Traffic Statistics Output for JSON, CSV, and TSV on page 3989](#)

### Google Protocol Buffer (GPB)

Support for the Google Protocol Buffer (GPB) streaming format has been added in Junos OS Release 13.2X51-D15. This streaming format provides:

- Support for nine types of messages, based on resource type (system-wide or interface-specific).
- Sends messages in a hierarchical format.
- You can generate other stream format messages (JSON, CSV, TSV) from GPB formatted messages.
- Includes a 8-byte message header. See [Table 422 on page 3986](#) for more information.

[Table 422 on page 3986](#) describes the GPB stream format message header.

**Table 422: GPB Stream Format Message Header Information**

| Byte Position | Field                   |
|---------------|-------------------------|
| 0 to 3        | Length of message       |
| 4             | Message version         |
| 5 to 7        | Reserved for future use |

The following GPB prototype file (**analytics.proto**) provides details about the streamed data:

```
package analytics;

// Traffic statistics related info
message TrafficStatus {
    optional uint32      status          = 1;
    optional uint32      poll_interval   = 2;
}

// Queue statistics related info
message QueueStatus {
    optional uint32      status          = 1;
    optional uint32      poll_interval   = 2;
    optional uint64      lt_high         = 3;
    optional uint64      lt_low          = 4;
    optional uint64      dt_high         = 5;
    optional uint64      dt_low          = 6;
}

message LinkStatus {
    optional uint64      speed           = 1;
    optional uint32      duplex          = 2;
    optional uint32      mtu             = 3;
    optional bool        state           = 4;
    optional bool        auto_negotiation= 5;
}

message InterfaceInfo {
    optional uint32      snmp_index      = 1;
    optional uint32      index           = 2;
    optional uint32      slot            = 3;
    optional uint32      port            = 4;
    optional uint32      media_type      = 5;
    optional uint32      capability      = 6;
    optional uint32      porttype        = 7;
```

```

}

message InterfaceStatus {
    optional LinkStatus      link          = 1;
    optional QueueStatus     queue_status  = 2;
    optional TrafficStatus   traffic_status = 3;
}

message QueueStats {
    optional uint64          timestamp     = 1;
    optional uint64          queue_depth   = 2;
    optional uint64          latency       = 3;
}

message TrafficStats {
    optional uint64          timestamp     = 1;
    optional uint64          rxpkt         = 2;
    optional uint64          rxucpkt       = 3;
    optional uint64          rxmcpkt       = 4;
    optional uint64          rxbcpkt       = 5;
    optional uint64          rxpps         = 6;
    optional uint64          rxbyte        = 7;
    optional uint64          rxbps         = 8;
    optional uint64          rxrcerr       = 9;
    optional uint64          rxdropkt      = 10;
    optional uint64          txpkt         = 11;
    optional uint64          txucpkt       = 12;
    optional uint64          txmcpkt       = 13;
    optional uint64          txbcpkt       = 14;
    optional uint64          txpps         = 15;
    optional uint64          txbyte        = 16;
    optional uint64          txbps         = 17;
    optional uint64          txrcerr       = 18;
    optional uint64          txdropkt      = 19;
}

message InterfaceStats {
    optional TrafficStats     traffic_stats = 1;
    optional QueueStats       queue_stats  = 2;
}

//Interface message
message Interface {
    required string           name          = 1;
    optional bool             deleted       = 2;
    optional InterfaceInfo    information   = 3;
    optional InterfaceStats    stats        = 4;
    optional InterfaceStatus   status       = 5;
}

message SystemInfo {
    optional uint64           boot_time     = 1;
    optional string           model_info    = 2;
    optional string           serial_no     = 3;
    optional uint32           max_ports     = 4;
    optional string           collector     = 5;
    repeated string           interface_list = 6;
}

message SystemStatus {
    optional QueueStatus       queue_status = 1;

```

```

        optional TrafficStatus      traffic_status = 2;
    }

    //System message
    message System {
        required string              name           = 1;
        optional bool                deleted        = 2;
        optional SystemInfo          information    = 3;
        optional SystemStatus        status        = 4;
    }

    message AnRecord {
        optional uint64              timestamp     = 1;
        optional System              system        = 2;
        repeated Interface           interface     = 3;
    }

```

### JavaScript Object Notation (JSON)

The JavaScript Object Notation (JSON) streaming format supports the following data:

- Queue statistics data. For example:

```

{"record-type":"queue-stats","time":1383453988263,"router-id":"qfx5100-switch",
"port":"xe-0/0/18","latency":0,"queue-depth":208}

```

See [Table 423 on page 3989](#) for more information about queue statistics output fields.

- Traffic statistics. For example:

```

{"record-type":"traffic-stats","time":1383453986763,"router-id":"qfx5100-switch",
"port":"xe-0/0/16","rxpkt":26524223621,"rxpps":8399588,"rxbyte":3395100629632,
"rxbps":423997832,"rxdrop":0,"rxerr":0,"txpkt":795746503,"txpps":0,"txbyte":101855533467,
"txbps":0,"txdrop":0,"txerr":0}

```

See [Table 424 on page 3989](#) for more information about traffic statistics output fields.

### Comma-separated Values (CSV)

The Comma-separated Values (CSV) streaming format supports the following data:

- Queue statistics. For example:

```

q,1383454067604,qfx5100-switch,xe-0/0/18,0,208

```

See [Table 423 on page 3989](#) for more information about queue statistics output fields.

- Traffic statistics. For example:

```

t,1383454072924,qfx5100-switch,xe-0/0/19,1274299748,82950,163110341556,85603312,0,0,
27254178291,8300088,3488534810679,600002408,27268587050,3490379142400

```

See [Table 424 on page 3989](#) for more information about traffic statistics output fields.

### Tab-separated Values (TSV)

The Tab-separated Values (TSV) streaming format supports the following data:

- Queue statistics. For example:

```

q      585870192561703872      qfx5100-switch      xe-0/0/18      (null)
208    2

```

See [Table 423 on page 3989](#) for more information about queue statistics output fields.

- Traffic statistics. For example:

```
t          1383454139025  qfx5100-switch  xe-0/0/19      1279874033      82022
163823850036  84801488        0              0              27811618258      8199630
3559887126455  919998736       27827356915    3561901685120
```

See [Table 424 on page 3989](#) for more information about traffic statistics output fields.

### Queue Statistics Output for JSON, CSV, and TSV

[Table 423 on page 3989](#) describes the output fields for streamed queue statistics data in the order they appear.

**Table 423: Streamed Queue Statistics Data Output Fields**

| Field       | Description                                                                                                                                                    |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| record-type | Type of statistics. Displayed as: <ul style="list-style-type: none"> <li>• <b>queue-stats</b> (JSON format)</li> <li>• <b>q</b> (CSV or TSV format)</li> </ul> |
| time        | Time (in Unix epoch format) at which the statistics were captured.                                                                                             |
| router-id   | ID of the network analytics host device.                                                                                                                       |
| port        | Name of the physical port configured for network analytics.                                                                                                    |
| latency     | Traffic queue latency in milliseconds.                                                                                                                         |
| queue depth | Depth of the traffic queue in bytes.                                                                                                                           |

### Traffic Statistics Output for JSON, CSV, and TSV

[Table 424 on page 3989](#) describes the output fields for streamed traffic statistics data in the order they appear.

**Table 424: Streamed Traffic Statistics Data Output Fields**

| Field       | Description                                                                                                                                                      |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| record-type | Type of statistics. Displayed as: <ul style="list-style-type: none"> <li>• <b>traffic-stats</b> (JSON format)</li> <li>• <b>t</b> (CSV or TSV format)</li> </ul> |
| time        | Time (in Unix epoch format) at which the statistics were captured.                                                                                               |
| router-id   | ID of the network analytics host device.                                                                                                                         |
| port        | Name of the physical port configured for network analytics.                                                                                                      |
| rxpkt       | Total packets received.                                                                                                                                          |

Table 424: Streamed Traffic Statistics Data Output Fields (*continued*)

| Field  | Description                                      |
|--------|--------------------------------------------------|
| rxpps  | Total packets received per second.               |
| rxbyte | Total bytes received.                            |
| rxbps  | Total bytes received per second.                 |
| rxdrop | Total incoming packets dropped.                  |
| rxerr  | Total packets with errors.                       |
| txpkt  | Total packets transmitted.                       |
| txpps  | Total packets transmitted per second.            |
| txbyte | Total bytes transmitted.                         |
| txbps  | Total bytes transmitted per second.              |
| txdrop | Total transmitted bytes dropped.                 |
| txerr  | Total transmitted packets with errors (dropped). |

**Related  
Documentation**

- [Network Analytics Overview on page 3978](#)
- [Prototype File for the Google Protocol Buffer Stream Format on page 3993](#)
- [address \(Analytics Collector\) on page 4166](#)
- [collector \(Analytics\) on page 4171](#)
- [show analytics collector on page 4346](#)

## Understanding Enhanced Analytics Local File Output

The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. You enable network analytics by configuring queue or traffic statistics monitoring, or both. In addition, you can configure a local file for storing the traffic and queue statistics records.



**NOTE:** This topic describes the local file output in Junos OS Release 13.2X51-D15 and later. For information about local file output from earlier releases, see the [monitor start \(Analytics\)](#) topic.

Beginning in Junos OS Release 13.2X51-D15, the traffic and queue monitoring statistics can be stored locally in a single file. The following example shows the output from the **monitor start** command.

```
root@qfx5100-33> monitor start an
root@qfx5100-33>
*** an ***
q,1393947567698432,qfx5100-33,xe-0/0/19,1098572,1373216
q,1393947568702418,qfx5100-33,xe-0/0/19,1094912,1368640
q,1393947569703415,qfx5100-33,xe-0/0/19,1103065,1378832
t,1393947569874528,qfx5100-33,xe-0/0/16,12603371884,12603371884,0,0,
8426023,1613231610488,8628248712,0,3,5916761,5916761,0,0,0,757345408,0,0,0
t,1393947569874528,qfx5100-33,xe-0/0/18,12601953614,12601953614,0,0,
8446737,1613050071660,8649421552,0,5,131761619,131761619,0,0,84468,
16865487232,86495888,0,0
t,1393947569874528,qfx5100-33,xe-0/0/19,126009250,126009250,0,0,84469,
16129184128,86496392,0,0,12584980342,12584980342,0,0,8446866,1610877487744,
8649588432,12593703960,0
q,1393947575698402,qfx5100-33,xe-0/0/19,1102233,1377792
q,1393947576701398,qfx5100-33,xe-0/0/19,1107724,1384656
```

See [Table 425 on page 3991](#) for queue statistics output, and [Table 426 on page 3992](#) for traffic statistics output. The fields in the tables are listed in the order they appear in the output example.

**Table 425: Output Fields for Queue Statistics in Local Analytics File**

| Field                 | Description                                                                      | Example in Output       |
|-----------------------|----------------------------------------------------------------------------------|-------------------------|
| Record type           | Type of statistics (queue or traffic monitoring)                                 | <b>q</b>                |
| Time (microseconds)   | Unix epoch (or Unix time) in microseconds at which the statistics were captured. | <b>1393947567698432</b> |
| Router ID             | ID of the network analytics host device.                                         | <b>qfx5100-33</b>       |
| Port                  | Name of the physical port configured for network analytics.                      | <b>xe-0/0/19</b>        |
| Latency (nanoseconds) | Traffic queue latency in nanoseconds.                                            | <b>1098572</b>          |
| Queue depth (bytes)   | Depth of the traffic queue in bytes.                                             | <b>1373216</b>          |

Table 426: Output Fields for Traffic Statistics in Local Analytics File

| Field               | Description                                                                      | Example in Output       |
|---------------------|----------------------------------------------------------------------------------|-------------------------|
| Record type         | Type of statistics (queue or traffic monitoring)                                 | <b>t</b>                |
| Time (microseconds) | Unix epoch (or Unix time) in microseconds at which the statistics were captured. | <b>1393947569874528</b> |
| Router ID           | ID of the network analytics host device.                                         | <b>qfx5100-33</b>       |
| Port                | Name of the physical port configured for network analytics.                      | <b>xe-0/0/16</b>        |
| rxpkt               | Total packets received.                                                          | <b>12603371884</b>      |
| rxucpkt             | Total unicast packets received.                                                  | <b>12603371884</b>      |
| rxmcpkt             | Total multicast packets received.                                                | <b>0</b>                |
| rxbcpkt             | Total broadcast packets received.                                                | <b>0</b>                |
| rxpps               | Total packets received per second.                                               | <b>8426023</b>          |
| rxbyte              | Total octets received.                                                           | <b>1613231610488</b>    |
| rxbps               | Total bytes received per second.                                                 | <b>8628248712</b>       |
| rxdroppkt           | Total incoming packets dropped.                                                  | <b>0</b>                |
| rxrcerr             | CRC/Align errors received.                                                       | <b>3</b>                |
| txpkt               | Total packets transmitted.                                                       | <b>5916761</b>          |
| txucpkt             | Total unicast packets transmitted.                                               | <b>5916761</b>          |
| txmcpkt             | Total multicast packets transmitted.                                             | <b>0</b>                |
| txbcpkt             | Total broadcast packets transmitted.                                             | <b>0</b>                |
| txpps               | Total packets transmitted per second.                                            | <b>0</b>                |
| txbyte              | Total octets transmitted.                                                        | <b>757345408</b>        |
| txbps               | Bytes per second transmitted.                                                    | <b>0</b>                |
| txdroppkt           | Total transmitted packets dropped.                                               | <b>0</b>                |
| txrcerr             | CRC/Align errors transmitted.                                                    | <b>0</b>                |



- Related Documentation**
- [Network Analytics Overview on page 3978](#)
  - [analytics on page 4167](#)

## Prototype File for the Google Protocol Buffer Stream Format

The Google Protocol Buffer (GPB) stream format is used for streaming monitoring statistics data to a remote collector in a single AnRecord message.

The **analytics.proto** file provides a template for the GPB stream format. This file can be used for writing your analytics server application.

To download the GPB prototype file, go to:

[http://www.juniper.net/techpubs/en\\_US/junos13.2/topics/reference/proto-files/analytics-proto.txt](http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/proto-files/analytics-proto.txt)

- Related Documentation**
- [Network Analytics Overview on page 3978](#)
  - [analytics on page 4167](#)
  - [export-profiles on page 4173](#)



## CHAPTER 67

# Configuration

- [Configuration Examples on page 3995](#)
- [Configuration Tasks on page 4039](#)
- [Configuration Statements: SNMP on page 4075](#)
- [Configuration Statements: Analyzers and Port Mirroring on page 4155](#)
- [Configuration Statements for Network Analytics on page 4165](#)
- [Configuration Statements: sFlow Technology on page 4184](#)
- [Configuration Statements: Ethernet OAM Connectivity Fault Management on page 4194](#)
- [Configuration Statements: Ethernet OAM Link Fault Management on page 4217](#)
- [Configuration Statements: RPM on page 4236](#)
- [Configuration Statements: Uplink Failure Detection on page 4257](#)

### Configuration Examples

---

- [Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995](#)
- [Example: Configuring Mirroring for Local Monitoring of Employee Resource Use on EX4300 Switches on page 4000](#)
- [Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches on page 4006](#)
- [Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX4300 Switches on page 4016](#)
- [Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches on page 4023](#)
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027](#)
- [Example: Configuring Enhanced Network Analytics Features on page 4029](#)

### Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches

sFlow technology is a networking monitoring technology for high-speed switched or routed networks. It is a technology that is based on statistical sampling. You can configure sFlow technology to continuously monitor traffic at wire speed on all interfaces simultaneously. sFlow data can be used to provide network traffic visibility information.

You can specify sampling rates for ingress and egress packets. Junos OS fully supports the sFlow standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.

This example describes how to configure and use sFlow technology to monitor network traffic.

- [Requirements on page 3996](#)
- [Overview and Topology on page 3996](#)
- [Configuration on page 3997](#)
- [Verification on page 3999](#)

---

### Requirements

This example uses the following hardware and software components:

- One EX Series switch
- Junos OS Release 9.3 or later for EX Series switches

---


### Overview and Topology

sFlow technology samples network packets and sends the samples to a monitoring station. You can specify sampling rates for ingress and egress packets. The information gathered is used to create a network traffic visibility picture.

An sFlow monitoring system consists of an sFlow agent embedded in the switch and a centralized collector. The sFlow agent runs on the switch. It combines interface counters and flow samples and sends them across the network to the sFlow collector.

[Figure 54 on page 3997](#) depicts the basic elements of the sFlow system.

Figure 54: sFlow Technology Monitoring System



### Configuration

To configure sFlow technology, perform the following tasks:

#### CLI Quick Configuration

To quickly configure sFlow technology, copy the following commands and paste them into the switch terminal window:

```
[edit protocols]
set sflow collector 10.204.32.46 udp-port 5600
set sflow interfaces ge-0/0/0
set sflow polling-interval 20
set sflow sample-rate egress 1000
```

**Step-by-Step  
Procedure**

To configure sFlow technology:

1. Configure the IP address and UDP port of the collector:

```
[edit protocols]  
user@switch# set sflow collector 10.204.32.46 udp-port 5600
```



**NOTE:** You can configure a maximum of 4 collectors.

The default UDP port is 6343.

2. Enable sFlow technology on a specific interface:

```
[edit protocols sflow]  
user@switch# set interfaces ge-0/0/0
```



**NOTE:** You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a link aggregation group (LAG) interface, but you can enable it on the member interfaces of a LAG.

3. Specify in seconds how often the sFlow agent polls the interface:

```
[edit protocols sflow]  
user@switch# set polling-interval 20
```



**NOTE:** The polling interval can be specified as a global parameter also. Specify 0 if you do not want to poll the interface.

4. Specify the rate at which egress packets must be sampled:

```
[edit protocols sflow]  
user@switch# set sample-rate egress 1000
```



**NOTE:** You can specify both egress and ingress sampling rates. If you set only the egress sampling rate, the ingress sampling rate will be disabled.



**NOTE:** We recommend that you configure the same sampling rates on all the ports on a line card. If you configure different sampling rates are different, the lowest value is used for all ports. You could still configure different rates on different line cards.

**Results** Check the results of the configuration:

```
[edit protocols sflow]
user@switch# show
polling-interval 20;
sample-rate egress 1000;
collector 10.204.32.46 {
  udp-port 5600;
}
interfaces ge-0/0/0.0;
```

### Verification

To confirm that the configuration is correct, perform these tasks:

- [Verifying That sFlow Technology Is Configured Properly on page 3999](#)
- [Verifying That sFlow Technology Is Enabled on the Specified Interface on page 3999](#)
- [Verifying the sFlow Collector Configuration on page 4000](#)

#### *Verifying That sFlow Technology Is Configured Properly*

**Purpose** Verify that sFlow technology is configured properly.

**Action** Use the **show sflow** command:

```
user@switch> show sflow
sFlow: Enabled
Sample limit: 300 packets/second
Polling interval: 20 seconds
Sample rate egress: 1:1000: Enabled
Sample rate ingress: 1:2048: Disabled
Agent ID: 10.204.96.222
```



**NOTE:** The sampling limit cannot be configured and is set to 300 packets/second per FPC.

**Meaning** The output shows that sFlow technology is enabled and specifies the values for the sampling limit, polling interval, and the egress sampling rate.

#### *Verifying That sFlow Technology Is Enabled on the Specified Interface*

**Purpose** Verify that sFlow technology is enabled on the specified interfaces and display the sampling parameters.

**Action** Use the **show sflow interface** command:

```
user@switch> show sflow interface
```

| Interface  | Status           | Sample rate    | Adapted sample rate | Polling-interval |
|------------|------------------|----------------|---------------------|------------------|
|            | Egress Ingress   | Egress Ingress | Egress Ingress      |                  |
| ge-0/0/0.0 | Enabled Disabled | 1000 2048      | 1000 2048           | 20               |

**Meaning** The output indicates that sFlow technology is enabled on the ge-0/0/0.0 interface with an egress sampling rate of 1000, a disabled ingress sampling rate, and a polling interval of 20 seconds.

#### *Verifying the sFlow Collector Configuration*

**Purpose** Verify the sFlow collector's configuration.

**Action** Use the **show sflow collector** command:

```
user@switch> show sflow collector
```

| Collector address | Udp-port | No. of samples |
|-------------------|----------|----------------|
| 10.204.32.46      | 5600     | 1000           |
| 10.204.32.76      | 3400     | 1000           |

**Meaning** The output displays the IP address of the collectors and the UDP ports. It also displays the number of samples.

**Related Documentation**

- [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\) on page 4049](#)
- [Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch on page 3969](#)

### Example: Configuring Mirroring for Local Monitoring of Employee Resource Use on EX4300 Switches



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

EX4300 switches enable you to configure mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN

You can analyze the mirrored traffic by using a protocol analyzer application installed on a system connected to the local destination interface (or running on a remote monitoring station if you are sending mirrored traffic to an analyzer VLAN).



This example describes how to configure local mirroring on an EX4300 switch. This example describes how to configure the switch to mirror traffic entering interfaces connected to employee computers to an analyzer output interface on the same switch.

- [Requirements on page 4001](#)
- [Overview and Topology on page 4001](#)
- [Mirroring All Employee Traffic for Local Analysis on page 4002](#)
- [Mirroring Employee-to-Web Traffic for Local Analysis on page 4003](#)
- [Verification on page 4005](#)

### Requirements

---

This example uses the following hardware and software components:

- One EX4300 switch
- Junos OS Release 13.2X50-D10, or later for EX Series switches

Before you configure mirroring, be sure you have an understanding of mirroring concepts. For information about mirroring, see [“Understanding Port Mirroring and Analyzers on EX4300 Switches” on page 3964](#).

### Overview and Topology

---

This topic includes two examples that describe how to mirror traffic entering ports on the switch to a destination interface on the same switch (local mirroring). The first example shows how to mirror all traffic entering the ports connected to employee computers. The second example shows the same scenario, but includes a filter to mirror only the employee traffic going to the Web.

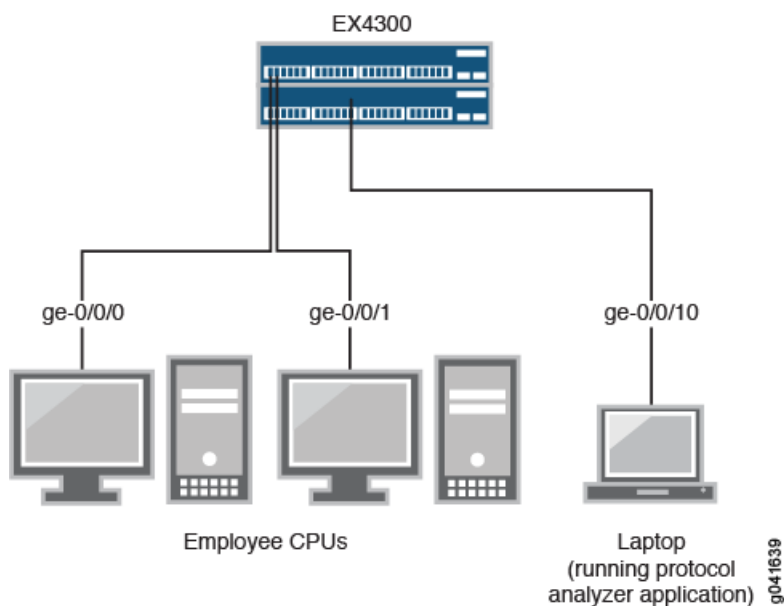
The interfaces ge-0/0/0 and ge-0/0/1 serve as connections for employee computers. The interface ge0/0/10 is reserved for analysis of mirrored traffic. Connect a PC running a protocol analyzer application to the analyzer output interface to analyze the mirrored traffic.



**NOTE:** Multiple ports mirrored to one interface can cause buffer overflow and dropped packets.

Both examples use the network topology shown in [Figure 55 on page 4002](#).

Figure 55: Network Topology for Local Mirroring Example



### Mirroring All Employee Traffic for Local Analysis

To configure mirroring for all employee traffic for local analysis, perform these tasks:

#### CLI Quick Configuration

To quickly configure local mirroring for ingress traffic to the two ports connected to employee computers, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 unit 0 family ethernet-switching
set interfaces ge-0/0/1 unit 0 family inet 192.1.1.1/24
set interfaces ge-0/0/10 unit 0 family ethernet-switching
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/0.0
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/1.0
set forwarding-options analyzer employee-monitor output interface ge-0/0/10.0
```

#### Step-by-Step Procedure

To configure an analyzer called **employee-monitor** and specify the input (source) interfaces and the analyzer output interface:

1. Configure each interface connected to employee computers as an input interface for the analyzer **employee-monitor**:  
  

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
```
2. Configure the output analyzer interface for the analyzer **employee-monitor**. This will be the destination interface for the mirrored packets:  
  

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor output interface ge-0/0/10.0
```

**Results** Check the results of the configuration:

```
[edit]
user@switch> show
```

```

forwarding-options {
  analyzer employee-monitor {
    input {
      ingress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;}
      }
    output {
      interface {
        ge-0/0/10.0;
      }
    }
  }
}

```

### Mirroring Employee-to-Web Traffic for Local Analysis

To configure mirroring for employee to Web traffic, perform these tasks:

#### CLI Quick Configuration

To quickly configure local mirroring of traffic from the two ports connected to employee computers, filtering so that only traffic to the external Web is mirrored, copy the following commands and paste them into the switch terminal window:

```

[edit]
set forwarding-options port-mirroring instance employee-web-monitor output interface
ge-0/0/10.0
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
source-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp then accept
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then
port-mirroring-instance employee-web-monitor
set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee

```

#### Step-by-Step Procedure

To configure local mirroring of employee to Web traffic from the two ports connected to employee computers:

1. Configure the local analyzer interface:

```

[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching

```

2. Configure the **employee-web-monitor** output instance (the input to the instance comes from the action of the filter):

```

[edit forwarding-options port-mirroring]
user@switch# set instance employee-web-monitor output interface ge-0/0/10.0

```

3. Configure a firewall filter called **watch-employee** to send mirrored copies of employee requests to the Web to the **employee-web-monitor** instance. Accept all traffic to and from the corporate subnet (destination or source address of **192.0.2.16/28**). Send mirrored copies of all packets destined for the Internet (**destination port 80**) to the **employee-web-monitor** instance.

```

[edit firewall family ethernet-switching]

```

```
user@switch# set filter watch-employee term employee-to-corp from destination-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp from source-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp then accept
user@switch# set filter watch-employee term employee-to-web from destination-port 80
user@switch# set filter watch-employee term employee-to-web then
port-mirroring-instance employee-web-monitor
4. Apply the watch-employee filter to the appropriate ports:

[edit interfaces]
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee
```

**Results** Check the results of the configuration:

```
[edit]
user@switch> show
forwarding-options {
  port-mirroring {
    instance {
      employee-web-monitor {
        family ethernet-switching {
          output {
            interface ge-0/0/10.0;
          }
        }
      }
    }
  }
}
...
firewall family ethernet-switching {
  filter watch-employee {
    term employee-to-corp {
      from {
        destination-address 192.0.2.16/28;
        source-address 192.0.2.16/28;
      }
      then accept {
      }
    }
    term employee-to-web {
      from {
        destination-port 80;
      }
      then port-mirroring-instance employee-web-monitor;
    }
  }
}
...
interfaces {
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan members [employee-vlan, voice-vlan];
        filter {
```

```

        input watch-employee;
    }
}
}
ge-0/0/1 {
    family ethernet-switching {
        filter {
            input watch-employee;
        }
    }
}
}
}

```

### Verification

To confirm that the configuration is correct, perform these tasks:

- [Verifying That the Analyzer Has Been Correctly Created on page 4005](#)
- [Verifying That The Port-Mirroring Instance Is Configured Properly on page 4005](#)

#### *Verifying That the Analyzer Has Been Correctly Created*

**Purpose** Verify that the analyzer **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces, and appropriate output interface.

**Action** You can use the **show forwarding-options analyzer** command to verify that the analyzer is configured properly.

```

user@switch> show forwarding-options analyzer
Analyzer name           : employee-monitor
Mirror rate             : 1
Maximum packet length   : 0
State                   : up
Ingress monitored interfaces : ge-0/0/0.0
Ingress monitored interfaces : ge-0/0/1.0
Output interface        : ge-0/0/10.0

```

**Meaning** This output shows that the analyzer **employee-monitor** has a ratio of 1 (mirroring every packet, the default setting), the maximum size of the original packet that was mirrored (0 indicates the entire packet), the state of the configuration (is up indicates that the analyzer is mirroring the traffic entering the ge-0/0/0, and ge-0/0/1 interfaces, and sending the mirrored traffic to the ge-0/0/10 interface). If the state of the output interface is down or if the output interface is not configured, the value of state will be **down** and the analyzer will not be programmed for mirroring.

#### *Verifying That The Port-Mirroring Instance Is Configured Properly*

**Purpose** Verify that the port-mirroring instance **employee-web-monitor** has been configured properly on the switch with the appropriate input interfaces.

**Action** You can verify that the port-mirroring instance is configured properly by using the **show forwarding-options port-mirroring** command.

```
user@switch> show forwarding-options port-mirroring
Instance Name: employee-web-monitor
Instance Id: 3
Input parameters:
  Rate           : 1
  Run-length     : 0
  Maximum-packet-length : 0
Output parameters:
  Family      State      Destination      Next-hop
  ethernet-switching up      ge-0/0/10.0
```

**Meaning** This output shows that the **employee-web-monitor** instance has a ratio of 1 (mirroring every packet, the default), the maximum size of the original packet that was mirrored (0 indicates an entire packet), the state of the configuration is up and port mirroring is programmed, and that mirrored traffic from the firewall filter action is sent out on interface ge-0/0/10.0. If the state of the output interface is down or if the interface is not configured, the value for state will be down and port mirroring will not be programmed for mirroring.

- Related Documentation**
- [Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches on page 4006](#)
  - [Configuring Mirroring on EX4300 Switches to Analyze Traffic \(CLI Procedure\) on page 4043](#)
  - [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 4047](#)
  - [Understanding Port Mirroring and Analyzers on EX4300 Switches on page 3964](#)

## Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [“Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches” on page 4006](#). For ELS details see: *Getting Started with Enhanced Layer 2 Software*.

EX4300 switches enable you to configure mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on EX4300 switches

You can analyze the mirrored traffic by using a protocol analyzer application running on a remote monitoring station if you are sending mirrored traffic to an analyzer VLAN.

This topic includes two related examples that describe how to mirror traffic entering ports on the switch to the **remote-analyzer** VLAN so that you can perform analysis from a remote monitoring station. The first example shows how to mirror all traffic entering the ports connected to employee computers. The second example shows the same scenario but includes a filter to mirror only the employee traffic going to the Web.



**BEST PRACTICE:** Mirror only necessary packets to reduce potential performance impact. We recommend that you:

- Disable your configured mirroring sessions when you are not using them.
- Specify individual interfaces as input to analyzers rather than specifying all interfaces as input.
- Limit the amount of mirrored traffic by using firewall filters.

This example describes how to configure remote mirroring:

- [Requirements on page 4007](#)
- [Overview and Topology on page 4007](#)
- [Mirroring All Employee Traffic for Remote Analysis on page 4008](#)
- [Mirroring Employee-to-Web Traffic for Remote Analysis on page 4011](#)
- [Verification on page 4015](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.2X50-D10 or later for EX Series switches
- An EX4300 switch connected to another EX4300 switch

The diagram shows an EX4300 Virtual Chassis connected to an EX4300 destination switch.

Before you configure remote mirroring, be sure that:

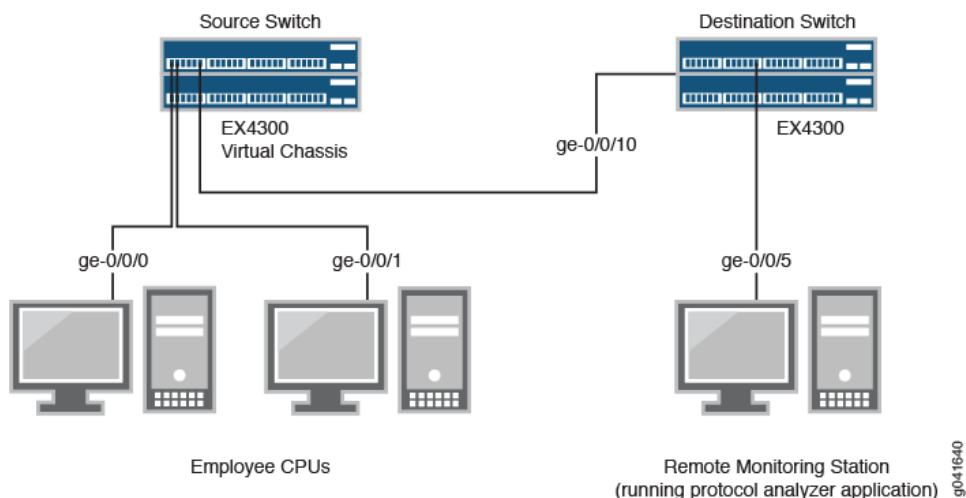
- You have an understanding of mirroring concepts.
- The interfaces that the analyzer will use as input interfaces have been configured on the switch.

### Overview and Topology

This topic includes two related examples that describe how to configure mirroring to the **remote-analyzer** VLAN so that analysis can be performed from a remote monitoring station. The first example shows how to configure a switch to mirror all traffic from employee computers. The second example shows the same scenario, but the setup includes a filter to mirror only the employee traffic going to the Web.

[Figure 56 on page 4008](#) shows the network topology for both these example scenarios.

Figure 56: Remote Mirroring Network Topology Example



In this example:

- Interface ge-0/0/0 is a Layer 2 interface, and interface ge-0/0/1 is a Layer 3 interface (both interfaces on the source switch) that serve as connections for employee computers.
- Interface ge-0/0/10 is a Layer 2 interface that connects the source switch to the destination switch.
- Interface ge-0/0/5 is a Layer 2 interface that connects the destination switch to the remote monitoring station.
- VLAN **remote-analyzer** is configured on all switches in the topology to carry the mirrored traffic.

### Mirroring All Employee Traffic for Remote Analysis

To configure an analyzer for remote traffic analysis for all incoming and outgoing employee traffic, perform these tasks:

#### CLI Quick Configuration

To quickly configure an analyzer for remote traffic analysis for incoming and outgoing employee traffic, copy the following commands and paste them into the switch terminal window:

- Copy and paste the following commands in the source switch terminal window:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/0.0
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/1.0
set forwarding-options analyzer employee-monitor input egress interface ge-0/0/0.0
set forwarding-options analyzer employee-monitor input egress interface ge-0/0/1.0
set forwarding-options analyzer employee-monitor output vlan remote-analyzer
```

- Copy and paste the following commands in the destination switch terminal window:



```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set interfaces ge-0/0/5 unit 0 family ethernet-switching interface-mode trunk
set forwarding-options analyzer employee-monitor input ingress vlan remote-analyzer
set forwarding-options analyzer employee-monitor output interface ge-0/0/5.0
```

**Step-by-Step Procedure** To configure basic remote port mirroring:

1. On the source switch:

- Configure the VLAN ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the interface on the network port connected to the destination switch for trunk mode and associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```

- Configure the **employee-monitor** analyzer:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
user@switch# set instance employee-monitor input egress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input egress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor output vlan remote-analyzer
```

2. On the destination switch:

- Configure the VLAN ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the interface on the destination switch for trunk mode and associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```

- Configure the interface connected to the destination switch for trunk mode:

```
[edit interfaces]
user@switch# set ge-0/0/5 unit 0 family ethernet-switching interface-mode trunk
```

- Configure the **employee-monitor** analyzer:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor input ingress vlan remote-analyzer
user@switch# set analyzer employee-monitor output interface ge-0/0/5.0
```

**Results** Check the results of the configuration on the source switch:

```
[edit]
user@switch> show
forwarding-options {
  analyzer employee-monitor {
```

```
input {
  ingress {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
  }
  egress {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
  }
}
output {
  vlan {
    remote-analyzer;
  }
}
}
interfaces {
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members 999;
        }
      }
    }
  }
}
}
vlangs {
  remote-analyzer {
    vlan-id 999;
    interface {
      ge-0/0/10.0
    }
  }
}
}
```

Check the results of the configuration on the destination switch:

```
[edit]
user@switch> show
interfaces {
  ge0/0/5 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
      }
    }
  }
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
```

```

        members 999;
    }
}
}
}
vpls {
    remote-analyzer {
        vlan-id 999;
        interface {
            ge-0/0/10.0
        }
    }
}
}
forwarding-options {
    analyzer employee-monitor {
        input {
            ingress {
                vlan remote-analyzer;
            }
        }
        output {
            interface {
                ge-0/0/5.0;
            }
        }
    }
}
}
}

```

### Mirroring Employee-to-Web Traffic for Remote Analysis

To configure port mirroring for remote traffic analysis of employee- to- Web traffic, perform these tasks:

#### CLI Quick Configuration

To quickly configure port mirroring to mirror employee traffic to the external Web, copy the following commands and paste them into the switch terminal window:

- Copy and paste the following commands in the source switch terminal window:
 

```

[edit]
set forwarding-options port-mirroring instance employee-web-monitor output vlan 999
set vpls remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
source-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp then
accept
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then
port-mirror-instance employee-web-monitor
set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee
      
```

- Copy and paste the following commands in the destination switch terminal window:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set interfaces ge-0/0/5 unit 0 family ethernet-switching interface-mode trunk
set forwarding-options analyzer employee-web-monitor input ingress vlan remote-analyzer
set forwarding-options analyzer employee-web-monitor output interface ge-0/0/5.0
```

#### Step-by-Step Procedure

To configure port mirroring of all traffic from the two ports connected to employee computers to the **remote-analyzer** VLAN for use from a remote monitoring station:

1. On the source switch:

- Configure the **employee-web-monitor** port mirroring instance:

```
[edit ]
user@switch# set interfaces ge-0/0/10 unit 0 family ethernet-switching port mode trunk
user@switch# set forwarding-options port-mirroring instance employee-web-monitor
output vlan 999
```

- Configure the VLAN ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the interface to associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```

- Configure the firewall filter called **watch-employee**:

```
[edit firewall family ethernet-switching]
user@switch# set filter (Firewall Filters) watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp from source-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp then accept
user@switch# set filter watch-employee term employee-to-web from destination-port
80
user@switch# set filter watch-employee term employee-to-web then port-mirror-instance
employee-web-monitor
```

- Apply the firewall filter to the employee interfaces:

```
[edit interfaces]
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee
```

2. On the destination switch:

- Configure the VLAN ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the interface on the destination switch for trunk mode and associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```

- Configure the interface connected to the destination switch for trunk mode:

```
[edit interfaces]
user@switch# set ge-0/0/5 unit 0 family ethernet-switching interface-mode trunk
```

- Configure the **employee-monitor** analyzer:

```
[edit forwarding-options port-mirroring]
user@switch# set instance employee-web-monitor input ingress vlan remote-analyzer
user@switch# set instance employee-web-monitor output interface ge-0/0/5.0
```

**Results** Check the results of the configuration on the source switch:

```
[edit]
user@switch> show
interfaces {
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members remote-analyzer;
        }
      }
    }
  }
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        filter {
          input watch-employee;
        }
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        filter {
          input watch-employee;
        }
      }
    }
  }
}
firewall {
  family ethernet-switching {
    filter watch-employee {
      term employee-to-corp {
        from {
          source-address {
            192.0.2.16/28;
          }
          destination-address {
            192.0.2.16/28;
          }
        }
        then accept;
      }
    }
  }
}
```

```
        term employee-to-web {
            from {
                destination-port 80;
            }
            then port-mirror-instance employee-web-monitor;
        }
    }
}
forwarding-options {
    analyzer employee-web-monitor {
        output {
            vlan {
                999;
            }
        }
    }
}
vlangs {
    remote-analyzer {
        vlan-id 999;
    }
}
```

Check the results of the configuration on the destination switch:

```
[edit]
user@switch> show
vlangs {
    remote-analyzer {
        vlan-id 999;
    }
}
interfaces {
    ge-0/0/10 {
        unit 0 {
            family ethernet-switching {
                interface-mode trunk;
                vlan {
                    members remote-analyzer;
                }
            }
        }
    }
    ge-0/0/5 {
        unit 0 {
            family ethernet-switching {
                interface-mode trunk;
            }
        }
    }
}
forwarding-options {
    port-mirroring {
        instance employee-web-monitor {
            input {
                ingress {
```

```

        vlan remote-analyzer;
    }
}
output {
    interface {
        ge-0/0/5.0;
    }
}
}
}

```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Analyzer Has Been Correctly Created on page 4015](#)

#### *Verifying That the Analyzer Has Been Correctly Created*

**Purpose** Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.

**Action** You can verify the analyzer is configured as expected by using the **show forwarding-options analyzer** command. To view previously created analyzers that are disabled, go to the J-Web interface.

To verify that the analyzer is configured as expected while monitoring all employee traffic on the source switch, run the **show analyzer** command on the source switch. The following output is displayed for this configuration example:

```

user@switch> show forwarding-options analyzer
Analyzer name           : employee-monitor
Mirror rate             : 1
Maximum packet length   : 0
State                   : up
Ingress monitored interfaces : ge-0/0/0.0
Ingress monitored interfaces : ge-0/0/1.0
Egress monitored interfaces : ge-0/0/0.0
Egress monitored interfaces : ge-0/0/1.0
Output VLAN             : default-switch/remote-analyzer

```

**Meaning** This output shows that the **employee-monitor** instance has a ratio of 1 (mirroring every packet, the default), the maximum size of the original packet that was mirrored (0 indicates the entire packet), the state of the configuration is up (which indicates the proper state and that the analyzer is programmed, and is mirroring the traffic entering ge-0/0/0 and ge-0/0/1 and is sending the mirrored traffic to the VLAN called **remote-analyzer**). If the state of the output interface is down or if the output interface is not configured, the value of state will be down and the analyzer will not be programmed for mirroring.

**Related Documentation** • [Example: Configuring Mirroring for Local Monitoring of Employee Resource Use on EX4300 Switches on page 4000](#)

- [Configuring Mirroring on EX4300 Switches to Analyze Traffic \(CLI Procedure\) on page 4043](#)
- [Understanding Port Mirroring and Analyzers on EX4300 Switches on page 3964](#)

## Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX4300 Switches



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

EX4300 switches enable you to configure mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on EX4300 switches

You can analyze the mirrored traffic by using a protocol analyzer application running on a remote monitoring station if you are sending mirrored traffic to an analyzer VLAN.

This topic includes an example that describes how to mirror traffic entering ports on the switch to the **remote-analyzer** VLAN through a transit switch, so that you can perform analysis from a remote monitoring station.



**BEST PRACTICE:** Mirror only necessary packets to reduce potential performance impact. We recommend that you:

- Disable your configured mirroring sessions when you are not using them.
- Specify individual interfaces as input to analyzers rather than specifying all interfaces as input.
- Limit the amount of mirrored traffic by using firewall filters.

This example describes how to configure remote mirroring through a transit switch:

- [Requirements on page 4017](#)
- [Overview and Topology on page 4017](#)
- [Mirroring All Employee Traffic for Remote Analysis Through a Transit Switch on page 4018](#)
- [Verification on page 4022](#)



## Requirements

This example uses the following hardware and software components:

- An EX4300 switch connected to another EX4300 switch through a third EX4300 switch
- Junos OS Release 13.2X50-D10 or later for EX Series switches

Before you configure remote mirroring, be sure that:

- You have an understanding of mirroring concepts.
- The interfaces that the analyzer will use as input interfaces have been configured on the switch.

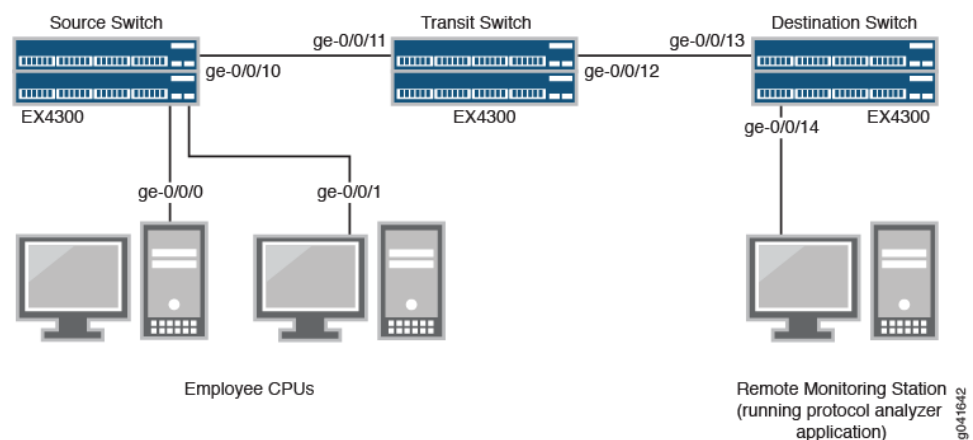
## Overview and Topology

This example describes how to mirror traffic entering ports on the switch to the **remote-analyzer** VLAN through a transit switch so that you can perform analysis from a remote monitoring station. The example shows how to configure a switch to mirror all traffic from employee computers to a remote analyzer.

In this configuration, an analyzer session is required on the destination switch to mirror incoming traffic from the analyzer VLAN to the egress interface to which the remote monitoring station is connected. You must disable MAC learning on the transit switch for the **remote-analyzer** VLAN so that MAC learning is disabled for all member interfaces of the **remote-analyzer** VLAN on the transit switch.

Figure 57 on page 4017 shows the network topology for this example.

**Figure 57: Remote Mirroring Through a Transit Switch Network—Sample Topology**



In this example:

- Interface ge-0/0/0 is a Layer 2 interface, and interface ge-0/0/1 is a Layer 3 interface (both interfaces on the source switch) that serve as connections for employee computers.
- Interface ge-0/0/10 is a Layer 2 interface that connects to the transit switch.
- Interface ge-0/0/11 is a Layer 2 interface on the transit switch.
- Interface ge-0/0/12 is a Layer 2 interface on the transit switch and connects to the destination switch.
- Interface ge-0/0/13 is a Layer 2 interface on the destination switch .
- Interface ge-0/0/14 is a Layer 2 interface on the destination switch and connects to the remote monitoring station.
- VLAN **remote-analyzer** is configured on all switches in the topology to carry the mirrored traffic.

### Mirroring All Employee Traffic for Remote Analysis Through a Transit Switch

To configure mirroring for remote traffic analysis through a transit switch, for all incoming and outgoing employee traffic, perform these tasks:

#### **CLI Quick Configuration**

To quickly configure mirroring for remote traffic analysis through a transit switch, for incoming and outgoing employee traffic, copy the following commands and paste them into the switch terminal window:

- Copy and paste the following commands in the source switch (monitored switch) terminal window:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/0.0
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/1.0
set forwarding-options analyzer employee-monitor input egress interface ge-0/0/0.0
set forwarding-options analyzer employee-monitor input egress interface ge-0/0/1.0
set forwarding-options analyzer employee-monitor output vlan remote-analyzer
```

- Copy and paste the following commands in the transit switch window:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/11 unit 0 family ethernet-switching interface-mode trunk
set vlans remote-analyzer interface ge-0/0/11
set interfaces ge-0/0/12 unit 0 family ethernet-switching interface-mode trunk
set vlans remote-analyzer interface ge-0/0/12
set vlans remote-analyzer no-mac-learning
```

- Copy and paste the following commands in the destination switch window:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/13 unit 0 family ethernet-switching interface-mode trunk
set vlans remote-analyzer interface ge-0/0/13 ingress
set interfaces ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk
```

```
set forwarding-options analyzer employee-monitor input ingress vlan remote-analyzer
set forwarding-options analyzer employee-monitor output interface ge-0/0/14.0
```

### Step-by-Step Procedure

To configure remote mirroring through a transit switch:

1. On the source switch:

- Configure the VLAN ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the interfaces on the network port connected to transit switch for trunk mode and associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```

- Configure the **employee-monitor** analyzer:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor input egress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input egress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor output vlan remote-analyzer
```

2. On the transit switch:

- Configure the VLAN ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the **ge-0/0/11** interface for trunk mode, associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/11 unit 0 family ethernet-switching interface-mode trunk
```

- Configure the **ge-0/0/12** interface for trunk mode, associate it with the **remote-analyzer** VLAN, and set the interface for egress traffic only:

```
[edit interfaces]
user@switch# set ge-0/0/12 unit 0 family ethernet-switching interface-mode trunk
user@switch# set vlans remote-analyzer interface ge-0/0/12
```

- Configure the **no-mac-learning** option for the **remote-analyzer** VLAN to disable MAC learning on all interfaces that are members of the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set vlans remote-analyzer no-mac-learning
```

3. On the destination switch:

- Configure the VLAN ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the **ge-0/0/13** interface for trunk mode, associate it with the **remote-analyzer** VLAN, and set the interface for ingress traffic only:

```
[edit interfaces]
user@switch# set ge-0/0/13 unit 0 family ethernet-switching interface-mode trunk
```

```
user@switch# set vlans remote-analyzer interface ge-0/0/13 ingress
```

- Configure the interface connected to the remote monitoring station for trunk mode:

```
[edit interfaces]
```

```
user@switch# set ge-0/0/14 unit 0 family ethernet-switching interface-mode trunk
```

- Configure the **employee-monitor** analyzer:

```
[edit forwarding-options]
```

```
user@switch# set analyzer employee-monitor input ingress vlan remote-analyzer
```

```
user@switch# set analyzer employee-monitor output interface ge-0/0/14.0
```

**Results** Check the results of the configuration on the source switch:

```
[edit]
user@switch> show
forwarding-options {
  analyzer employee-monitor {
    input {
      ingress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
      }
      egress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
      }
    }
    output {
      vlan {
        remote-analyzer;
      }
    }
  }
}
vlans {
  remote-analyzer {
    vlan-id 999;
  }
}
interfaces {
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          member 999;
        }
      }
    }
  }
}
```

Check the results of the configuration on the transit switch:

```
[edit]
```

```

user@switch> show
vpls {
  remote-analyzer {
    vlan-id 999;
    interface {
      ge-0/0/11.0 {
      }
      ge-0/0/12.0 {
      }
    }
    no-mac-learning;
  }
}
interfaces {
  ge-0/0/11 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
      }
    }
  }
  ge-0/0/12 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
      }
    }
  }
}
}

```

Check the results of the configuration on the destination switch:

```

[edit]
user@switch> show
vpls {
  remote-analyzer {
    vlan-id 999;
    interface {
      ge-0/0/13.0 {
        ingress;
      }
    }
  }
}
interfaces {
  ge-0/0/13 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
      }
    }
  }
  ge-0/0/14 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
      }
    }
  }
}

```

```
    }  
  }  
}  
forwarding-options {  
  analyzer employee-monitor {  
    input {  
      ingress {  
        vlan remote-analyzer;  
      }  
    }  
    output {  
      interface {  
        ge-0/0/14.0;  
      }  
    }  
  }  
}
```

---

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Analyzer Has Been Correctly Created on page 4022](#)

#### *Verifying That the Analyzer Has Been Correctly Created*

**Purpose** Verify that the analyzer named **employee-monitor** has been created on the switch with the appropriate input interfaces and the appropriate output interface.

**Action** You can verify whether the analyzer is configured as expected by using the **show analyzer** command. To view previously created analyzers that are disabled, go to the J-Web interface.

To verify that the analyzer is configured as expected while monitoring all employee traffic on the source switch, run the **show analyzer** command on the source switch. The following output is displayed for this example configuration:

```
user@switch> show forwarding-options analyzer  
Analyzer name           : employee-monitor  
Mirror rate             : 1  
Maximum packet length   : 0  
State                   : up  
Ingress monitored interfaces : ge-0/0/0.0  
Ingress monitored interfaces : ge-0/0/1.0  
Egress monitored interfaces : ge-0/0/0.0  
Egress monitored interfaces : ge-0/0/1.0  
Output vlan             : default-switch/remote-analyzer
```

**Meaning** This output shows that the **employee-monitor** analyzer has a ratio of 1 (mirroring every packet, the default), is mirroring the traffic entering ge-0/0/0 and ge-0/0/1, and sending the mirrored traffic to the analyzer **remote-analyzer**.

- Related Documentation**
- [Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches on page 4006](#)
  - [Example: Configuring Mirroring for Local Monitoring of Employee Resource Use on EX4300 Switches on page 4000](#)
  - [Configuring Mirroring on EX4300 Switches to Analyze Traffic \(CLI Procedure\) on page 4043](#)
  - [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 4047](#)
  - [Understanding Port Mirroring and Analyzers on EX4300 Switches on page 3964](#)

## Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches

Ethernet interfaces on EX Series switches and Junos OS for EX Series switches support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM).

This example describes how to enable and configure OAM CFM on a Gigabit Ethernet interface:

- [Requirements on page 4023](#)
- [Overview and Topology on page 4023](#)
- [Configuring Ethernet OAM Connectivity Fault Management on Switch 1 on page 4023](#)
- [Configuring Ethernet OAM Connectivity Fault Management on Switch 2 on page 4024](#)
- [Verification on page 4026](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 10.2 or later for EX Series switches
- Two EX Series switches connected by a point-to-point Gigabit Ethernet link

### Overview and Topology

CFM can be used to monitor the physical link between two switches. In the following example, two switches are connected by a point-to-point Gigabit Ethernet link. The link between these two switches is monitored using CFM.

### Configuring Ethernet OAM Connectivity Fault Management on Switch 1

- CLI Quick Configuration**
- To quickly configure Ethernet OAM CFM, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain]
set name-format character-string
set maintenance-domain private level 0
set maintenance-association private-ma
set continuity-check hold-interval 1s
```

**Step-by-Step  
Procedure**

To enable and configure OAM CFM on switch 1:

1. Specify the maintenance domain name format:  

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain]  
user@switch1# set name-format character-string
```
2. Specify the maintenance domain name and the maintenance domain level:  

```
[edit protocols oam ethernet connectivity-fault-management]  
user@switch1# set maintenance-domain private level 0
```
3. Create a maintenance association:  

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain  
private]  
user@switch1# set maintenance-association private-ma
```
4. Enable the continuity check protocol and specify the continuity check hold interval:  

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain  
private maintenance-association private-ma]  
user@switch1# set continuity-check hold-interval 1s
```
5. Configure the maintenance association end point (MEP):  

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain  
private maintenance-association private-ma]  
user@switch1# set mep 100 interface ge-1/0/1 auto-discovery direction down
```

**Results**

Check the results of the configuration.

```
[edit]  
user@switch1 > show  
  
protocols {  
  oam {  
    ethernet {  
      connectivity-fault-management {  
        maintenance-domain private {  
          level 0;  
          maintenance-association private-ma {  
            continuity-check {  
              interval 1s;  
            }  
            mep 100 {  
              interface ge-1/0/1;  
              auto-discovery;  
              direction down;  
            }  
          }  
        }  
      }  
    }  
  }  
}
```

---

**Configuring Ethernet OAM Connectivity Fault Management on Switch 2**

---

**CLI Quick  
Configuration**

To quickly configure Ethernet OAM CFM, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain]  
set name-format character-string
```



```

set maintenance-domain private level 0
set maintenance-association private-ma
set continuity-check hold-interval 1s

```

### Step-by-Step Procedure

The configuration on switch 2 mirrors that on switch 2.

1. Specify the maintenance domain name format:  

```

[edit protocols oam ethernet connectivity-fault-management]
user@switch2# set name-format character-string

```
2. Specify the maintenance domain name and the maintenance domain level:  

```

[edit protocols oam ethernet connectivity-fault-management]
user@switch2# set maintenance-domain private level 0

```
3. Create a maintenance association:  

```

[edit protocols oam ethernet connectivity-fault-management maintenance-domain
private]
user@switch2# set maintenance-association private-ma

```
4. Enable the continuity check protocol and specify the continuity check hold interval:  

```

[edit protocols oam ethernet connectivity-fault-management maintenance-domain
private maintenance-association private-ma]
user@switch2# set continuity-check hold-interval 1s

```
5. Configure the maintenance association end point (MEP)  

```

[edit protocols oam ethernet connectivity-fault-management maintenance-domain
private maintenance-association private-ma]
user@switch2# set mep 200 interface ge-0/2/5 auto-discovery direction down

```

### Results

Check the results of the configuration.

```

[edit]
user@switch2 > show

protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain private {
          level 0;
          maintenance-association private-ma {
            continuity-check {
              interval 1s;
            }
            mep 200 {
              interface ge-0/2/5;
              auto-discovery;
              direction down;
            }
          }
        }
      }
    }
  }
}

```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That OAM CFM Has Been Configured Properly on page 4026](#)

### *Verifying That OAM CFM Has Been Configured Properly*

**Purpose** Verify that OAM CFM has been configured properly.

**Action** Use the `show oam ethernet connectivity-fault-management interfaces detail` command:

```
user@switch1# show oam ethernet connectivity-fault-management interfaces detail
```

## Sample Output

```
Interface name: ge-1/0/1.0, Interface status: Active, Link status: Up
Maintenance domain name: private, Format: string, Level: 0
Maintenance association name: private-ma, Format: string
Continuity-check status: enabled, Interval: 1ms, Loss-threshold: 3 frames
MEP identifier: 100, Direction: down, MAC address: 00:90:69:0b:4b:94
MEP status: running
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : yes
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : yes
Statistics:
  CCMs sent                                  : 76
  CCMs received out of sequence              : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received               : 0
  Valid out-of-order LBRs received           : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                             : 0
  LTRs sent                                  : 0
  LTRs received                             : 0
  Sequence number of next LTM request        : 0
Remote MEP count: 2
  Identifier  MAC address      State  Interface
  2001       00:90:69:0b:7f:71  ok    ge-0/2/5.0
```

**Meaning** When the output displays that continuity-check status is **enabled** and displays details of the remote MEP, it means that connectivity fault management (CFM) has been configured properly.

**Related Documentation**

- [Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch on page 3973](#)
- [Junos OS Network Interfaces Configuration Guide](#)

## Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches

Junos OS for EX Series switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example describes how to enable and configure OAM LFM on a Gigabit Ethernet interface:

- [Requirements on page 4027](#)
- [Overview and Topology on page 4027](#)
- [Configuring Ethernet OAM Link Fault Management on Switch 1 on page 4027](#)
- [Configuring Ethernet OAM Link Fault Management on Switch 2 on page 4028](#)
- [Verification on page 4029](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.4 or later for EX Series switches
- Two EX3200 or EX4200 switches connected directly

### Overview and Topology

Junos OS for EX Series switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example uses two EX4200 switches connected directly. Before you begin configuring Ethernet OAM LFM on two switches, connect the two switches directly through a trunk interface.

### Configuring Ethernet OAM Link Fault Management on Switch 1

**CLI Quick Configuration** To quickly configure Ethernet OAM LFM, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet link-fault-management]
set interface ge-0/0/0
set interface ge-0/0/0 link-discovery active
set interface ge-0/0/0 pdu-interval 800
set interface ge-0/0/0 remote-loopback
```

**Step-by-Step Procedure** To configure Ethernet OAM LFM on switch 1:

1. Enable IEEE 802.3ah OAM support on an interface:
 

```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface (OAM LFM) ge-0/0/0
```

2. Specify that the interface initiates the discovery process by configuring the link discovery mode to **active**:  

```
[edit protocols oam ethernet link-fault-management]  
user@switch1# set interface ge-0/0/0 link-discovery active
```
3. Set the periodic OAM PDU-sending interval (in milliseconds) to 800 on switch 1:  

```
[edit protocols oam ethernet link-fault-management]  
user@switch1# set interface pdu-interval 800
```
4. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Ensure that the remote DTE supports remote loopback mode. To set the remote DTE in loopback mode  

```
[edit protocols oam ethernet link-fault-management]  
user@switch1# set interface ge-0/0/0.0 remote-loopback
```

### Results

Check the results of the configuration:

```
[edit]  
user@switch1# show  
  
protocols {  
  oam {  
    ethernet {  
      link-fault-management {  
        interface ge-0/0/0 {  
          pdu-interval 800;  
          link-discovery active;  
          remote-loopback;  
        }  
      }  
    }  
  }  
}
```

---

### Configuring Ethernet OAM Link Fault Management on Switch 2

**CLI Quick Configuration** To quickly configure Ethernet OAM LFM on switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet link-fault-management ]  
set interface ge-0/0/1  
set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

**Step-by-Step Procedure** To configure Ethernet OAM LFM on switch 2:

1. Enable OAM on the peer interface on switch 2:  

```
[edit protocols oam ethernet link-fault-management]  
user@switch2# set interface ge-0/0/1
```
2. Enable remote loopback support for the local interface:  

```
[edit protocols oam ethernet link-fault-management]  
user@switch2# set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

**Results** Check the results of the configuration:

```
[edit]
```

```

user@switch2# show

protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-0/0/1 {
          negotiation-options {
            allow-remote-loopback;
          }
        }
      }
    }
  }
}

```

### Verification

#### *Verifying That OAM LFM Has Been Configured Properly*

**Purpose** Verify that OAM LFM has been configured properly.

**Action** Use the `show oam ethernet link-fault-management` command:

```
user@switch1#show oam ethernet link-fault-management
```

### Sample Output

```

Interface: ge-0/0/0.0
Status: Running, Discovery state: Send Any
Peer address: 00:19:e2:50:3b:e1
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote entity information:
Remote MUX action: forwarding, Remote parser action: forwarding
Discovery mode: active, Unidirectional mode: unsupported
Remote loopback mode: supported, Link events: supported
Variable requests: unsupported

```

**Meaning** When the output displays the MAC address and the discover state is **Send Any**, it means that OAM LFM has been configured properly.

**Related Documentation**

- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 4062](#)
- [Understanding Ethernet OAM Link Fault Management on page 3972](#)

### Example: Configuring Enhanced Network Analytics Features

This example shows how to configure the enhanced network analytics feature, including queue and traffic monitoring.

- [Requirements on page 4030](#)
- [Overview on page 4030](#)
- [Configuration on page 4031](#)
- [Verification on page 4035](#)

## Requirements

---

This example uses the following hardware and software components:

- A QFX5100 standalone switch
- A external streaming server to collect data
- Junos OS Release 13.2X51-D15 software
- TCP server software (for remote streaming servers)

Before you configure network analytics, be sure you have:

- Junos OS Release 13.2X51-D15 or later software installed and running on the QFX5100 switch.
- (Optional for streaming servers for the JSON, CSV, and TSV formats) TCP or UDP server software set up for processing records separated by a newline character (\n) on the remote streaming server.
- (Optional for streaming servers for the GPB format) TCP or UDP build streaming server using the **analytics.proto** file.
- All other network devices running.

## Overview

---

The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data using sophisticated algorithms, and captures the results in reports. Network administrators can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed.

You enable network analytics by first defining a resource profile template, and then applying the profile to the system (for a global configuration) or to individual interfaces.



**NOTE:** You can configure queue and traffic monitoring on physical network interfaces only; logical interfaces and Virtual Chassis physical (VCP) interfaces are not supported.

Disabling of the queue or traffic monitoring supersedes the configuration (enabling) of this feature. You disable monitoring by applying a resource profile that includes the **no-queue-monitoring** or **no-traffic-monitoring** configuration statement at the [edit services analytics resource-profiles] hierarchy level.

---

## Topology

In this example, the QFX5100 switch is connected to an external server used for streaming statistics data.

## Configuration

To configure the network analytics features, perform these tasks:

- [Configuring the Polling Interval for Queue and Traffic Monitoring on page 4031](#)
- [Configuring a Local Statistics File on page 4032](#)
- [Configuring and Applying a Resource Profile for the System on page 4032](#)
- [Configuring and Applying a Resource Profile for an Interface on page 4032](#)
- [Configuring an Export Profile and Collector for Streaming Data on page 4033](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
[edit]
set services analytics resource system polling-interval queue-monitoring 1000
set services analytics resource system polling-interval traffic-monitoring 5
set services analytics collector local file an.stats
set services analytics collector local file an files 3
set services analytics collector local file an size 10m
set services analytics resource-profiles sys-rp queue-monitoring
set services analytics resource-profiles sys-rp traffic-monitoring
set services analytics resource-profiles sys-rp depth-threshold high 999999 low 99
set services analytics resource system resource-profile sys-rp
set services analytics resource-profiles if-rp queue-monitoring
set services analytics resource-profiles if-rp traffic-monitoring
set services analytics resource-profiles if-rp latency-threshold high 2300 low 20
set services analytics resource interfaces xe-0/0/16 resource-profile if-rp
set services analytics resource interfaces xe-0/0/18 resource-profile if-rp
set services analytics resource interfaces xe-0/0/19 resource-profile if-rp
set services analytics export-profiles ep stream-format gpb
set services analytics export-profiles ep interface information
set services analytics export-profiles ep interface statistics queue
set services analytics export-profiles ep interface statistics traffic
set services analytics export-profiles ep interface status link
set services analytics export-profiles ep system information
set services analytics export-profiles ep system status queue
set services analytics export-profiles ep system status traffic
set services analytics collector address 10.94.198.11 port 50001 transport tcp export-profile
ep
set services analytics collector address 10.94.184.25 port 50013 transport udp
export-profile ep
```

### *Configuring the Polling Interval for Queue and Traffic Monitoring*

#### Step-by-Step Procedure

To configure the polling interval queue and traffic monitoring globally:

1. Configure the queue monitoring polling interval (in milliseconds) for the system:  

```
[edit]
set services analytics resource system polling-interval queue-monitoring 1000
```
2. Configure the traffic monitoring polling interval (in seconds) for the system:

```
[edit]
set services analytics resource system polling-interval traffic-monitoring 5
```

### *Configuring a Local Statistics File*

#### **Step-by-Step Procedure**

To configure a file for local statistics collection:

1. Configure the filename:  

```
[edit]
set services analytics collector local file an.stats
```
2. Configure the number of files:  

```
[edit]
set services analytics collector local file an files 3
```
3. Configure the file size:  

```
[edit]
set services analytics collector local file an size 10m
```

### *Configuring and Applying a Resource Profile for the System*

#### **Step-by-Step Procedure**

To define a resource profile template for queue and traffic monitoring resources:

1. Configure a resource profile and enable queue monitoring:  

```
[edit]
set services analytics resource-profiles sys-rp queue-monitoring
```
2. Enable traffic monitoring in the profile:  

```
[edit]
set services analytics resource-profiles sys-rp traffic-monitoring
```
3. Configure the depth-threshold (high and low values) for queue monitoring in the profile:  

```
[edit]
set services analytics resource-profiles sys-rp depth-threshold high 999999 low 99
```
4. Apply the resource profile template to the system resource type for a global configuration:  

```
[edit]
set services analytics resource system resource-profile sys-rp
```

### *Configuring and Applying a Resource Profile for an Interface*

#### **Step-by-Step Procedure**

You can configure queue and traffic monitoring for one or more specific interfaces. The interface-specific configuration supersedes the global (system) configuration. To define a resource profile template for queue and traffic monitoring resources for an interface:

1. Configure a resource profile and enable queue monitoring:  

```
[edit]
set services analytics resource-profiles if-rp queue-monitoring
```



2. Enable traffic monitoring in the profile:
 

```
[edit]
set services analytics resource-profiles if-rp traffic-monitoring
```
3. Configure the latency-threshold (high and low values) for queue monitoring in the profile:
 

```
[edit]
set services analytics resource-profiles if-rp latency-threshold high 2300 low 20
```
4. Apply the resource profile template to the interfaces resource type for specific interfaces:
 

```
[edit]
set services analytics resource interfaces xe-0/0/16 resource-profile if-rp
set services analytics resource interfaces xe-0/0/18 resource-profile if-rp
set services analytics resource interfaces xe-0/0/19 resource-profile if-rp
```

### *Configuring an Export Profile and Collector for Streaming Data*

#### **Step-by-Step Procedure**

To configure a collector (streaming server) for receiving monitoring data:

1. Create an export profile and specify the stream format:
 

```
[edit]
set services analytics export-profiles ep stream-format gpb
```
2. Configure the export profile to include interface information:
 

```
[edit]
set services analytics export-profiles ep interface information
```
3. Configure the export profile to include interface queue statistics:
 

```
[edit]
set services analytics export-profiles ep interface statistics queue
```
4. Configure the export profile to include interface traffic statistics:
 

```
[edit]
set services analytics export-profiles ep interface statistics traffic
```
5. Configure the export profile to include interface status link information:
 

```
[edit]
set services analytics export-profiles ep interface status link
```
6. Configure the export profile to include system information:
 

```
[edit]
set services analytics export-profiles ep system information
```
7. Configure the export profile to include system queue status:
 

```
[edit]
set services analytics export-profiles ep system status queue
```
8. Configure the export profile to include system traffic status:
 

```
[edit]
set services analytics export-profiles ep system status traffic
```

9. Configure the transport protocol for the collector addresses and apply an export profile:

```
[edit]
set services analytics collector address 10.94.198.11 port 50001 transport tcp
export-profile ep
set services analytics collector address 10.94.184.25 port 50013 transport udp
export-profile ep
```



**NOTE:** If you configure the `tcp` or `udp` option for the JSON, CSV, and TSV formats, you must also set up the TCP or UDP client software on the remote collector to process records that are separated by the newline character (`\n`) on the remote server.

If you configure the `tcp` or `udp` option for the GPB format, you must also set up the TCP or UDP build streaming server using the `analytics.proto` file.

**Results** Display the results of the configuration:

```
[edit services analytics]
user@switch# run show configuration
services {
  analytics {
    export-profiles {
      ep {
        stream-format gpb;
        interface {
          information;
          statistics {
            traffic;
            queue;
          }
          status {
            link;
          }
        }
        system {
          information;
          status {
            traffic;
            queue;
          }
        }
      }
    }
  }
  resource-profiles {
    sys-rp {
      queue-monitoring;
      traffic-monitoring;
      depth-threshold high 99999 low 99;
    }
  }
}
```

```

    if-rp {
        queue-monitoring;
        traffic-monitoring;
        latency-threshold high 2300 low 20;
    }
}
resource {
    system {
        resource-profile sys-rp;
        polling-interval {
            traffic-monitoring 5;
            queue-monitoring 1000;
        }
    }
    interfaces {
        xe-0/0/16 {
            resource-profile if-rp;
        }
        xe-0/0/18 {
            resource-profile if-rp;
        }
        xe-0/0/19 {
            resource-profile if-rp;
        }
    }
}
collector {
    local {
        file an size 10m files 3;
    }
    address 10.94.184.25 {
        port 50013 {
            transport udp {
                export-profile ep;
            }
        }
    }
    address 10.94.198.11 {
        port 50001 {
            transport tcp {
                export-profile ep;
            }
        }
    }
}
}
}

```

### Verification

Confirm that the configuration is correct and works as expected by performing these tasks:

- [Verifying the Network Analytics Configuration on page 4036](#)
- [Verifying the Network Analytics Status on page 4036](#)

- [Verifying the Collector Configuration on page 4037](#)
- [Verifying Queue Statistics on page 4037](#)
- [Verifying Traffic Statistics on page 4038](#)

### ***Verifying the Network Analytics Configuration***

**Purpose** Verify the configuration for network analytics.

**Action** From operational mode, enter the **show analytics configuration** command to display the traffic and queue monitoring configuration.

```
user@host> show analytics configuration
Traffic monitoring status is enabled
Traffic monitoring polling interval : 5 seconds
Queue monitoring status is enabled
Queue monitoring polling interval : 1000 milliseconds
Queue depth high threshold : 99999 bytes
Queue depth low threshold : 99 bytes
```

| Interface | Traffic<br>Statistics | Queue<br>Statistics | Queue depth<br>threshold |     | Latency<br>threshold |     |
|-----------|-----------------------|---------------------|--------------------------|-----|----------------------|-----|
|           |                       |                     | High                     | Low | High                 | Low |
|           |                       |                     | (bytes)                  |     | (nanoseconds)        |     |
| xe-0/0/16 | enabled               | enabled             | n/a                      | n/a | 2300                 | 20  |
| xe-0/0/18 | enabled               | enabled             | n/a                      | n/a | 2300                 | 20  |
| xe-0/0/19 | enabled               | enabled             | n/a                      | n/a | 2300                 | 20  |

**Meaning** The output displays the traffic and queue monitoring configuration information on the switch.

### ***Verifying the Network Analytics Status***

**Purpose** Verify the network analytics operational status of the switch.

**Action** From operational mode, enter the **show analytics status global** command to display global traffic and queue monitoring status.

```
user@host> show analytics status global
Traffic monitoring status is enabled
Traffic monitoring polling interval : 5 seconds
Queue monitoring status is enabled
Queue monitoring polling interval : 1000 milliseconds
Queue depth high threshold : 99999 bytes
Queue depth low threshold : 99 bytes
```

From operational mode, enter the **show analytics status** command to display both the interface and global queue monitoring status.

```
user@host> show analytics status
Traffic monitoring status is enabled
Traffic monitoring polling interval : 5 seconds
Queue monitoring status is enabled
Queue monitoring polling interval : 1000 milliseconds
Queue depth high threshold : 99999 bytes
Queue depth low threshold : 99 bytes
```

| Interface | Traffic<br>Statistics | Queue<br>Statistics | Queue depth<br>threshold |     | Latency<br>threshold |     |
|-----------|-----------------------|---------------------|--------------------------|-----|----------------------|-----|
|           |                       |                     | High                     | Low | High                 | Low |
|           |                       |                     | (bytes)                  |     | (nanoseconds)        |     |
| xe-0/0/16 | enabled               | enabled             | n/a                      | n/a | 2300                 | 20  |
| xe-0/0/18 | enabled               | enabled             | n/a                      | n/a | 2300                 | 20  |
| xe-0/0/19 | enabled               | enabled             | n/a                      | n/a | 2300                 | 20  |

**Meaning** The output displays the global and interface status of traffic and queue monitoring on the switch.

### *Verifying the Collector Configuration*

**Action** Verify the configuration for the collector for streamed data is working.

From operational mode, enter the **show analytics collector** command to display the streaming servers configuration.

```
user@host> show analytics collector
Address      Port  Transport Stream format State      Sent
10.94.184.25 50013 udp      gpb        n/a        484
10.94.198.11 50001 tcp      gpb        In progress 0
```

**Meaning** The output displays the collector configuration.



**NOTE:** The connection state of a port configured with the **udp** transport protocol is always displayed as **n/a**.

### *Verifying Queue Statistics*

**Purpose** Verify that queue statistics collection is working.

**Action** From operational mode, enter the **show analytics queue-statistics** command to display the queue statistics.

```
user@host> show analytics queue-statistics
CLI issued at 2014-03-04 15:37:03.116018
Time                Interface  Queue-depth  Latency
                   (bytes)      (nanoseconds)
00:00:00.412371 ago xe-0/0/19    1384656      1107724
00:00:01.412395 ago xe-0/0/19    1375712      1100569
00:00:02.415366 ago xe-0/0/19    1385280      1108224
00:00:03.417395 ago xe-0/0/19    1381744      1105395
00:00:04.411392 ago xe-0/0/19    1368432      1094745
00:00:05.414387 ago xe-0/0/19    1374880      1099904
00:00:06.414365 ago xe-0/0/19    1373632      1098905
00:00:07.416386 ago xe-0/0/19    1370096      1096076
00:00:08.413384 ago xe-0/0/19    1377168      1101734
00:00:09.415379 ago xe-0/0/19    1370720      1096576
00:00:10.418374 ago xe-0/0/19    1381120      1104896
00:00:11.410376 ago xe-0/0/19    1383408      1106726
00:00:12.412372 ago xe-0/0/19    1382576      1106060
00:00:13.417371 ago xe-0/0/19    1387152      1109721
00:00:14.411368 ago xe-0/0/19    1375296      1100236
---(more)---
```

**Meaning** The output displays queue-statistics information, with the latest record at the top of the report.

#### *Verifying Traffic Statistics*

**Purpose** Verify that traffic statistics collection is working.

**Action** From operational mode, enter the **show analytics traffic-statistics** command to display the traffic statistics.

```

user@host> show analytics traffic-statistics
CLI issued at 2014-03-04 15:37:52.047136
Time: 00:00:02.252377 ago, Physical interface: xe-0/0/19
Traffic Statistics:
  Receive          Transmit
Total octets:      15044882432      1502607382656
Total packets:     117538143       11739120146
Unicast packet:    117538143       11739120146
Multicast packets: 0
Broadcast packets: 0
Octets per second: 86488360       8649309384
Packets per second: 84461        8446590
CRC/Align errors:  0
Packets dropped:   0       11760298455
Time: 00:00:02.252377 ago, Physical interface: xe-0/0/18
Traffic Statistics:
  Receive          Transmit
Total octets:      1504619929836    15782818944
Total packets:     11754843131     123303273
Unicast packet:    11754843131     123303273
Multicast packets: 0
Broadcast packets: 0
Octets per second: 8649134008     86487816
Packets per second: 8446458      84461
CRC/Align errors:  5
Packets dropped:   0
Time: 00:00:02.252377 ago, Physical interface: xe-0/0/16
Traffic Statistics:
  Receive          Transmit
Total octets:      1504801437048    757345408
Total packets:     11756261156     5916761
Unicast packet:    11756261156     5916761
Multicast packets: 0
Broadcast packets: 0
Octets per second: 7910619496      0
Packets per second: 7725214       0
CRC/Align errors:  3
Packets dropped:   0

```

**Meaning** The output displays traffic-statistics information.

- Related Documentation**
- [Network Analytics Overview on page 3978](#)
  - [analytics on page 4167](#)
  - [show analytics status on page 4354](#)
  - [show analytics collector on page 4346](#)

## Configuration Tasks

- [Configuring SNMP \(J-Web Procedure\) on page 4040](#)
- [Configuring Mirroring on EX4300 Switches to Analyze Traffic \(CLI Procedure\) on page 4043](#)
- [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 4047](#)
- [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\) on page 4049](#)

- [Configuring Real-Time Performance Monitoring \(J-Web Procedure\) on page 4050](#)
- [Configuring the Interface for RPM Timestamping for Client/Server on an EX Series Switch \(CLI Procedure\) on page 4057](#)
- [Configuring Ethernet OAM Connectivity Fault Management \(CLI Procedure\) on page 4058](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 4062](#)
- [Configuring Interfaces for Uplink Failure Detection \(CLI Procedure\) on page 4064](#)
- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 4065](#)
- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4066](#)
- [Configuring an Iterator Profile on a Switch \(CLI Procedure\) on page 4067](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 4068](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4069](#)
- [Configuring Queue Monitoring on page 4070](#)
- [Configuring Traffic Monitoring on page 4072](#)
- [Configuring a Local File for Network Analytics Data on page 4073](#)
- [Configuring a Remote Collector for Streaming Analytics Data on page 4074](#)

## Configuring SNMP (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You can use the J-Web interface to define system identification information, create SNMP communities, create SNMP trap groups, and configure health monitor options for EX Series switches.

To configure SNMP features:

1. Select **Configure > Services > SNMP**.
2. Enter information into the configuration page for SNMP as described in [Table 427 on page 4041](#).
3. To apply the configuration click **Apply**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.



Table 427: SNMP Configuration Page

| Field                                   | Function                                                                                                                                                                                                                                                                                                                                                                | Your Action                                                                                   |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Identification                          |                                                                                                                                                                                                                                                                                                                                                                         |                                                                                               |
| Contact Information                     | Free-form text string that specifies an administrative contact for the system.                                                                                                                                                                                                                                                                                          | Type contact information for the administrator of the system (such as name and phone number). |
| System Description                      | Free-form text string that specifies a description for the system.                                                                                                                                                                                                                                                                                                      | Type information that describes the system                                                    |
| Local Engine ID                         | Provides an administratively unique identifier of an SNMPv3 engine for system identification.<br><br>The local engine ID contains a prefix and a suffix. The prefix is formatted according to specifications defined in RFC 3411. The suffix is defined by the local engine ID. Generally, the local engine ID suffix is the MAC address of Ethernet management port 0. | Type the MAC address of Ethernet management port 0.                                           |
| System Location                         | Free-form text string that specifies the location of the system.                                                                                                                                                                                                                                                                                                        | Type location information for the system (lab name or rack name, for example).                |
| System Override Name                    | Free-form text string that overrides the system hostname.                                                                                                                                                                                                                                                                                                               | Type the hostname of the system.                                                              |
| <b>Communities</b>                      |                                                                                                                                                                                                                                                                                                                                                                         |                                                                                               |
| To add a community, click <b>Add</b>    |                                                                                                                                                                                                                                                                                                                                                                         |                                                                                               |
| Community Name                          | Specifies the name of the SNMP community.                                                                                                                                                                                                                                                                                                                               | Type the name of the community being added.                                                   |
| Authorization                           | Specifies the type of authorization (either read-only or read-write) for the SNMP community being configured.                                                                                                                                                                                                                                                           | Select the authorization (either read-only or read-write) from the list.                      |
| <b>Traps</b>                            |                                                                                                                                                                                                                                                                                                                                                                         |                                                                                               |
| To add a trap group, click <b>Add</b> . |                                                                                                                                                                                                                                                                                                                                                                         |                                                                                               |
| Trap Group Name                         | Specifies the name of the SNMP trap group being configured.                                                                                                                                                                                                                                                                                                             | Type the name of the group being added.                                                       |

Table 427: SNMP Configuration Page (*continued*)

| Field                    | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Categories               | Specifies which trap categories are added to the trap group being configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <ul style="list-style-type: none"> <li>To generate traps for authentication failures, select <b>Authentication</b>.</li> <li>To generate traps for chassis and environment notifications, select <b>Chassis</b>.</li> <li>To generate traps for configuration changes, select <b>Configuration</b>.</li> <li>To generate traps for link-related notifications (up-down transitions), select <b>Link</b>.</li> <li>To generate traps for remote operation notifications, select <b>Remote operations</b>.</li> <li>To generate traps for remote network monitoring (RMON), select <b>RMON alarm</b>.</li> <li>To generate traps for routing protocol notifications, select <b>Routing</b>.</li> <li>To generate traps on system warm and cold starts, select <b>Startup</b>.</li> <li>To generate traps on Virtual Router Redundancy Protocol (VRRP) events (such as new-master or authentication failures), select <b>VRRP events</b>.</li> </ul> |
| Targets                  | Specifies one or more hostnames or IP addresses for the systems to receive SNMP traps generated by the trap group being configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <ol style="list-style-type: none"> <li>Enter the hostname or IP address, in dotted decimal notation, of the target system to receive the SNMP traps.</li> <li>Click <b>Add</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Health Monitoring        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Enable Health Monitoring | <p>Enables the SNMP health monitor on the switch. The health monitor periodically (over the time you specify in the interval field) checks the following key indicators of switch health:</p> <ul style="list-style-type: none"> <li>Percentage of file storage used</li> <li>Percentage of Routing Engine CPU used</li> <li>Percentage of Routing Engine memory used</li> <li>Percentage of memory used for each system process</li> <li>Percentage of CPU used by the forwarding process</li> <li>Percentage of memory used for temporary storage by the forwarding process</li> </ul> | <p>Select the check box to enable the health monitor and configure options. Clear the check box to disable the health monitor.</p> <p><b>NOTE:</b> If you select the <b>Enable Health Monitoring</b> check box and do not specify options, then SNMP health monitoring is enabled with default values.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Interval                 | <p>Specifies the sampling frequency, in seconds, over which the key health indicators are sampled and compared with the rising and falling thresholds.</p> <p>For example, if you configure the interval as 100 seconds, the values are checked every 100 seconds.</p>                                                                                                                                                                                                                                                                                                                   | <p>Enter an interval time, in seconds, from <b>1</b> through <b>2147483647</b>.</p> <p>The default value is 300 seconds (5 minutes).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

Table 427: SNMP Configuration Page (*continued*)

| Field             | Function                                                                                                                                                                                                                                                                                                       | Your Action                                                                                                                                                                          |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rising Threshold  | <p>Specifies the value at which SNMP generates an event (trap and system log message) when the value of a sampled indicator is increasing.</p> <p>For example, if the rising threshold is 90 (the default), SNMP generates an event when the value of any key indicator reaches or exceeds 90 percent.</p>     | <p>Enter a value from <b>0</b> through <b>100</b>. The default value is <b>90</b>.</p>                                                                                               |
| Falling Threshold | <p>Specifies the value at which SNMP generates an event (trap and system log message) when the value of a sampled indicator is decreasing.</p> <p>For example, if the falling threshold is 80 (the default), SNMP generates an event when the value of any key indicator falls back to 80 percent or less.</p> | <p>Enter a value from <b>0</b> through <b>100</b>. The default value is <b>80</b>.</p> <p><b>NOTE:</b> The falling threshold value must be less than the rising threshold value.</p> |

- Related Documentation**
- [Monitoring System Process Information on page 850](#)
  - [Monitoring System Properties on page 846](#)

## Configuring Mirroring on EX4300 Switches to Analyze Traffic (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Port Mirroring to Analyze Traffic (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

EX4300 switches enable you to configure mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN



**BEST PRACTICE:** Mirror only necessary packets to reduce potential performance impact. We recommend that you:

- Disable your configured mirroring configurations when you are not using them.
- Specify individual interfaces as input to analyzers rather than specifying all interfaces as input.
- Limit the amount of mirrored traffic by using firewall filters.



**NOTE:** If you want to create additional analyzers without deleting the existing analyzers, then disable the existing analyzers by using the `disable analyzer analyzer-name` statement from the command-line interface or the J-Web configuration page for mirroring.



**NOTE:** Interfaces used as output for an analyzer must be configured under the `ethernet-switching` family.

- [Configuring an Analyzer for Local Traffic Analysis on page 4044](#)
- [Configuring an Analyzer for Remote Traffic Analysis on page 4044](#)
- [Configuring Port Mirroring on page 4045](#)

---

### Configuring an Analyzer for Local Traffic Analysis

---

To mirror interface traffic or VLAN traffic on the switch to an interface on the switch (by using analyzers):

1. Choose a name for the analyzer and specify the input:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name input ingress interface interface-name
```

For example, create an analyzer called **employee-monitor** for which the input traffic is packets entering interfaces ge-0/0/0.0 and ge-0/0/1.0:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
[edit forwarding-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
```

2. Configure the destination interface for the mirrored packets:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name output interface interface-name
```

For example, configure ge-0/0/10.0 as the destination interface for the **employee-monitor** analyzer:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor output interface ge-0/0/10.0
```

---

### Configuring an Analyzer for Remote Traffic Analysis

---

To mirror traffic that is traversing interfaces or a VLAN on the switch to a VLAN for analysis from a remote location (by using analyzers):

1. Configure a VLAN to carry the mirrored traffic:

```
[edit]
user@switch# set vlans analyzer-name vlan-id vlan-ID
```

For example, define an analyzer VLAN called **remote-analyzer** and assign it a VLAN ID of **999**:

```
[edit]
user@switch# set vlans remote-analyzer vlan-id 999
```

2. Set the uplink module interface that is connected to the distribution switch to trunk mode and associate it with the analyzer VLAN:

```
[edit]
user@switch# set interfaces interface-name unit 0 family ethernet-switching interface-mode
trunk vlan members vlan-ID
```

For example, set the interface ge-0/1/1 to trunk mode and associate it with the analyzer VLAN ID 999:

```
[edit]
user@switch# set interfaces ge-0/1/1 unit 0 family ethernet-switching interface-mode trunk
vlan members 999
```

3. Configure the analyzer:

- a. Define an analyzer and specify the traffic to be mirrored:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name input ingress interface interface-name
```

For example, define the **employee-monitor** analyzer for which traffic to be mirrored is packets entering interfaces ge-0/0/0.0 and ge-0/0/1.0:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
[edit forwarding-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
```

- b. Specify the analyzer VLAN as the output for the analyzer:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name output vlan vlan-ID
```

For example, specify the **remote-analyzer** VLAN as the output analyzer for the **employee-monitor** analyzer:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor output vlan 999
```

## Configuring Port Mirroring

To filter packets to be mirrored to a port-mirroring instance, create the instance and then use it as the action in the firewall filter. You can use firewall filters in both local and remote mirroring configurations.

If the same port-mirroring instance is used in multiple filters or terms, the packets are copied to the analyzer output port or analyzer VLAN only once.

To filter mirrored traffic, create a port-mirroring instance under the **[edit forwarding-options]** hierarchy level, and then create a firewall filter. The filter can use any of the available match conditions and must have **port-mirror-instance *instance-name*** as an action. This action in the firewall filter configuration provides the input to the port-mirroring instance.

To configure a port-mirroring instance with firewall filters:

1. Configure the port-mirroring instance name (here, **employee-monitor**) and the output:

- a. For local analysis, set the output to the local interface to which you will connect the computer running the protocol analyzer application:

```
[edit forwarding-options]
```

```
user@switch# set port-mirroring instance employee-monitor output interface ge-0/0/10.0
```

- b. For remote analysis, set the output to the **remote-analyzer** VLAN:

```
[edit forwarding-options]
user@switch# set port-mirroring instance employee-monitor output vlan 999
```

2. Create a firewall filter by using any of the available match conditions and assign **employee-monitor** to the **port-mirror-instance** action:

This step shows a firewall filter **example-filter**, with two terms (**no-analyzer** and **to-analyzer**):

- a. Create the first term to define the traffic that should not pass through to the port-mirroring instance **employee-monitor**:

```
[edit firewall family ethernet-switching]
user@switch# set filter (Firewall Filters) example-filter term no-analyzer from
source-address ip-address
[edit firewall family ethernet-switching]
user@switch# set filter example-filter term no-analyzer from destination-address
ip-address
[edit firewall family ethernet-switching]
user@switch# set filter example-filter term no-analyzer then accept
```

- b. Create the second term to define the traffic that should pass through to the port-mirroring instance **employee-monitor**:

```
[edit firewall family ethernet-switching]
user@switch# set filter example-filter term to-analyzer from destination-port 80
[edit firewall family ethernet-switching]
user@switch# set filter example-filter term to-analyzer then port-mirror-instance
employee-monitor
[edit firewall family ethernet-switching]
user@switch# set filter example-filter term to-analyzer then accept
```

3. Apply the firewall filter to the interfaces or VLAN that provide input to the port-mirroring instance:

```
[edit]
user@switch# set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input
example-filter
[edit]
user@switch# set vlan (802.1Q Tagging) remote-analyzer filter input example-filter
```

#### Related Documentation

- [Example: Configuring Mirroring for Local Monitoring of Employee Resource Use on EX4300 Switches on page 4000](#)
- [Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches on page 4006](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Understanding Port Mirroring on EX Series Switches](#)
- [Firewall Filters for EX Series Switches Overview on page 4696](#)

## Configuring Port Mirroring to Analyze Traffic (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

EX Series switches allow you to configure port mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on EX2200, EX3200, EX3300, EX4200, EX4300, EX4500, EX6200 switches
- Packets exiting a VLAN on EX8200 switches

To configure port mirroring on an EX Series switch using the J-Web interface:

1. Select **Configure > Security > Port Mirroring**.

The top of the screen displays analyzer details such as the name, status, analyzer port, ratio, and loss priority.

The bottom of the screen lists ingress and egress ports of the selected analyzer.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—Add an analyzer. Enter information as specified in [Table 428 on page 4048](#).
- **Edit**—Modify details of the selected analyzer. Enter information as specified in [Table 428 on page 4048](#).
- **Delete**—Delete the selected analyzer.
- **Enable/Disable**—Enable or disable the selected analyzer (toggle).



**NOTE:** On EX2200, EX3200, EX4200, and EX4500 switches, only one analyzer can be enabled at a time. On EX8200 switches, a maximum of seven analyzers can be enabled. On EX4300 switches a maximum of four Analyzers/Port Mirror instances can be enabled.



**NOTE:** When an analyzer is deleted or disabled, any filter association is removed.

Table 428: Port Mirroring Configuration Settings

| Field                                                                                 | Function                                                                                                                                                                                                                                                                                                                                                        | Your Action                                                                                                                                                                                                                                                                  |
|---------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Analyzer Name                                                                         | Specifies the name of the analyzer.                                                                                                                                                                                                                                                                                                                             | Type a name for the analyzer.                                                                                                                                                                                                                                                |
| Ratio<br><br><b>NOTE:</b> This option is not supported on EX4300 switches.            | Specifies the ratio of packets to be mirrored. For example: <ul style="list-style-type: none"> <li>A ratio of 1 sends copies of all packets.</li> <li>A ratio of 2047 sends copies of 1 out of every 2047 packets.</li> </ul>                                                                                                                                   | Enter a number from 0 through 2047.                                                                                                                                                                                                                                          |
| Loss Priority<br><br><b>NOTE:</b> This option is not supported on EX4300 switches.    | Specifies the loss priority of the mirrored packets.<br><br>By default, the switch applies a lower priority to mirrored data than to regular port-to-port data—mirrored traffic is dropped in preference to regular traffic when capacity is exceeded.<br><br>For port-mirroring configurations with output to an analyzer VLAN, set the loss priority to high. | Keep the default of low, unless the output is to a VLAN.                                                                                                                                                                                                                     |
| Analyzer Port                                                                         | Specifies a local interface or VLAN to which mirrored packets are sent.<br><br><b>NOTE:</b> A VLAN must have only one associated interface to be specified as an analyzer interface.                                                                                                                                                                            | Click <b>Select</b> . In the Select Analyzer Port/VLAN window, select either port or VLAN as the <b>Analyzer Type</b> . Next, select the required port or VLAN. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the port (interface) from the list. |
| Analyzer Type<br><br><b>NOTE:</b> This option is supported only on EX4300 switches.   | Specifies the analyzer type.                                                                                                                                                                                                                                                                                                                                    | Select the <b>Analyzer Type</b> from the list.                                                                                                                                                                                                                               |
| No Filter check<br><br><b>NOTE:</b> This option is supported only on EX4300 switches. | Enable this option to skip checking for filters on port-mirroring instance.                                                                                                                                                                                                                                                                                     | To enable this option, select the check box.                                                                                                                                                                                                                                 |
| Ingress                                                                               | Specifies interfaces or VLANs for which entering traffic is mirrored.                                                                                                                                                                                                                                                                                           | Click <b>Add</b> . For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.<br><br>Click <b>Remove</b> to delete an ingress interface or VLAN.                                                                                  |
| Egress                                                                                | Specifies interfaces for which exiting traffic is mirrored.                                                                                                                                                                                                                                                                                                     | Click <b>Add</b> and select <b>Port</b> or <b>VLAN</b> . For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.<br><br>Click <b>Remove</b> to remove egress interfaces.                                                       |



## Related Documentation

- *Configuring Port Mirroring to Analyze Traffic (CLI Procedure)*
- *Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches*
- *Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches*
- *Understanding Port Mirroring on EX Series Switches*

## Configuring sFlow Technology for Network Monitoring (CLI Procedure)

sFlow technology is a network monitoring technology for high-speed switched or routed networks. It is a technology that is based on statistical sampling. You can configure sFlow technology to continuously monitor traffic at wire speed on all interfaces simultaneously. Junos OS fully supports the sFlow standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.

To configure sFlow features:

1. Configure the IP address and the UDP port of the collector:

```
[edit protocols]
user@switch# set sflow collector ip-address udp-port port-number
```

The default UDP port is 6343,

2. Enable sFlow technology on a specific interface:

```
[edit protocols sflow]
user@switch# set interfaces interface-name
```



**NOTE:** You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a link aggregation group (LAG), but you can enable it on the member interfaces of a LAG.

3. Specify in seconds how often the sFlow agent polls interfaces:

```
[edit protocols sflow]
user@switch# set polling-interval seconds
```



**NOTE:** Specify 0 if you do not want to poll the interface.

4. Specify the rate at which packets must be sampled. You can specify either an egress or an ingress sampling rate, or both.



**NOTE:** We recommend that you configure the same sampling rates on all the ports on a line card. If you configure different sampling rates on different ports, the lowest value is used for all ports. You could still configure different rates on different line cards.

To specify an egress sampling rate:

```
[edit protocols sflow]
user@switch# set sample-rate egress number
```

To specify an ingress sampling rate:

```
[edit protocols sflow]
user@switch# set sample-rate ingress number
```

5. To configure the polling interval and the egress and ingress sampling rates at the interface level:

```
[edit protocols sflow interfaces interface-name]
user@switch# set polling-interval seconds
[edit protocols sflow interfaces]
user@switch# set sample-rate egress number
[edit protocols sflow interfaces]
user@switch# set sample-rate ingress number
```



**NOTE:** The interface-level configuration overrides the global configuration.

6. To specify an IP address to be used as the agent ID for the sFlow agent:

```
[edit protocols sflow]
user@switch# set agent-id ip-address
```

7. To specify the source IP address to be used for sFlow datagrams:

```
[edit protocols sflow]
user@switch# set source-ip ip-address
```

#### Related Documentation

- [Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995](#)
- [Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch on page 3969](#)

## Configuring Real-Time Performance Monitoring (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

Real-time performance monitoring (RPM) in EX Series switches enables you to configure and send probes to a specified target and monitor the analyzed results to determine packet loss, round-trip time, and jitter. Jitter is the difference in relative transit time between two consecutive probes. You can set up probe owners and configure one or more performance probe tests under each probe owner.

The ways in which you can use RPM include:

- Monitor time delays between devices.
- Monitor time delays at the protocol level.

- Set thresholds to trigger SNMP traps when threshold values are exceeded. You can configure thresholds for round-trip time, ingress or egress delay, standard deviation, jitter, successive lost probes, and total lost probes per test.
- Determine automatically whether a path exists between a host switch and its configured Border Gateway Protocol (BGP) neighbors. You can view the results of the discovery using an SNMP client.
- Use the history of the most recent 50 probes to analyze trends in your network and predict future needs.

Probes collect packets per destination and per application, including PING Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets.

EX Series switches support the following tests and probe types:

- Ping tests:
  - ICMP echo
  - ICMP timestamp
- HTTP tests:
  - HTTP get (not available for BGP RPM services)
- UDP and TCP tests with user-configured ports:
  - UDP echo
  - TCP connection
  - UDP timestamp

To account for latency in the communication of probe messages, you can enable timestamping of the probe packets. You must configure both the requester and the responder to timestamp the RPM packets. The RPM features provides an additional configuration option to set one-way hardware timestamps. Use one-way timestamps when you want information about one-way, rather than round-trip, times for packets to traverse the network between the requester and the responder.



#### NOTE:

- EX Series switches support hardware timestamps for UDP and ICMP probes. EX Series switches do not support hardware timestamps for HTTP or TCP probes.
- If the responder does not support hardware timestamps, RPM can only report the round-trip measurements, it cannot calculate round-trip jitter.
- In EX Series switches timestamps apply only to IPv4 traffic.

To configure RPM using the J-Web interface:

1. Select **Troubleshoot > RPM > Configure RPM**.
2. In the **Configure RPM** page, enter information as specified in [Table 429 on page 4052](#).
  - a. Click **Add** to set up the **Owner Name** and **Performance Probe Tests**.
  - b. Select a probe owner from **Probe Owners** list and click **Delete** to remove the selected probe owner
  - c. Double-click one of the probe owners in **Probe Owners** list to display the list of performance probe tests.
  - d. Double-click one of the performance probe tests to edit the test parameters.
3. Enter the **Maximum Number of Concurrent Probes** and specify the **Probe Servers**.
4. Click **Apply** to apply the RPM probe settings.

**Table 429: RPM Probe Owner, Concurrent Probes, and Probe Servers Configuration Fields**

| Field                               | Function                                                                                                                                                             | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Probe Owners                        | Identifies a owner for whom one or more RPM tests are configured. In most implementations, the owner name identifies a network on which a set of tests is being run. | <ol style="list-style-type: none"> <li>1. Click <b>Add</b> and type an owner name.</li> <li>2. In <b>Performance Probe Tests</b>, click <b>Add</b> to define the RPM test parameters. See <a href="#">Table 430 on page 4053</a> for information on configuring RPM test parameters.</li> <li>3. Click <b>OK</b> to save the settings or <b>Cancel</b> to exit from the window without saving the changes.</li> </ol>                                                          |
| Maximum Number of Concurrent Probes | Specifies the maximum number of concurrent probes allowed.                                                                                                           | Type a number from 1 through 500.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Probe Servers                       | Specifies the servers that act as receivers and transmitters for the probes.                                                                                         | Set up the following servers: <ul style="list-style-type: none"> <li>• TCP Probe Server—Specifies the port on which the device is to receive and transmit TCP probes. Type the number 7 (a standard TCP port number) or a port number from 49160 through 65535.</li> <li>• UDP Probe Server—Specifies the port on which the device is to receive and transmit UDP probes. Type the number 7 (a standard TCP port number) or a port number from 49160 through 65535.</li> </ul> |

Table 430: Performance Probe Tests Configuration Fields

| Field                      | Function                                                                                                                                         | Your Action                                                                                                                                                                                                                                                                                             |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Identification</b>      |                                                                                                                                                  |                                                                                                                                                                                                                                                                                                         |
| Test Name                  | Identifies the RPM test.                                                                                                                         | Type a test name.                                                                                                                                                                                                                                                                                       |
| Target (Address or URL)    | Specifies the IP address or the URL of the probe target.                                                                                         | Type the IP address in dotted decimal notation or the URL of the probe target. If the target is a URL, type a fully formed URL that includes <b>http://</b> .                                                                                                                                           |
| Source Address             | Specifies the IP address to be used as the probe source address.                                                                                 | Type the source address to be used for the probe. If you do not supply this value, the packet uses the outgoing interface's address as the probe source address.                                                                                                                                        |
| Routing Instance           | Specifies the routing instance over which the probe is sent.                                                                                     | Type the routing instance name. The routing instance applies only to <b>icmp-ping</b> and <b>icmp-ping-timestamp</b> probe types. The default routing instance is <b>inet.0</b> .                                                                                                                       |
| History Size               | Specifies the number of probe results to be saved in the probe history.                                                                          | Type a number from 0 through 255. The default history size is 50.                                                                                                                                                                                                                                       |
| <b>Request Information</b> |                                                                                                                                                  |                                                                                                                                                                                                                                                                                                         |
| Probe Type                 | Specifies the type of probe to send as part of the test.                                                                                         | Select a probe type from the list: <ul style="list-style-type: none"> <li>• <b>http-get</b></li> <li>• <b>http-get-metadata</b></li> <li>• <b>icmp-ping</b></li> <li>• <b>icmp-ping-timestamp</b></li> <li>• <b>tcp-ping</b></li> <li>• <b>udp-ping</b></li> <li>• <b>udp-ping-timestamp</b></li> </ul> |
| Interval                   | Sets the wait time (in seconds) between probe transmissions.                                                                                     | Type a number from 1 through 255 .                                                                                                                                                                                                                                                                      |
| Test Interval              | Sets the wait time (in seconds) between tests.                                                                                                   | Type a number from 0 through 86400 .                                                                                                                                                                                                                                                                    |
| Probe Count                | Sets the total number of probes to be sent for each test.                                                                                        | Type a number from 1 through 15.                                                                                                                                                                                                                                                                        |
| Moving Average Size        | Specifies the number of samples to be used in the statistical calculation operations to be performed across a number of the most recent samples. | Type a number from 0 through 255.                                                                                                                                                                                                                                                                       |

Table 430: Performance Probe Tests Configuration Fields (*continued*)

| Field                           | Function                                                                                                                                                                                                                                                                                                           | Your Action                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Destination Port                | Specifies the TCP or UDP port to which probes are sent.<br><br>To use TCP or UDP probes, you must configure the remote server as a probe receiver. Both the probe server and the remote server must be Juniper Networks network devices configured to receive and transmit RPM probes on the same TCP or UDP port. | Type the number 7 (a standard TCP or UDP port number) or a port number from 49160 through 65535. |
| DSCP Bits                       | Specifies the Differentiated Services code point (DSCP) bits. This value must be a valid 6-bit pattern.                                                                                                                                                                                                            | Type a valid 6-bit pattern.                                                                      |
| Data Size                       | Specifies the size (in bytes) of the data portion of the ICMP probes.                                                                                                                                                                                                                                              | Type a number from 0 through 65507.                                                              |
| Data Fill                       | Specifies the hexadecimal value of the data portion of the ICMP probes.                                                                                                                                                                                                                                            | Type a hexadecimal value from 1h through 800h .                                                  |
| <b>Hardware Timestamp</b>       |                                                                                                                                                                                                                                                                                                                    |                                                                                                  |
| One Way Hardware Timestamp      | Enables one-way hardware timestamp.                                                                                                                                                                                                                                                                                | To enable timestamping, select the check box.                                                    |
| Destination Interface           | Enables hardware timestamp on the specified interface.                                                                                                                                                                                                                                                             | Select an interface from the list.                                                               |
| <b>Maximum Probe Thresholds</b> |                                                                                                                                                                                                                                                                                                                    |                                                                                                  |
| Successive Lost Probes          | Sets the number of probes that can be lost successively, if exceeded, triggers a probe failure and generates a system log message.                                                                                                                                                                                 | Type a number from 0 through 15.                                                                 |
| Lost Probes                     | Sets the number of probes that can be lost , if exceeded, triggers a probe failure and generates a system log message.                                                                                                                                                                                             | Type a number from 0 through 15.                                                                 |
| Round Trip Time                 | Sets the round-trip time (in microseconds), from the switch to the remote server, if exceeded, triggers a probe failure and generates a system log message.                                                                                                                                                        | Type a number from 0 through 60000000.                                                           |
| Jitter                          | Sets the jitter (in microseconds), if exceeded, triggers a probe failure and generates a system log message.                                                                                                                                                                                                       | Type a number from 0 through 60000000.                                                           |

Table 430: Performance Probe Tests Configuration Fields (*continued*)

| Field                              | Function                                                                                                                                                     | Your Action                                                                                                                                                           |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Standard Deviation                 | Sets the maximum allowable standard deviation (in microseconds), if exceeded, triggers a probe failure and generates a system log message.                   | Type a number from 0 through 60000000.                                                                                                                                |
| Egress Time                        | Sets the one-way time (in microseconds), from the switch to the remote server, if exceeded, triggers a probe failure and generates a system log message.     | Type a number from 0 through 60000000.                                                                                                                                |
| Ingress Time                       | Sets the one-way time (in microseconds), from the remote server to the switch, if exceeded, triggers a probe failure and generates a system log message.     | Type a number from 0 through 60000000 (microseconds).                                                                                                                 |
| Jitter Egress Time                 | Sets the outbound-time jitter (in microseconds), if exceeded triggers a probe failure and generates a system log message.                                    | Type a number from 0 through 60000000.                                                                                                                                |
| Jitter Ingress Time                | Sets the inbound-time jitter (in microseconds), if exceeded, triggers a probe failure and generates a system log message.                                    | Type a number from 0 and 60000000.                                                                                                                                    |
| Egress Standard Deviation          | Sets the maximum allowable standard deviation of outbound times (in microseconds), if exceeded, triggers a probe failure and generates a system log message. | Type a number from 0 through 60000000.                                                                                                                                |
| Ingress Standard Deviation         | Sets the maximum allowable standard deviation of inbound times (in microseconds), if exceeded, triggers a probe failure and generates a system log message.  | Type a number from 0 through 60000000.                                                                                                                                |
| <b>Traps</b>                       |                                                                                                                                                              |                                                                                                                                                                       |
| Egress Jitter Exceeded             | Generates SNMP traps when the threshold for jitter in outbound time is exceeded.                                                                             | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Egress Standard Deviation Exceeded | Generates SNMP traps when the threshold for standard deviation in outbound times is exceeded.                                                                | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |

Table 430: Performance Probe Tests Configuration Fields (*continued*)

| Field                               | Function                                                                                        | Your Action                                                                                                                                                           |
|-------------------------------------|-------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Egress Time Exceeded                | Generates SNMP traps when the threshold for maximum outbound time is exceeded.                  | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Ingress Jitter Exceeded             | Generates SNMP traps when the threshold for jitter in inbound time is exceeded.                 | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Ingress Standard Deviation Exceeded | Generates SNMP traps when the threshold for standard deviation in inbound times is exceeded.    | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Ingress Time Exceeded               | Generates SNMP traps when the threshold for maximum inbound time is exceeded.                   | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Jitter Exceeded                     | Generates SNMP traps when the threshold for jitter in round-trip time is exceeded.              | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Probe Failure                       | Generates SNMP traps when the threshold for the number of successive lost probes is exceeded.   | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| RTT Exceeded                        | Generates SNMP traps when the threshold for maximum round-trip time is exceeded.                | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Standard Deviation Exceeded         | Generates SNMP traps when the threshold for standard deviation in round-trip times is exceeded. | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Test Completion                     | Generates SNMP traps when a test is completed.                                                  | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Test Failure                        | Generates SNMP traps when the threshold for the total number of lost probes is exceeded.        | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |



- Related Documentation**
- [Configuring SNMP \(J-Web Procedure\) on page 4040](#)
  - [Viewing Real-Time Performance Monitoring Information on page 4268](#)

## Configuring the Interface for RPM Timestamping for Client/Server on an EX Series Switch (CLI Procedure)

Use real-time performance monitoring (RPM) to configure active probes to track and monitor traffic across the network and to investigate network problems. To configure basic RPM probes on the EX Series switch, you must configure the probe owner, the test, and the specific parameters of the RPM probe.

You can also set a timestamp to improve the measurement of latency or jitter. The probe is timestamped by the device originating the probe (the RPM client). If you do not enable hardware timestamps, the timer values are set. You should configure both the RPM client (the requester) and the RPM server (the responder) to timestamp the RPM packets. However, if the RPM server does not support hardware timestamps, RPM can only report the round-trip measurements.

Timestamps apply only to IPv4 traffic.

You can enable hardware timestamps for the following RPM probe types:

- **icmp-ping**
- **icmp-ping-timestamp**
- **udp-ping**
- **udp-ping-timestamp**

To configure RPM probes and enable hardware timestamping:

1. Specify the probe owner:
 

```
[edit services rpm]
user@switch# set probe owner
```
2. Specify a test name. A test represents the range of probes over which the standard deviation, average, and jitter are calculated.
 

```
[edit services rpm probe owner]
user@switch# set test test-name
```
3. Specify the packet and protocol contents of the probe:
 

```
[edit services rpm probe owner test test-name]
user@switch# set probe-type type
```
4. Specify the destination IPv4 address to be used for the probes:
 

```
[edit services rpm probe owner test test-name]
user@switch# set target address
```
5. Specify the number of probes within a test:
 

```
[edit services rpm probe owner test test-name]
user@switch# set probe-count count
```
6. Specify the time, in seconds, to wait between sending packets:
 

```
[edit services rpm probe owner test test-name]
```

```
user@switch# set probe-interval interval
```

7. Specify the time, in seconds, to wait between tests:

```
[edit services rpm probe owner test test-name]  
user@switch# set test-interval interval
```

8. Specify the source IP address to be used for probes. If the source IP address is not one of the switch's assigned addresses, the packet uses the outgoing interface's address as its source.

```
[edit services rpm probe owner test test-name]  
user@switch# set source-address address
```

9. Specify the value of the Differentiated Services (DiffServ) field within the IP header. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern.

```
[edit services rpm probe owner test test-name]  
user@switch# set dscp-code-point dscp-bits
```

10. If you are using ICMP probes, specify the size of the data portion of ICMP probes:

```
[edit services rpm probe owner test test-name]  
user@switch# set data-size size
```

11. Enable hardware timestamping of RPM probe messages:

```
[edit services rpm probe owner test test-name]  
user@switch# set hardware-timestamp
```

**Related  
Documentation**

- [Configuring Real-Time Performance Monitoring \(J-Web Procedure\) on page 4050](#)
- [Understanding Real-Time Performance Monitoring on EX Series Switches on page 3960](#)

## Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)

Ethernet interfaces on Juniper Networks EX Series Ethernet Switches and Juniper Networks Junos OS for EX Series switches support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM).

This topic describes these tasks:

1. [Creating the Maintenance Domain on page 4058](#)
2. [Configuring the Maintenance Domain MIP Half Function on page 4059](#)
3. [Creating a Maintenance Association on page 4059](#)
4. [Configuring the Continuity Check Protocol on page 4060](#)
5. [Configuring a Maintenance Association End Point on page 4060](#)
6. [Configuring a Connectivity Fault Management Action Profile on page 4061](#)
7. [Configuring the Linktrace Protocol on page 4061](#)

### Creating the Maintenance Domain

---

A maintenance domain comprises network entities such as operators, providers, and customers. To enable connectivity fault management (CFM) on an Ethernet interface, you must create a maintenance domains, maintenance associations, and MEPs.

To create a maintenance domain:

1. Specify a name for the maintenance domain:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set maintenance-domain domain-name
```

2. Specify a format for the maintenance domain name. If you specify **none**, no name is configured:

- A plain ASCII character string
- A domain name service (DNS) format
- A media access control (MAC) address plus a two-octet identifier in the range 0 through 65,535
- **none**

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set name-format format
```

For example, to specify the name format as MAC address plus a two-octet identifier:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set name-format mac+2oct
```

3. Configure the maintenance domain level, which is used to indicate the nesting relationship between this domain and other domains. Use a value from 0 through 7:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set level level
```

### Configuring the Maintenance Domain MIP Half Function

MIP Half Function (MHF) divides the maintenance association intermediate point (MIP) functionality into two unidirectional segments, improves visibility with minimal configuration, and improves network coverage by increasing the number of points that can be monitored. MHF extends monitoring capability by responding to loop-back and link-trace messages to help isolate faults. Whenever a MIP is configured, the MIP half function value for all maintenance domains and maintenance associations must be the same.

To configure the MIP half function:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set mip-half-function (none | default | explicit)
```

### Creating a Maintenance Association

In a CFM maintenance domain, each service instance is called a maintenance association.

To create a maintenance association:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set maintenance-association ma-name
```

## Configuring the Continuity Check Protocol

---

The continuity check protocol is used for fault detection by a maintenance association end point (MEP) within a maintenance association. The MEP periodically sends continuity check multicast messages. The receiving MEPs use the continuity check messages (CCMs) to build a MEP database of all MEPs in the maintenance association.

To configure the continuity check protocol:

1. Enable the continuity check protocol:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain  
domain-name maintenance-association ma-name]  
user@switch# set continuity-check
```

2. Specify the continuity check hold interval. The hold interval is the number of minutes to wait before flushing the MEP database if no updates occur. The default value is 10 minutes.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain  
domain-name maintenance-association ma-name continuity-check]  
user@switch# set hold-interval number
```

3. Specify the CCM interval. The interval is the time between the transmission of CCMs. You can specify 10 minutes (10m), 1 minute (1m), 10 seconds (10s), 1 second (1s), 100 milliseconds (100ms), or 10 milliseconds (10ms).

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain  
domain-name maintenance-association ma-name continuity-check]  
user@switch# set interval number
```

4. Specify the number of CCMs (that is, protocol data units) that can be lost before the MEP is marked as down. The default number of protocol data units (PDUs) is 3.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain  
domain-name maintenance-association ma-name continuity-check]  
user@switch# set loss-threshold number
```

## Configuring a Maintenance Association End Point

---

To configure a maintenance association end point:

1. Specify an ID for the MEP. The value can be from 1 through 8191.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain  
domain-name maintenance-association ma-name]  
user@switch# set mep mep-id]
```

2. Enable maintenance endpoint automatic discovery if you want to have the MEP accept continuity check messages (CCMs) from all remote MEPs of the same maintenance association:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain  
domain-name maintenance-association ma-name mep mep-id]  
user@switch# set auto-discovery
```

3. You can specify that CFM packets (CCMs) be transmitted only in one direction for the MEP, that is, the direction be set as **down** so that CCMs are transmitted only out of (not into) the interface configured on this MEP.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain  
domain-name maintenance-association ma-name mep mep-id]  
user@switch# set direction down
```

- Specify the logical interface to which the MEP is attached. It can be either an access interface or a trunk interface. If you specify a trunk interface, the VLAN associated with that interface must have a VLAN ID.



**NOTE:** You cannot associate an access interface that belongs to multiple VLANs with the MEP.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name maintenance-association ma-name mep mep-id]
user@switch# set interface interface-name
```

- You can configure a remote MEP from which CCMs are expected. If autodiscovery is not enabled, the remote MEP must be configured under the **mep** statement. If the remote MEP is not configured under the **mep** statement, the CCMs from the remote MEP are treated as errors.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name maintenance-association ma-name mep mep-id]
user@switch# set remote-mep mep-id
```

### Configuring a Connectivity Fault Management Action Profile

You can configure an action profile and specify the action to be taken when any of the configured events occur. Alternatively, you can configure an action profile and specify default actions when connectivity to a remote MEP fails.

To configure an action profile:

- Specify a name for an action profile:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set action-profile profile-name
```

- Configure the action of the action profile:

```
[edit protocols oam ethernet connectivity-fault-management action-profile
profile-name]
user@switch# set action interface-down
```

- Configure one or more events under the action profile, the occurrence of which will trigger the corresponding action to be taken:

```
[edit protocols oam ethernet connectivity-fault-management action-profile
profile-name]
user@switch# set event event
```

See [Junos OS Network Interfaces Configuration Guide](#)

### Configuring the Linktrace Protocol

The linktrace protocol is used for path discovery between a pair of maintenance points. Linktrace messages are triggered by an administrator using the **traceroute** command to verify the path between a pair of MEPs under the same maintenance association. Linktrace messages can also be used to verify the path between a MEP and a MIP under the same maintenance domain.

To configure the linktrace protocol:

1. Configure the linktrace path age timer. If no response to a linktrace request is received, the request and response entries are deleted after the age timer expires:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set linktrace age time
```

2. Configure the number of linktrace reply entries to be stored per linktrace request:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set linktrace path-database-size path-database-size
```

#### Related Documentation

- [Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches on page 4023](#)
- [Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch on page 3973](#)
- [Junos OS Network Interfaces Configuration Guide](#)

## Configuring Ethernet OAM Link Fault Management (CLI Procedure)

Ethernet OAM link fault management (LFM) can be used for physical link-level fault detection and management. The IEEE 802.3ah LFM works across point-to-point Ethernet links either directly or through repeaters.

To configure Ethernet OAM LFM using the CLI:

1. Enable IEEE 802.3ah OAM support on an interface:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name
```



**NOTE:** You can configure Ethernet OAM LFM on aggregated interfaces.



**NOTE:** The remaining steps are optional. You can choose which of these features to configure for Ethernet OAM LFM on your switch.

2. Specify whether the interface or the peer initiates the discovery process by configuring the link discovery mode to **active** or **passive** (**active** = interface initiates; **passive** = peer initiates):

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name link-discovery active
```

3. Configure a periodic OAM PDU-sending interval (in milliseconds) for fault detection:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface pdu-interval interval
```

4. Specify the number of OAM PDUs that an interface can miss before the link between peers is considered down:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name pdu-threshold threshold-value
```

5. Configure event threshold values on an interface for the local errors that trigger the sending of link event TLVs:

- Set the threshold value (in seconds) for sending frame-error events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-error count
```

- Set the threshold value (in seconds) for sending frame-period events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-period count
```

- Set the threshold value (in seconds) for sending frame-period-summary events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-period-summary count
```

- Set the threshold value (in seconds) for sending symbol-period events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds symbol-period count
```



**NOTE:** You can disable the sending of link event TLVs.

To disable the sending of link event TLVs:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name negotiation-options no-allow-link-events
```

6. Create an action profile to define event fault flags and thresholds to be taken when the link fault event occurs. Then apply the action profile to one or more interfaces. (You can also apply multiple action profiles to a single interface.)

- a. Name the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name
```

- b. Specify actions to be taken by the system when the link fault event occurs:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name action syslog
user@switch# set action-profile profile-name action link-down
```

- c. Specify events for the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name event link-adjacency-loss
```



**NOTE:** For each action profile, you must specify at least one link event and one action. The actions are taken only when all of the events in the action profile are true. If more than one action is specified, all actions are executed. You can set a low threshold for a specific action such as logging the error and set a high threshold for another action such as system logging.

7. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Set the remote DTE in loopback

mode (the remote DTE must support remote-loopback mode) and then enable remote loopback support for the local interface.

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name remote-loopback
user@switch# set interface interface-name negotiation-options allow-remote-loopback
```

**Related  
Documentation**

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027](#)
- [Understanding Ethernet OAM Link Fault Management on page 3972](#)

## Configuring Interfaces for Uplink Failure Detection (CLI Procedure)

You can configure uplink failure detection on EX Series switches to help ensure balanced traffic flow. Using this feature, switches can monitor and detect link failure on uplink interfaces and can propagate the failure to downlink interfaces so that servers connected to those downlink interfaces can switch over to secondary interfaces.

Follow these configuration guidelines:

- You can configure a maximum of 48 groups for each switch.
- You can configure a maximum of 48 uplink interfaces and 48 downlink interfaces in each group.
- You can configure physical links and logical links in separate groups.
- Ensure that all the interfaces in the group are up. If the interfaces are down, uplink failure detection does not work.



**NOTE:** Routed VLAN interfaces (RVIs) cannot be configured as uplink interfaces to be monitored.

To configure uplink failure detection on a switch:

1. Specify a name for the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name
```

2. Add an uplink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-monitor interface-name
```

3. Repeat Step 2 for adding each uplink interface to the group.



**NOTE:** An interface can be configured as link-to-monitor in multiple groups.

4. Add a downlink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-disable interface-name
```

5. Repeat Step 4 for adding each downlink interface to the group.





**NOTE:** After you have configured a group, use the [show uplink-failure-detection group \*group-name\*](#) command to verify that all interfaces in the group are up.

**Related Documentation**

- [Verifying That Uplink Failure Detection Is Working Correctly on page 4268](#)
- [Understanding Uplink Failure Detection on page 3977](#)

## Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements (CLI Procedure)

Ethernet frame delay measurement is a useful tool for providing performance statistics or supporting or challenging service-level agreements (SLAs). By default, Ethernet frame delay measurement uses software for timestamping and delay calculations. You can configure an EX Series switch to perform and display Ethernet frame delay measurements on Ethernet interfaces. The switches support software-assisted timestamping.

Before you can begin configuring MEP interfaces to support Ethernet frame delay measurements on switches, ensure that you have:

- Configured Operation, Administration, and Maintenance (OAM) connectivity fault management (CFM) correctly
- Enabled distributed periodic packet management (PPM) (distributed PPM is enabled by default)

To configure MEP interfaces on switches to support Ethernet frame delay measurements:

1. Enable the Ethernet frame delay measurement by issuing the [monitor ethernet delay-measurement](#) operational mode command. In this command, you must specify one measurement type (either one-way or two-way measurement), and you must specify either the unicast MAC address of the peer MEP or its numeric identifier.

Optionally, you can also specify the following parameters:

- Number of frames to send to the peer MEP (**count *count***)
- Number of seconds to wait between sending frames (**wait *time***)
- Priority value of the delay measurement request frame (**priority *value***)
- Size of the data in the data TLV of the request packet (**size *value***)
- Suppression of the insertion of the session ID TLV in the request packet (**no-session-id-tlv**)

```
user@switch> monitor ethernet delay-measurement maintenance-domain md-name
maintenance-association ma-name one-way mep remote-mep-id count count wait
time priority value size value no-session-id-tlv
```

**Related Documentation**

- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4066](#)

- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4069](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 4068](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 3974](#)

## Configuring One-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure)

Ethernet frame delay measurement is a useful tool for providing performance statistics or supporting or challenging service-level agreements (SLAs). You can configure the frame delay measurements in either a one-way mode or a two-way (round-trip) mode to gather frame delay statistics. For one-way Ethernet frame delay measurement, clocks at the local and remote MEPs need to be synchronized. However, clock synchronization is not required for two-way Ethernet frame delay measurement.

Before you begin configuring one-way Ethernet frame delay measurements on two EX Series switches, ensure that you have:

- Configured Operation, Administration, and Maintenance (OAM) connectivity fault management (CFM) correctly on both the switches
- Synchronized the system clocks of both the switches

To configure one-way Ethernet frame delay measurements:

1. Configure the maintenance domain, maintenance association, and MEP ID on both the switches.
2. From either switch, start a one-way Ethernet frame delay measurement:

```
user@switch> monitor ethernet delay-measurement maintenance-domain md-name
maintenance-association ma-name one-way mep remote-mep-id count count wait
time
```

You can view the result on the other switch:

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md-name maintenance-association ma-name local-mep mep-id
remote-mep mep-id
```

### Related Documentation

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 4065](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4069](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 4068](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 3974](#)

## Configuring an Iterator Profile on a Switch (CLI Procedure)

Ethernet frame delay measurement provides fine control to operators for triggering delay measurement on a given service and can be used to monitor service-level agreements (SLAs). You can create an iterator profile with its parameters to periodically transmit SLA measurement packets in the form of ITU-Y.1731-compliant frames for two-way delay measurement.

To create an iterator profile:

1. Specify a name for an SLA iterator profile—for example, i1:

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@switch# edit sla-iterator-profiles i1
```

2. (Optional) Configure the cycle time, which is the time (in milliseconds) between back-to-back transmissions of SLA frames.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles i1]
user@switch# set cycle-time cycle-time-value
```

3. (Optional) Configure the iteration period, which indicates the maximum number of cycles per iteration (the number of connections registered to an iterator cannot exceed this value).

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles i1]
user@switch# set iteration-period iteration-period-value
```

4. Configure the measurement type as two-way delay measurement.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles i1]
user@switch# set measurement-type two-way-delay
```

5. (Optional) Configure the calculation weight for delay.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles i1]
user@switch# set calculation-weight delay delay-value
```

6. (Optional) Configure the calculation weight for delay variation.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles i1]
user@switch# set calculation-weight delay-variation delay-variation-value
```

7. Configure a remote MEP with the iterator profile.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
md-name maintenance-association ma-name mep mep-id remote-mep
remote-mep-id]
user@switch# set sla-iterator-profiles i1
```

### Related Documentation

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 4065](#)

- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4066](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4069](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 4068](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 3974](#)

## Triggering an Ethernet Frame Delay Measurement Session on a Switch

To trigger Ethernet frame delay measurement, use the **monitor ethernet delay-measurement** operational command and specify the following values:

- Either one-way (**one-way**) or two-way (**two-way**) measurement
- Either the MAC address (**remote-mac-address**) or the MEP ID (**mep**) of the remote host
- The maintenance domain (**maintenance-domain**)
- The maintenance association (**maintenance-association**)
- (Optional) Any or all of these options: **count**, **size**, **wait**, **no-session-id-tlv**, **priority**

For example:

```
user@switch> monitor ethernet delay-measurement one-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10 size 50 wait 5 no-session-id-tlv
priority 1
```

### Related Documentation

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 4065](#)
- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4066](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4069](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 3974](#)

## Configuring Two-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure)

Ethernet frame delay measurement is a useful tool for providing performance statistics or supporting or challenging service-level agreements (SLAs). You can configure the frame delay measurements in either a one-way mode or a two-way (round-trip) mode to gather frame delay statistics. For one-way Ethernet frame delay measurement, clocks at the local and remote MEPs need to be synchronized. However, clock synchronization is not required for two-way Ethernet frame delay measurement.

Before you begin configuring two-way Ethernet frame delay measurements on two EX Series switches, ensure that you have:

- Configured Operation, Administration, and Maintenance (OAM) connectivity fault management (CFM) correctly on both the switches

To configure two-way Ethernet frame delay measurements:

1. Configure the maintenance domain, maintenance association, and MEP ID on both the switches.
2. From either switch, start a two-way Ethernet frame delay measurement:

```
user@switch> monitor ethernet delay-measurement maintenance-domain md-name
maintenance-association ma-name two-way mep remote-mep-id count count wait
time
```

You can view the result on the other switch:

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md-name maintenance-association ma-name local-mep mep-id
remote-mep mep-id
```

### Related Documentation

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 4065](#)
- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4066](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 4068](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 3974](#)

## Configuring Queue Monitoring

Network analytics queue monitoring provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data using sophisticated algorithms, and captures the results in reports. You can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed.

You enable queue monitoring by first defining a resource profile template, and then applying the profile to the system (for a global configuration) or to individual interfaces.



**NOTE:** You can configure queue monitoring on physical network interfaces only; logical interfaces and Virtual Chassis physical (VCP) interfaces are not supported.



**NOTE:** This procedure requires Junos OS Release 13.2X51-D15 or later to be installed on your device.

To configure queue monitoring on a QFX Series standalone switch:

1. Configure the queue monitoring polling interval (in milliseconds) globally (for the system):

[edit]

**set services analytics resource system polling-interval queue-monitoring *interval***

2. Configure a resource profile for the system, and enable queue monitoring:

[edit]

**set services analytics resource-profiles *profile-name* queue-monitoring**

3. Configure high and low values of the depth-threshold (in bytes) for queue monitoring in the system profile:

[edit]

**set services analytics resource-profiles *profile-name* depth-threshold high *number* low *number***

For both high and low values, the range is from 1 to 1,250,000,000 bytes, and the default value is 0 bytes.



**NOTE:** You can configure either the depth-threshold or latency threshold for the system, but not both.

4. Apply the resource profile template to the system for a global configuration:

[edit]

**set services analytics resource system resource-profile *profile-name***

5. Configure an interface-specific resource profile and enable queue monitoring for the interface:

```
[edit]
set services analytics resource-profiles profile-name queue-monitoring
```

6. Configure the latency-threshold (high and low values) for queue monitoring in the interface-specific profile:

```
[edit]
set services analytics resource-profiles profile-name latency-threshold high number
low number
```

For both high and low values, the range is from 1 to 100,000,000 nanoseconds, and the default value is 1,000,000 nanoseconds.



**NOTE:** You can configure either the depth-threshold or latency threshold for interfaces, but not both.

7. Apply the resource profile template for interfaces to one or more interfaces:

```
[edit]
set services analytics resource interfaces interface-name resource-profile profile-name
```



**NOTE:** If a conflict arises between the system and interface configurations, the interface-specific configuration supersedes the global (system) configuration.

#### Related Documentation

- [Network Analytics Overview on page 3978](#)
- [Example: Configuring Enhanced Network Analytics Features on page 4029](#)
- [analytics on page 4167](#)

## Configuring Traffic Monitoring

Network analytics queue monitoring provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data using sophisticated algorithms, and captures the results in reports. You can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed.

You enable traffic monitoring by first defining a resource profile template, and then applying the profile to the system (for a global configuration) or to individual interfaces.



**NOTE:** You can configure traffic monitoring on physical network interfaces only; logical interfaces and Virtual Chassis physical (VCP) interfaces are not supported.



**NOTE:** This procedure requires Junos OS Release 13.2X51-D15 or later to be installed on your device.

To configure traffic monitoring on a QFX Series standalone switch:

1. Configure the traffic monitoring polling interval (in seconds) for the system:

[edit]

**set services analytics resource system polling-interval traffic-monitoring *interval***

2. Configure a resource profile for the system, and enable traffic monitoring in the profile:

[edit]

**set services analytics resource-profiles *profile-name* traffic-monitoring**

3. Apply the resource profile to the system for a global configuration:

[edit]

**set services analytics resource system resource-profile *profile-name***

4. Configure a resource profile for interfaces, and enable traffic monitoring in the profile:

[edit]

**set services analytics resource-profiles *profile-name* traffic-monitoring**



**NOTE:** If a conflict arises between the system and interface configurations, the interface-specific configuration supersedes the global (system) configuration.

5. Apply the resource profile template to one or more interfaces:

[edit]

**set services analytics resource interfaces *interface-name* resource-profile *profile-name***



- Related Documentation**
- [Network Analytics Overview on page 3978](#)
  - [Example: Configuring Enhanced Network Analytics Features on page 4029](#)
  - [analytics on page 4167](#)

## Configuring a Local File for Network Analytics Data

The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data using sophisticated algorithms, and captures the results in reports. Network administrators can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed.

To save the queue and traffic statistics data in a local file, you must configure a filename to store it.



**NOTE:** This procedure requires Junos OS Release 13.2X51-D15 or later to be installed on your device.

To configure a local file for storing queue and traffic monitoring statistics:

1. Configure a filename:

[edit]

**set services analytics collector local file *filename***

There is no default filename. If you do not configure a filename, network analytics statistics are not saved locally.

2. Configure the number of files (from 2 to 1000 files):

[edit]

**set services analytics collector local file *filename* files *number***

3. Configure the file size (from 10 to 4095 MB) in the format of xm:

[edit]

**set services analytics collector local file an size *size***

- Related Documentation**
- [Network Analytics Overview on page 3978](#)
  - [Example: Configuring Enhanced Network Analytics Features on page 4029](#)
  - [analytics on page 4167](#)

## Configuring a Remote Collector for Streaming Analytics Data

The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data using sophisticated algorithms, and captures the results in reports. Network administrators can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed.

You can configure an export profile to define the stream format and type of data, and one or more remote servers (collectors) to receive streaming network analytics data.



**NOTE:** This procedure requires Junos OS Release 13.2X51-D15 or later to be installed on your device.

To configure a collector for receiving streamed analytics data:

1. Create an export profile and specify the stream format:

```
[edit]
set services analytics export-profiles profile-name stream-format format
```

2. Configure the export profile to include interface information:

```
[edit]
set services analytics export-profiles profile-name interface information
```

3. Configure the export profile to include interface queue statistics:

```
[edit]
set services analytics export-profiles profile-name interface statistics queue
```

4. Configure the export profile to include interface traffic statistics:

```
[edit]
set services analytics export-profiles profile-name interface statistics traffic
```

5. Configure the export profile to include interface status link information:

```
[edit]
set services analytics export-profiles profile-name interface status link
```

6. Configure the export profile to include system information:

```
[edit]
set services analytics export-profiles profile-name system information
```

7. Configure the export profile to include system queue status:

```
[edit]
set services analytics export-profiles profile-name system status queue
```

8. Configure the export profile to include system traffic status:

```
[edit]
set services analytics export-profiles profile-name system status traffic
```

9. Configure the transport protocol for the collector addresses and apply the export profile:

[edit]

set services analytics collector address *ip-address* port *port* transport *protocol*  
export-profile *profile-name*

set services analytics collector address *ip-address* port *port* transport *protocol*  
export-profile *profile-name*



**NOTE:** If you configure the `tcp` or `udp` option for the JSON, CSV, and TSV formats, you must also set up the TCP or UDP client software on the remote collector to process records that are separated by the newline character (`\n`) on the remote server.

If you configure the `tcp` or `udp` option for the GPB format, you must also set up the TCP or UDP build streaming server using the `analytics.proto` file.

#### Related Documentation

- [Network Analytics Overview on page 3978](#)
- [Example: Configuring Enhanced Network Analytics Features on page 4029](#)
- [analytics on page 4167](#)

## Configuration Statements: SNMP

- [\[edit snmp\] Configuration Statement Hierarchy on EX Series Switches on page 4078](#)
- [address \(SNMP\) on page 4085](#)
- [address-mask on page 4086](#)
- [agent-address on page 4086](#)
- [alarm \(SNMP RMON\) on page 4087](#)
- [authorization on page 4088](#)
- [bucket-size on page 4088](#)
- [categories on page 4089](#)
- [client-list on page 4089](#)
- [client-list-name on page 4090](#)
- [clients on page 4090](#)
- [commit-delay on page 4091](#)
- [community \(SNMP\) on page 4092](#)
- [community \(SNMP RMON\) on page 4093](#)
- [community-name \(SNMP\) on page 4094](#)
- [contact \(SNMP\) on page 4095](#)
- [description \(SNMP\) on page 4096](#)
- [description \(SNMP RMON\) on page 4097](#)
- [destination-port on page 4097](#)
- [engine-id \(SNMP\) on page 4098](#)

- [event \(SNMP\) on page 4099](#)
- [falling-event-index on page 4099](#)
- [falling-threshold on page 4100](#)
- [falling-threshold on page 4101](#)
- [falling-threshold-interval on page 4102](#)
- [filter-duplicates on page 4102](#)
- [filter-interfaces on page 4103](#)
- [group \(Configuring Access Privileges\) on page 4104](#)
- [group \(Associating a Security Name\) on page 4105](#)
- [health-monitor on page 4105](#)
- [history on page 4106](#)
- [interface \(SNMP RMON History\) on page 4107](#)
- [interface \(SNMP\) on page 4107](#)
- [interval on page 4108](#)
- [interval \(SNMP Health Monitor\) on page 4108](#)
- [interval \(SNMP RMON\) on page 4109](#)
- [location \(SNMP\) on page 4109](#)
- [logical-system \(SNMP\) on page 4110](#)
- [message-processing-model on page 4111](#)
- [name on page 4111](#)
- [nonvolatile on page 4112](#)
- [notify on page 4113](#)
- [notify-filter \(Applying to the Management Target\) on page 4114](#)
- [notify-filter \(Configuring the Profile Name\) on page 4114](#)
- [notify-view on page 4115](#)
- [oid \(SNMP\) on page 4115](#)
- [oid on page 4116](#)
- [owner on page 4116](#)
- [parameters on page 4117](#)
- [port \(SNMP\) on page 4117](#)
- [read-view on page 4118](#)
- [request-type on page 4119](#)
- [rising-event-index on page 4120](#)
- [rising-threshold \(SNMP Health Monitor\) on page 4121](#)
- [rising-threshold \(SNMP RMON\) on page 4122](#)
- [rmon on page 4122](#)
- [rmon on page 4123](#)

- [routing-instance \(SNMP\) on page 4124](#)
- [routing-instance \(SNMPv3\) on page 4125](#)
- [sample-type on page 4125](#)
- [security-level \(Defining Access Privileges\) on page 4126](#)
- [security-level \(Generating SNMP Notifications\) on page 4127](#)
- [security-model \(Access Privileges\) on page 4128](#)
- [security-model \(Group\) on page 4129](#)
- [security-model \(SNMP Notifications\) on page 4130](#)
- [security-name \(Security Group\) on page 4131](#)
- [security-name \(Community String\) on page 4132](#)
- [security-name \(SNMP Notifications\) on page 4133](#)
- [security-to-group on page 4134](#)
- [snmp on page 4134](#)
- [snmp on page 4135](#)
- [snmp-community on page 4136](#)
- [source-address \(SNMP\) on page 4137](#)
- [startup-alarm on page 4138](#)
- [syslog-subtag on page 4138](#)
- [tag \(SNMPv3\) on page 4139](#)
- [tag-list on page 4139](#)
- [target-address on page 4140](#)
- [target-parameters on page 4141](#)
- [targets on page 4142](#)
- [traceoptions \(SNMP\) on page 4143](#)
- [trap-group on page 4145](#)
- [trap-options on page 4146](#)
- [type \(SNMP RMON\) on page 4147](#)
- [type \(SNMPv3\) on page 4148](#)
- [v3 on page 4149](#)
- [vacm on page 4151](#)
- [variable on page 4152](#)
- [version \(SNMP\) on page 4152](#)
- [view \(Associating a MIB View with a Community\) on page 4153](#)
- [view \(Configuring a MIB View\) on page 4154](#)
- [write-view on page 4155](#)

## [edit snmp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit snmp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*

This topic lists:

- [Supported Statements in the \[edit snmp\] Hierarchy Level on page 4078](#)
- [Unsupported Statements in the \[edit snmp\] Hierarchy Level on page 4083](#)

### Supported Statements in the [edit snmp] Hierarchy Level

---

The following hierarchy shows the **[edit snmp]** configuration statements supported on EX Series switches:

```
snmp {
  client-list list-name {
    address {
      restrict;
    }
  }
  community community-name {
    authorization (read-only | read-write);
    client-list-name list-name;
    clients {
      address <restrict>;
    }
    routing-instance instance-name;
    routing-instance instance-name {
      client-list-name list-name;
      clients {
        address <restrict>;
      }
    }
  }
  view view-name;
}
contact contact-information;
description description;
engine-id {
  (local engine-id | use-default-ip-address | use-mac-address);
}
filter-duplicates;
filter-interfaces {
  interfaces
  all-internal-interfaces;
```

```

interface 1;
interface 2;
}
health-monitor {
  falling-threshold percentage;
  idp {
    falling-threshold;
    interval seconds;
    rising-threshold;
  }
  interval seconds;
  rising-threshold percentage;
}
interface [ interface-names ];
location location;
name system-name;
nonvolatile {
  commit-delay seconds;
}
rmon {
  alarm index {
    description description;
    falling-event-index index;
    falling-threshold integer;
    falling-threshold-interval seconds;
    interval seconds;
    request-type (get-next-request | get-request | walk-request);
    rising-event-index index;
    rising-threshold integer;
    sample-type (absolute-value | delta-value);
    startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
    syslog-subtag text-string;
    variable oid-variable;
  }
  event index {
    community community-name;
    description description;
    type (log | log-and-trap | none | snmptrap);
  }
  history index {
    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
  }
}
routing-instance-access {
  access-list {
    routing-instance-name <restrict>;
  }
}
traceoptions {
  file <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}

```

```
}
trap-group group-name {
  categories {
    authentication;
    chassis;
    configuration;
    link;
    otn-alarms {
      alarm-name;
    }
    remote-operations;
    rmon-alarm;
    routing;
    services;
    sonet-alarms {
      alarm-name;
    }
    startup;
    vrrp-events;
  }
  destination-port port-number;
  routing-instance instance-name;
  routing-instance instance-name;
  targets {
    address;
  }
  version (all | v1 | v2);
}
trap-options {
  agent-address outgoing-interface;
  enterprise-oid;
  routing-instance instance-name;
  routing-instance instance-name {
    source-address (address | lo0);
  }
  source-address address;
}
v3 {
  ... the v3 subhierarchy appears after the main [edit snmp] hierarchy level ...
}
view view-name {
  oid object-identifier <exclude | include>;
}
}

snmp {
  v3 {
    notify name {
      tag tag-name;
      type (inform | trap);
    }
    notify-filter profile-name {
      oid oid <exclude | include>;
    }
    snmp-community community-index {
      community-name community-name;
    }
  }
}
```



```

context context-name;
security-name security-name;
tag tag-name;
}
target-address target-address-name {
  address address;
  address-mask address-mask;
  routing-instance routing-instance-name;
  port port-number;
  retry-count number;
  routing-instance routing-instance-name;
  tag-list tag-list;
  target-parameters parameter-name;
  timeout seconds;
}
target-parameters parameter-name {
  notify-filter profile-name;
  parameters {
    message-processing-model (v1 | v2c | v3);
    security-level (authentication | none | privacy);
    security-model (usm | v1 | v2c);
    security-name security-name;
  }
}
usm {
  local-engine {
    user username {
      authentication-md5 {
        authentication-key password;
        authentication-password password;
      }
      authentication-none;
      authentication-sha {
        authentication-key password;
        authentication-password password;
      }
      privacy-3des {
        privacy-password password;
      }
      privacy-aes128 {
        privacy-password password;
      }
      privacy-des {
        privacy-password password;
      }
      privacy-none;
    }
  }
  remote-engine engine-id {
    user username {
      authentication-md5 {
        authentication-key password;
        authentication-password password;
      }
      authentication-none;
      authentication-sha {

```

```
        authentication-key
        authentication-password password;
    }
    privacy-3des {
        privacy-password password;
    }
    privacy-aes128 {
        privacy-password password;
    }
    privacy-des {
        privacy-password password;
    }
    privacy-none;
}
}
}
vacm {
    access {
        group group-name {
            context-prefix prefix {
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        context-match (exact | prefix);
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
    default-context-prefix prefix {
        security-model (any | usm | v1 | v2c) {
            security-level (authentication | none | privacy) {
                context-match (exact | prefix);
                notify-view view-name;
                read-view view-name;
                write-view view-name;
            }
        }
    }
}
}
}
security-to-group {
    security-model (usm | v1 | v2c) {
        security-name security-name {
            group group-name;
        }
    }
}
}
}
```

---

### Unsupported Statements in the [edit snmp] Hierarchy Level

---

All statements in the **[edit snmp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 431: Unsupported [edit snmp] Configuration Statements on EX Series Switches**

| Statement | Hierarchy |
|-----------|-----------|
|-----------|-----------|

**NOTE:** Variables, such as *community-name*, are not shown in the statements or hierarchies.

---

Table 431: Unsupported [edit snmp] Configuration Statements on EX Series Switches (*continued*)

| Statement      | Hierarchy                                       |
|----------------|-------------------------------------------------|
| logical-system | [edit snmp community]<br>[edit snmp trap-group] |

Table 431: Unsupported [edit snmp] Configuration Statements on EX Series Switches (*continued*)

| Statement                   | Hierarchy                                                 |
|-----------------------------|-----------------------------------------------------------|
|                             | [edit snmp trap-options]<br>[edit snmp v3 target-address] |
| logical-systems-trap-filter | [edit snmp]                                               |

- Related Documentation**
- [Configuring SNMP \(J-Web Procedure\) on page 4040](#)
  - *Network Management Administration Guide for Routing Devices*

## address (SNMP)

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | address <i>address</i> ;                                                                                                  |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-address <i>target-address-name</i> ]                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches. |
| <b>Description</b>              | Specify the SNMP target address.                                                                                          |
| <b>Options</b>                  | <b>address</b> —IPv4 address of the system to receive traps or informs. You must specify an address, not a hostname.      |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Address</i></li> </ul>                                        |

## address-mask

---

|                                 |                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>address-mask address-mask;</code>                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 target-address target-address-name]</code>                                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 on the QFX Series. |
| <b>Description</b>              | Define and verify the source addresses for a group of target addresses for SNMP traps and informs.                                                                                            |
| <b>Options</b>                  | <b>address-mask</b> —Define a range of addresses.                                                                                                                                             |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Address Mask</i></li></ul>                                                                                                         |

## agent-address

---

|                                 |                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>agent-address outgoing-interface;</code>                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit snmp trap-options]</code>                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                              |
| <b>Description</b>              | Set the agent address of all SNMPv1 traps generated by this router or switch. Currently, the only option is <b>outgoing-interface</b> , which sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.                                                                                        |
| <b>Options</b>                  | <b>outgoing-interface</b> —Value of the agent address of all SNMPv1 traps generated by this router or switch. The <b>outgoing-interface</b> option sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.<br><b>Default:</b> disabled (the agent address is not specified in SNMPv1 traps). |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Agent Address for SNMP Traps</i></li></ul>                                                                                                                                                                                                                                  |

## alarm (SNMP RMON)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>alarm <i>index</i> {     description <i>description</i>;     falling-event-index <i>index</i>;     falling-threshold <i>integer</i>;     falling-threshold-interval <i>seconds</i>;     interval <i>seconds</i>;     request-type (get-next-request   get-request   walk-request);     rising-event-index <i>index</i>;     rising-threshold <i>integer</i>;     sample-type (absolute-value   delta-value);     startup-alarm (falling-alarm   rising-alarm   rising-or-falling alarm);     syslog-subtag <i>syslog-subtag</i>;     variable <i>oid-variable</i>; }</pre> |
| <b>Hierarchy Level</b>          | [edit snmp rmon]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure RMON alarm entries.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><i>index</i>—Identifies this alarm entry as an integer.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring an RMON Alarm Entry and Its Attributes</i></li> <li>• <a href="#">event (SNMP) on page 4099</a></li> <li>• <i>Configuring RMON Alarms and Events</i></li> <li>• <i>RMON MIB Event, Alarm, Log, and History Control Tables</i></li> <li>• <i>Monitoring RMON MIB Tables</i></li> <li>• <i>Understanding RMON</i></li> </ul>                                                                                                                                                                                              |

## authorization

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authorization <i>authorization</i>;</code>                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit snmp community <i>community-name</i> ]                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Set the access authorization for SNMP <b>Get</b> , <b>GetBulk</b> , <b>GetNext</b> , and <b>Set</b> requests.                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>authorization</i> —Access authorization level: <ul style="list-style-type: none"><li>• <b>read-only</b>—Enable <b>Get</b>, <b>GetNext</b>, and <b>GetBulk</b> requests.</li><li>• <b>read-write</b>—Enable all requests, including <b>Set</b> requests. You must configure a view to enable <b>Set</b> requests.</li></ul> <b>Default:</b> read-only |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Communities</a></li></ul>                                                                                                                                                                                                                                                          |

## bucket-size

---

|                                 |                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bucket-size <i>number</i>;</code>                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">snmp rmon history</a> ]                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                        |
| <b>Description</b>              | Configure the sampling of Ethernet statistics for network fault diagnosis, planning, and performance tuning.                                                                                |
| <b>Default</b>                  | 50                                                                                                                                                                                          |
| <b>Options</b>                  | <i>number</i> —Number of discrete samples of Ethernet statistics requested.                                                                                                                 |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP (J-Web Procedure) on page 4040</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul> |



## categories

|                                 |                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>categories {<br/>    category;<br/>}</code>                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <code>[edit snmp trap-group group-name]</code>                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                            |
| <b>Description</b>              | Define the types of traps that are sent to the targets of the named trap group.                                                                                                                                                                                      |
| <b>Default</b>                  | If you omit the <b>categories</b> statement, all trap types are included in trap notifications.                                                                                                                                                                      |
| <b>Options</b>                  | <b>category</b> —Name of a trap type: <b>authentication</b> , <b>chassis</b> , <b>configuration</b> , <b>link</b> , <b>remote-operations</b> , <b>rmon-alarm</b> , <b>routing</b> , <b>services</b> , <b>sonet-alarms</b> , <b>startup</b> , or <b>vrmp-events</b> . |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring SNMP Trap Groups</i></li> </ul>                                                                                                                                                                              |

## client-list

|                                 |                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>client-list client-list-name {<br/>    ip-addresses;<br/>}</code>                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit snmp]</code>                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for QFX Series switches. |
| <b>Description</b>              | Define a list of SNMP clients.                                                                                                                                                                  |
| <b>Options</b>                  | <b>client-list-name</b> —Name of the client list.<br><br><b>ip-addresses</b> —IP addresses of the SNMP clients to be added to the client list,                                                  |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Adding a Group of Clients to an SNMP Community</i></li> </ul>                                                                                       |

## client-list-name

---

|                                 |                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>client-list-name</code> <i>client-list-name</i> ;                                                               |
| <b>Hierarchy Level</b>          | [edit snmp community <i>community-name</i> ]                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches. |
| <b>Description</b>              | Add a client list or prefix list to an SNMP community.                                                                |
| <b>Options</b>                  | <i>client-list-name</i> —Name of the client list or prefix list.                                                      |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Adding a Group of Clients to an SNMP Community</i></li></ul>               |

## clients

---

|                                 |                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>clients {<br/>    address &lt;restrict&gt;;<br/>}</pre>                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit snmp community <i>community-name</i> ]                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                  |
| <b>Description</b>              | Specify the IPv4 or IPv6 addresses of the SNMP client hosts that are authorized to use this community.                                                                                                                                                                                                     |
| <b>Default</b>                  | If you omit the <b>clients</b> statement, all SNMP clients using this community string are authorized to access the router.                                                                                                                                                                                |
| <b>Options</b>                  | <i>address</i> —Address of an SNMP client that is authorized to access this router. You must specify an address, not a hostname. To specify more than one client, include multiple <i>address</i> options.<br><br><i>restrict</i> —(Optional) Do not allow the specified SNMP client to access the router. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring SNMP Communities</i></li></ul>                                                                                                                                                                                                                      |

---

## commit-delay

---

|                                 |                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>commit-delay <i>seconds</i>;</code>                                                                                       |
| <b>Hierarchy Level</b>          | [edit snmp nonvolatile]                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.       |
| <b>Description</b>              | Configure the timer for the SNMP <b>Set</b> reply and start of the commit.                                                      |
| <b>Options</b>                  | <b><i>seconds</i></b> —Delay between an affirmative SNMP <b>Set</b> reply and start of the commit.<br><b>Default:</b> 5 seconds |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Commit Delay Timer</i></li></ul>                                     |

## community (SNMP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>community <i>community-name</i> {<br/>    authorization <i>authorization</i>;<br/>    client-list-name <i>client-list-name</i>;<br/>    clients {<br/>        address restrict;<br/>    }<br/>    view <i>view-name</i>;<br/>}</pre>                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Define an SNMP community. An SNMP community authorizes SNMP clients based on the source IP address of incoming SNMP request packets. A community also defines which MIB objects are available and the operations (read-only or read-write) allowed on those objects.</p> <p>The SNMP client application specifies an SNMP community name in <b>Get</b>, <b>GetBulk</b>, <b>GetNext</b>, and <b>Set</b> SNMP requests.</p> |
| <b>Default</b>                  | If you omit the <b>community</b> statement, all SNMP requests are denied.                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b><i>community-name</i></b>—Community string. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring SNMP Communities</i></li></ul>                                                                                                                                                                                                                                                                                                                                        |

---

## community (SNMP RMON)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>community <i>community-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit snmp rmon event <i>index</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | The trap group that is used when generating a trap (if <b>eventType</b> is configured to send traps). If that trap group has the <b>rmon-alarm</b> trap category configured, a trap is sent to all the targets configured for that trap group. The community string in the trap matches the name of the trap group (and hence, the value of <b>eventCommunity</b> ). If nothing is configured, traps are sent to each group with the <b>rmon-alarm</b> category set. |
| <b>Options</b>                  | <b>community-name</b> —Identifies the trap group that is used when generating a trap if the event is configured to send traps.                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring an RMON Event Entry and Its Attributes</i></li></ul>                                                                                                                                                                                                                                                                                                                                                          |

## community-name (SNMP)

---

|                            |                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>community-name <i>community-name</i>;</code>                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>     | <code>[edit snmp v3 snmp-community <i>community-index</i>]</code>                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11. for the QFX Series.                                                                                                                                                      |
| <b>Description</b>         | Define an SNMP community to authorize SNMPv1 or SNMPv2c clients in an SNMPv3 system. When you configure a community in SNMPv3, you can also specify a security name. The access privileges associated with the security name determine which MIB objects are available and which operations (read, write, or notify) are allowed on those objects. |
| <b>Options</b>             | <b><i>community-name</i></b> —Community string for an SNMPv1 or SNMPv2c community. If unconfigured, it is the same as the community index. If the name includes spaces, enclose the name in quotation marks (" ").                                                                                                                                 |



**NOTE:** Community names must be unique. You cannot configure the same community name at the `[edit snmp community]` and `[edit snmp v3 snmp-community community-index]` hierarchy levels.

The community name at the `[edit snmp v3 snmp-community community-index]` hierarchy level is encrypted and not displayed in the command-line interface (CLI).

---

|                                 |                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the SNMPv3 Community</i></li></ul>                                                 |

---

## contact (SNMP)

---

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>contact <i>contact</i>;</code>                                                                                      |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches. |
| <b>Description</b>              | Define the value of the MIB II <b>sysContact</b> object, which is the contact person for the managed system.              |
| <b>Options</b>                  | <b>contact</b> —Name of the contact person. If the name includes spaces, enclose it in quotation marks (" ").             |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the System Contact on a Device Running Junos OS</i></li></ul>      |

## description (SNMP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description <i>description</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Define the value of the MIB II sysDescription object, which is the description of the system being managed.                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | By default, the sysDescription object includes the following information:<br>Juniper Networks, Inc. <i>platform</i> , <i>build</i> , Build date: <i>date</i> UTC Copyright (c) <i>date-range</i><br>Juniper Networks, Inc.<br><br>For example:<br><br>sysDescr.0 = Juniper Networks, Inc. m7i internet router, kernel JUNOS<br>13.2-20130530_ib_13_3_psd.1, Build date: 2013-05-30 22:48:07<br>UTC Copyright (c) 1996-2013 Juniper Networks, Inc. |
| <b>Options</b>                  | <i>description</i> —System description. If the name includes spaces, enclose it in quotation marks (" ").                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the System Description on a Device Running Junos OS</i></li></ul>                                                                                                                                                                                                                                                                                                                          |



## description (SNMP RMON)


|                                 |                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>description <i>description</i>;</code>                                                                                                                |
| <b>Hierarchy Level</b>          | [edit snmp rmon alarm <i>index</i> ],<br>[edit snmp rmon event <i>index</i> ]                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                   |
| <b>Description</b>              | Text description of alarm or event.                                                                                                                         |
| <b>Options</b>                  | <b><i>description</i></b> —Text description of an alarm or event entry. If the description includes spaces, enclose it in quotation marks (" ").            |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Description</i></li> <li>• <i>Configuring an RMON Event Entry and Its Attributes</i></li> </ul> |

## destination-port

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>destination-port <i>port-number</i>;</code>                                                                         |
| <b>Hierarchy Level</b>          | [edit snmp trap-group]                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches. |
| <b>Description</b>              | Assign a trap port number other than the default.                                                                         |
| <b>Default</b>                  | If you omit this statement, the default port is 162.                                                                      |
| <b>Options</b>                  | <b><i>port-number</i></b> —SNMP trap port number.                                                                         |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring SNMP Trap Groups</i></li> </ul>                                   |

## engine-id (SNMP)

---

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | engine-id {<br>(local <i>engine-id-suffix</i>   use-default-ip-address   use-mac-address);<br>}                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | [edit snmp]                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | The local engine ID is defined as the administratively unique identifier of an SNMPv3 engine, and is used for identification, not for addressing. There are two parts of an engine ID: prefix and suffix. The prefix is formatted according to the specifications defined in RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i> . You can configure the suffix here. |
| <div> <b>NOTE:</b> SNMPv3 authentication and encryption keys are generated based on the associated passwords and the engine ID. If you configure or change the engine ID, you must commit the new engine ID before you configure SNMPv3 users. Otherwise the keys generated from the configured passwords are based on the previous engine ID.</div> <p>For the engine ID, we recommend using the MAC address of the management port.</p> |                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>local <i>engine-id-suffix</i></b> —Explicit setting for the engine ID suffix.<br><b>use-default-ip-address</b> —The engine ID suffix is generated from the default IP address.<br><b>use-mac-address</b> —The SNMP engine identifier is generated from the MAC address of the management interface on the router.<br><b>Default:</b> use-default-ip-address                                                                    |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <ul style="list-style-type: none"><li>Configuring the Local Engine ID</li></ul>                                                                                                                                                                                                                                                                                                                                                   |

## event (SNMP)

|                                 |                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>event <i>index</i> {     community <i>community-name</i>;     description <i>description</i>;     type <i>type</i>; }</pre>                            |
| <b>Hierarchy Level</b>          | [edit snmp <a href="#">rmon</a> ]                                                                                                                           |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                   |
| <b>Description</b>              | Configure RMON event entries.                                                                                                                               |
| <b>Options</b>                  | <p><i>index</i>—Identifier for a specific event entry.</p> <p>The remaining statements are explained separately.</p>                                        |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring an RMON Event Entry and Its Attributes</i></li> <li>• <a href="#">alarm on page 4087</a></li> </ul> |

## falling-event-index

|                                 |                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | falling-event-index <i>index</i> ;                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit snmp rmon <a href="#">alarm</a> <i>index</i> ]                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                       |
| <b>Description</b>              | The index of the event entry that is used when a falling threshold is crossed. If this value is zero, no event is triggered.                                                    |
| <b>Options</b>                  | <p><i>index</i>—Index of the event entry that is used when a falling threshold is crossed.</p> <p><b>Range:</b> 0 through 65,535</p> <p><b>Default:</b> 0</p>                   |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Falling Event Index or Rising Event Index</i></li> <li>• <a href="#">rising-event-index on page 4120</a></li> </ul> |

## falling-threshold

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>falling-threshold <i>percentage</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit snmp ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | The lower threshold is expressed as a percentage of the maximum possible value for the sampled variable. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval is greater than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is less than or equal to this threshold. After a falling event is generated, another falling event cannot be generated until the sampled value rises above this threshold and reaches the <b>rising-threshold</b> . |
| <b>Options</b>                  | <b><i>percentage</i></b> —The lower threshold for the alarm entry.<br><b>Range:</b> 1 through 100<br><b>Default:</b> 70 percent of the maximum possible value                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Falling Threshold or Rising Threshold</i></li><li>• <a href="#">rising-threshold (SNMP Health Monitor) on page 4121</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                      |

## falling-threshold

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>falling-threshold <i>integer</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <code>[edit snmp rmon <a href="#">alarm</a> <i>index</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | The lower threshold for the sampled variable. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval is greater than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is less than or equal to this threshold, and the associated startup-alarm value is equal to falling-alarm value or rising-or-falling-alarm value. After a falling event is generated, another falling event cannot be generated until the sampled value rises above this threshold and reaches the <b>rising-threshold</b> . |
| <b>Options</b>                  | <b><i>integer</i></b> —The lower threshold for the alarm entry.<br><b>Range:</b> -2,147,483,648 through 2,147,483,647<br><b>Default:</b> 20 percent less than <b>rising-threshold</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Falling Threshold or Rising Threshold</i></li> <li>• <a href="#">rising-threshold (SNMP RMON) on page 4122</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## falling-threshold-interval

---


|                                 |                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>falling-threshold-interval seconds;</code>                                                                                                                |
| <b>Hierarchy Level</b>          | [edit snmp rmon <a href="#">alarm index</a> ]                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                           |
| <b>Description</b>              | Interval between samples when the rising threshold is crossed. Once the alarm crosses the falling threshold, the regular sampling interval is used.             |
| <b>Options</b>                  | <b>seconds</b> —Time between samples, in seconds.<br><b>Range:</b> 1 through 2,147,483,647 seconds<br><b>Default:</b> 60 seconds                                |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Falling Threshold Interval</i></li><li>• <a href="#">interval (SNMP RMON) on page 4109</a></li></ul> |

## filter-duplicates

---

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>filter-duplicates;</code>                                                                                           |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches. |
| <b>Description</b>              | Filter duplicate <b>Get</b> , <b>GetNext</b> , or <b>GetBulk</b> SNMP requests.                                           |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Filtering Duplicate SNMP Requests</i></li></ul>                                |

## filter-interfaces

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>filter-interfaces {   all-internal-interfaces;   interfaces <i>interface-names</i>{     interface 1;     interface 2;   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.4.</p> <p>Statement introduced in Junos OS Release 9.4 for EX Series Switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Filter out information related to specific interfaces from the output of SNMP <b>Get</b> and <b>GetNext</b> requests performed on interface-related MIBs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>all-internal-interfaces</b>—Filters out information from SNMP <b>Get</b> and <b>GetNext</b> requests for internal interfaces that include interfaces belonging to internal routing tables, such as <code>_juniper_private1_</code>, <code>_juniper_private2_</code>, <code>_juniper_private3_</code>, and <code>_juniper_private4_</code>. In addition, this option filters out information about the physical interface corresponding to a logical interface with a unit number from the internal routing table.</p> <p><b>interfaces</b>—Specifies the interfaces to filter out from the output of SNMP <b>Get</b> and <b>GetNext</b> requests.</p> |
|                                 | <div>  <p><b>NOTE:</b> Starting with Release 12.1, Junos OS provides an except option (! operator) that enables you to filter out all interfaces except those interfaces that match the regular expressions prefixed with the ! mark.</p> </div>                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Filtering Interface Information Out of SNMP Get and GetNext Output</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## group (Configuring Access Privileges)

---

**Syntax**    `group group-name {  
                  (default-context-prefix | context-prefix context-prefix){  
                    security-model (any | usm | v1 | v2c) {  
                      security-level (authentication | none | privacy) {  
                        notify-view view-name;  
                        read-view view-name;  
                        write-view view-name;  
                      }  
                    }  
                  }  
                  }  
                  }`

**Hierarchy Level**    [edit snmp v3 vacm access]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Assign the security name to a group, and specify the SNMPv3 context applicable to the group. The **default-context-prefix** statement, when included, adds all the contexts configured on the device to the group, whereas the **context-prefix context-prefix** statement enables you to specify a context and to add that particular context to the group.

(Not applicable to the QFX Series.) When the context prefix is specified as default (for example, **context-prefix default**), the context associated with the master routing instance is added to the group. To specify a routing instance that is part of a logical system, specify it as **logical system/routing instance**. For example, to specify routing instance ri1 in logical system ls1, include **context-prefix ls1/ri1**.

The remaining statements under this hierarchy are explained separately.

**Options**    *group-name*—SNMPv3 group name created for the SNMPv3 group.

**Required Privilege Level**    snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**    • *Configuring the Group*



## group (Associating a Security Name)

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group group-name;</code>                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 vacm security-to-group security-model (usm   v1   v2c)<br/>security-name security-name]</code>                                                                             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Associate a security name with a group composed of users with the same access privileges. The security name is used during authentication of SNMP messages, and is mapped to a username.       |
| <b>Options</b>                  | <i>group-name</i> —Collection of SNMP security names that share the same SNMPv3 access privileges.                                                                                             |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring the Group</i></li> </ul>                                                                                                                 |

## health-monitor

|                                 |                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>health-monitor {   falling-threshold percentage;   interval seconds;   rising-threshold percentage; }</pre>                          |
| <b>Hierarchy Level</b>          | <code>[edit snmp]</code>                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                     |
| <b>Description</b>              | Configure health monitoring.<br><br>The remaining statements are explained separately.                                                    |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring Health Monitoring on Devices Running Junos OS</i></li> </ul>                        |

## history

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>history <i>history-index</i> {<br/>    <i>bucket-size number</i>;<br/>    <b>interface (SNMP RMON History)</b> <i>interface-name</i>;<br/>    <i>interval seconds</i>;<br/>    <b>owner</b> <i>owner-name</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <b>snmp rmon</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure RMON history group entries. This RMON feature can be used with the Simple Network Management Protocol (SNMP) agent in the switch to monitor all the traffic flowing among switches on all connected LAN segments. It collects statistics in accordance with user-configurable parameters.</p> <p>The history group controls the periodic statistical sampling of data from various types of networks. This group contains configuration entries that specify an interface, polling period, and other parameters. The <b>interface (SNMP RMON History)</b> <i>interface-name</i> statement is mandatory. Other statements in the history group are optional.</p> |
| <b>Default</b>                  | Not configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>history-index</b>—Identifies this history entry as an integer.</p> <p><b>Range:</b> 1 through 65535</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP (J-Web Procedure) on page 4040</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## interface (SNMP RMON History)

|                                 |                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interface <i>interface-name</i>;</code>                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit <a href="#">snmp rmon history</a> <i>history-index</i> ]                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Specify the interface to be monitored in the specified RMON history entry.<br><br>Only one interface can be specified for a particular RMON history index. There is a one-to-one relationship between the interface and the history index. The interface must be specified in order for the RMON history to be created. |
| <b>Options</b>                  | <i>interface-name</i> —Specify the interface to be monitored within the specified entry of the RMON history of Ethernet statistics.                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring SNMP (J-Web Procedure) on page 4040</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>                                                                                                                          |

## interface (SNMP)

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interface [ <i>interface-names</i> ];</code>                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the interfaces on which SNMP requests can be accepted.                                                                                                                               |
| <b>Default</b>                  | If you omit this statement, SNMP requests entering the router or switch through any interface are accepted.                                                                                    |
| <b>Options</b>                  | <i>interface-names</i> —Names of one or more logical interfaces.                                                                                                                               |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Interfaces on Which SNMP Requests Can Be Accepted</a></li> </ul>                                                          |

## interval

---

|                            |                                                                                      |
|----------------------------|--------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>interval seconds;</code>                                                       |
| <b>Hierarchy Level</b>     | [edit <a href="#">snmp rmon history</a> ]                                            |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.0 for EX Series switches.                 |
| <b>Description</b>         | Configure the interval over which data is to be sampled for the specified interface. |
| <b>Default</b>             | 1800 sec                                                                             |
| <b>Options</b>             | <i>seconds</i> —Interval at which data is to be sampled for the specified interface. |
| <b>Required Privilege</b>  | snmp—To view this statement in the configuration.                                    |
| <b>Level</b>               | snmp-control—To add this statement to the configuration.                             |

## interval (SNMP Health Monitor)

---

|                              |                                                                                                                                |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                | <code>interval seconds;</code>                                                                                                 |
| <b>Hierarchy Level</b>       | [edit <code>snmp health-monitor</code> ]                                                                                       |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.          |
| <b>Description</b>           | Interval between samples.                                                                                                      |
| <b>Options</b>               | <i>seconds</i> —Time between samples, in seconds.<br><b>Range:</b> 1 through 2147483647 seconds<br><b>Default:</b> 300 seconds |
| <b>Required Privilege</b>    | snmp—To view this statement in the configuration.                                                                              |
| <b>Level</b>                 | snmp-control—To add this statement to the configuration.                                                                       |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Configuring the Interval</i></li></ul>                                              |

## interval (SNMP RMON)

---

|                                 |                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interval seconds;</code>                                                                                                   |
| <b>Hierarchy Level</b>          | [edit snmp rmon <a href="#">alarm index</a> ]                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.        |
| <b>Description</b>              | Interval between samples.                                                                                                        |
| <b>Options</b>                  | <b>seconds</b> —Time between samples, in seconds.<br><b>Range:</b> 1 through 2,147,483,647 seconds<br><b>Default:</b> 60 seconds |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Interval</i></li> </ul>                                              |

## location (SNMP)

---

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>location location;</code>                                                                                           |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches. |
| <b>Description</b>              | Define the value of the MIB II <b>sysLocation</b> object, which is the physical location of the managed system.           |
| <b>Options</b>                  | <b>location</b> —Location of the local system. You must enclose the name within quotation marks (" ").                    |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the System Location for a Device Running Junos OS</i></li> </ul>  |

## logical-system (SNMP)

---

|                     |                                                                                                                                                                                                            |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax              | <code>logical-system <i>logical-system-name</i> {<br/>    <i>routing-instance routing-instance-name</i>;<br/>}</code>                                                                                      |
| Hierarchy Level     | <code>[edit snmp <b>community</b> <i>community-name</i>],<br/>[edit snmp <b>trap-group</b>],<br/>[edit snmp <b>trap-options</b>]<br/>[edit snmp <b>v3target-address</b> <i>target-address-name</i>]</code> |
| Release Information | Statement introduced in Junos OS Release 9.3<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                  |



**NOTE:** The `logical-system` statement replaces the `logical-router` statement, and is backward-compatible with Junos OS Release 8.3 and later.

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Description              | <p>Specify a logical system name for SNMP v1 and v2c clients.</p> <p>Include at the <code>[edit snmp trap-options]</code> hierarchy level to specify a logical-system address as the source address of an SNMP trap.</p> <p>Include at the <code>[edit snmp v3 target-address]</code> hierarchy level to specify a logical-system name as the destination address for an SNMPv3 trap or inform.</p> |
| Options                  | <p><b><i>logical-system-name</i></b>—Name of the logical system.</p> <p><b><i>routing-instance routing-instance-name</i></b>—Statement to specify a routing instance associated with the logical system.</p>                                                                                                                                                                                        |
| Required Privilege Level | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                            |
| Related Documentation    | <ul style="list-style-type: none"><li><i>Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community</i></li><li><i>Configuring the Trap Target Address</i></li></ul>                                                                                                                                                                                                                           |

## message-processing-model

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>message-processing-model (v1   v2c   v3);</code>                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 target-parameters <i>target-parameter-name</i> parameters]</code>                                                                                                          |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the message processing model to be used when generating SNMP notifications.                                                                                                          |
| <b>Options</b>                  | <b>v1</b> —SNMPv1 message process model.<br><b>v2c</b> —SNMPv2c message process model.<br><b>v3</b> —SNMPv3 message process model.                                                             |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Message Processing Model</i></li> </ul>                                                                                            |

## name

|                                 |                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>name <i>name</i>;</code>                                                                                                |
| <b>Hierarchy Level</b>          | <code>[edit snmp]</code>                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.     |
| <b>Description</b>              | Set the system name from the command-line interface.                                                                          |
| <b>Options</b>                  | <b><i>name</i></b> —System name override.                                                                                     |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the System Name</i></li> </ul>                                        |

## nonvolatile

---

|                                 |                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>nonvolatile {<br/>    <a href="#">commit-delay</a> <i>seconds</i>;<br/>}</code>                                                                      |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>The <a href="#">commit-delay</a> statement introduced in Junos OS Release 9.0 for EX Series switches. |
| <b>Description</b>              | Configure options for SNMP <b>Set</b> requests.<br><br>The statement is explained separately.                                                              |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Commit Delay Timer</i></li><li>• <a href="#">commit-delay on page 4091</a></li></ul>            |



## notify

---

|                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> notify <i>name</i> {     tag <i>tag-name</i>;     type (trap   inform); } </pre>                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit snmp v3]                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>type inform</b> option added in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                     |
| <b>Description</b>              | Select management targets for SNMPv3 notifications as well as the type of notifications. Notifications can be either traps or informs.                                                                                                                                                                            |
| <b>Options</b>                  | <p><b><i>name</i></b>—Name assigned to the notification.</p> <p><b><i>tag-name</i></b>—Notifications are sent to all targets configured with this tag.</p> <p><b><i>type</i></b>—Notification type is <b>trap</b> or <b>inform</b>. Traps are unconfirmed notifications. Informs are confirmed notifications.</p> |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Inform Notification Type and Target Address</i></li> <li>• <i>Configuring the SNMPv3 Trap Notification</i></li> </ul>                                                                                                                                 |

## notify-filter (Applying to the Management Target)

---

|                                 |                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>notify-filter <i>profile-name</i>;</code>                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 <a href="#">target-parameters</a> <i>target-parameters-name</i>]</code>                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.           |
| <b>Description</b>              | Configure the notify filter applied to a specific set of SNMPv3 target parameters. Target parameters are the message processing and security parameters for notifications sent to a target SNMP manager. |
| <b>Options</b>                  | <i>profile-name</i> —Name of the notify filter to apply to notifications.                                                                                                                                |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Applying the Trap Notification Filter</a></li></ul>                                                                                                  |

## notify-filter (Configuring the Profile Name)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>notify-filter <i>profile-name</i> {<br/>    <a href="#">oid</a> <i>oid</i> (include   exclude);<br/>}</code>                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3]</code>                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Specify a group of MIB objects for which you define access. The notify filter limits the type of traps or informs sent to the network management system.                                       |
| <b>Options</b>                  | <i>profile-name</i> —Name assigned to the notify filter.<br><br>The remaining statement is explained separately.                                                                               |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Trap Notification Filter</a></li><li>• <a href="#">oid (SNMP) on page 4115</a></li></ul>                                   |

## notify-view

|                                 |                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>notify-view <i>view-name</i>;</code>                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i>) security-model (any   usm   v1   v2c) security-level (authentication   none   privacy)]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                         |
| <b>Description</b>              | Associate the notify view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).                                                                                                       |
| <b>Options</b>                  | <b><i>view-name</i></b> —Name of the view to which the SNMP user group has access.                                                                                                                                     |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring MIB Views</i></li> <li>• <i>Configuring the Notify View</i></li> </ul>                                                                                         |

## oid (SNMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>oid <i>oid</i> (include   exclude);</code>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 notify-filter <i>profile-name</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Specify an object identifier (OID) used to represent a subtree of MIB objects. This OID is a prefix that the represented MIB objects have in common.                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b><code>exclude</code></b> —Exclude the subtree of MIB objects represented by the specified OID.<br><br><b><code>include</code></b> —Include the subtree of MIB objects represented by the specified OID.<br><br><b><i>oid</i></b> —Object identifier used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name. |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Trap Notification Filter</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                        |

## oid

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>oid <i>object-identifier</i> (exclude   include);</code>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit snmp view <i>view-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Specify an object identifier (OID) used to represent a subtree of MIB objects.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>exclude</b> —Exclude the subtree of MIB objects represented by the specified OID.<br><b>include</b> —Include the subtree of MIB objects represented by the specified OID.<br><b><i>object-identifier</i></b> —OID used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring MIB Views</a></li></ul>                                                                                                                                                                                                                                                                                                                                                  |

## owner

---

|                                 |                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>owner <i>owner-name</i>;</code>                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">snmp rmon history</a> ]                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                        |
| <b>Description</b>              | Specify the user or group responsible for this configuration.                                                                                                                               |
| <b>Options</b>                  | <b><i>owner-name</i></b> —The user or group responsible for this configuration.<br><b>Range:</b> 0 through 32 alphanumeric characters                                                       |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP (J-Web Procedure) on page 4040</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul> |

## parameters

---

|                                 |                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>parameters {   message-processing-model (v1   v2c   v3);   security-level (none   authentication   privacy);   security-model (usm   v1   v2c);   security-name security-name; }</pre>                   |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-parameters <i>target-parameters-name</i> ]                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> |
| <b>Description</b>              | <p>Configure a set of target parameters for message processing and security.</p> <p>The remaining statements are explained separately.</p>                                                                    |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Defining and Configuring the Trap Target Parameters</i></li> </ul>                                                                                                |

## port (SNMP)

---

|                                 |                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>port port-number;</code>                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-address <i>target-address-name</i> ]                                                                                                                                                     |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> |
| <b>Description</b>              | Configure a UDP port number for an SNMP target.                                                                                                                                                               |
| <b>Default</b>                  | If you omit this statement, the default port is 162.                                                                                                                                                          |
| <b>Options</b>                  | <i>port-number</i> —Port number for the SNMP target.                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Port</i></li> </ul>                                                                                                                               |

## read-view

---

|                                 |                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>read-view <i>view-name</i>;</code>                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [ <code>edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i>) security-model (any   usm   v1   v2c) security-level (authentication   none   privacy)</code> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                           |
| <b>Description</b>              | Associate the read-only view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).                                                                                                      |
| <b>Options</b>                  | <b><i>view-name</i></b> —The name of the view to which the SNMP user group has access.                                                                                                                                   |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Read View</i></li><li>• <i>Configuring MIB Views</i></li></ul>                                                                                                |

## request-type

---

|                                 |                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | request-type (get-next-request   get-request   walk-request);                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit snmp rmon <a href="#">alarm index</a> ]                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                        |
| <b>Description</b>              | Extend monitoring to a specific SNMP object instance ( <b>get-request</b> ), or extend monitoring to all object instances belonging to a MIB branch ( <b>walk-request</b> ), or extend monitoring to the next object instance after the instance specified in the configuration ( <b>get-next-request</b> ). |
| <b>Options</b>                  | <p><b>get-next-request</b>—Performs an SNMP get next request.</p> <p><b>get-request</b>—Performs an SNMP get request.</p> <p><b>walk-request</b>—Performs an SNMP walk request.</p> <p><b>Default:</b> walk-request</p>                                                                                      |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Request Type</i></li> <li>• <a href="#">variable on page 4152</a></li> </ul>                                                                                                                                                                     |

## rising-event-index

---

|                                 |                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rising-event-index <i>index</i>;</code>                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit snmp rmon <a href="#">alarm index</a> ]                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                     |
| <b>Description</b>              | Index of the event entry that is used when a rising threshold is crossed. If this value is zero, no event is triggered.                                                       |
| <b>Options</b>                  | <b><i>index</i></b> —Index of the event entry that is used when a rising threshold is crossed.<br><b>Range:</b> 0 through 65,535<br><b>Default:</b> 0                         |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Falling Event Index or Rising Event Index</i></li><li>• <a href="#">falling-event-index on page 4099</a></li></ul> |



## rising-threshold (SNMP Health Monitor)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rising-threshold <i>percentage</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <code>[edit snmp ]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | The upper threshold is expressed as a percentage of the maximum possible value for the sampled variable. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval is less than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is greater than or equal to this threshold. After a rising event is generated, another rising event cannot be generated until the sampled value falls below this threshold and reaches the <b>falling-threshold</b> . |
| <b>Options</b>                  | <b><i>percentage</i></b> —The lower threshold for the alarm entry.<br><b>Range:</b> 1 through 100<br><b>Default:</b> 80 percent of the maximum possible value                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">falling-threshold on page 4100</a></li> <li>• <i>Configuring the Falling Threshold or Rising Threshold</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                          |

## rising-threshold (SNMP RMON)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | rising-threshold <i>integer</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit snmp rmon <a href="#">alarm index</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Upper threshold for the sampled variable. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval is less than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is greater than or equal to this threshold, and the associated startup alarm value is equal to the falling alarm or rising or falling alarm value. After a rising event is generated, another rising event cannot be generated until the sampled value falls below this threshold and reaches the falling threshold. |
| <b>Options</b>                  | <i>integer</i> —The lower threshold for the alarm entry.<br><b>Range:</b> -2,147,483,648 through 2,147,483,647                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Falling Threshold or Rising Threshold</i></li><li>• <a href="#">falling-threshold on page 4101</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## rmon

---

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | rmon { ... }                                                                                                              |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches. |
| <b>Description</b>              | Configure Remote Monitoring.                                                                                              |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring an RMON Alarm Entry and Its Attributes</i></li></ul>               |

## rmon

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> rmon {     history <i>history-index</i> {         interface (SNMP RMON History) <i>interface-name</i>;         bucket-size <i>number</i>;         interval <i>seconds</i>;         owner <i>owner-name</i>;     } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">snmp</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>RMON is an existing feature of Junos OS.</p> <p>The RMON specification provides network administrators with comprehensive network fault diagnosis, planning, and performance tuning information. It delivers this information in nine groups of monitoring elements, each providing specific sets of data to meet common network monitoring requirements. Each group is optional, so that vendors do not need to support all the groups within the MIB.</p> <p>Junos OS supports RMON Statistics, History, Alarm, and Event groups. The EX Series documentation describes only the <b>rmon history</b> statement, which was added with this release.</p> <p>The statements are explained separately.</p> |
| <b>Default</b>                  | Disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring SNMP (J-Web Procedure) on page 4040</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## routing-instance (SNMP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>routing-instance <i>routing-instance-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit snmp <b>community</b> <i>community-name</i>],</code><br><code>[edit snmp <b>community</b> <i>community-name</i> logical-system <i>logical-system-name</i>],</code><br><code>[edit snmp <b>trap-group</b> <i>group</i>]</code>                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Added to the <code>[edit snmp <b>community</b> <i>community-name</i>]</code> hierarchy level in Junos OS Release 8.4.<br>Added to the <code>[edit snmp <b>community</b> <i>community-name</i> logical-system <i>logical-system-name</i>]</code> hierarchy level in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                 |
| <b>Description</b>              | <p>Specify a routing instance for SNMPv1 and SNMPv2 trap targets. All targets configured in the trap group use this routing instance.</p> <p>If the routing instance is defined within a logical system, include the <b>logical-system</b> <i>logical-system-name</i> statement at the <code>[edit snmp <b>community</b> <i>community-name</i>]</code> hierarchy level and specify the <b>routing-instance</b> statement under the <code>[edit snmp <b>community</b> <i>community-name</i> logical-system <i>logical system-name</i>]</code> hierarchy level.</p> |
| <b>Options</b>                  | <i>routing-instance-name</i> —Name of the routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring SNMP Trap Groups</i></li><li>• <i>Configuring the Source Address for SNMP Traps</i></li><li>• <i>Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community</i></li></ul>                                                                                                                                                                                                                                                                                                                             |

## routing-instance (SNMPv3)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>routing-instance <i>routing-instance-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 <a href="#">target-address</a> <i>target-address-name</i>]</code>                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify a routing instance for an SNMPv3 trap target.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b><i>routing-instance-name</i></b>—Name of the routing instance.</p> <p>To configure a routing instance within a logical system, specify the logical system name followed by the routing instance name. Use a slash ( / ) to separate the two names (for example, <b>test-ls/test-ri</b>). To configure the default routing instance on a logical system, specify the logical system name followed by <b>default</b> (for example, <b>test-ls/default</b>).</p> |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring the Trap Target Address</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                        |

## sample-type

|                                 |                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>sample-type (absolute-value   delta-value);</code>                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit snmp rmon <a href="#">alarm</a> <i>index</i>]</code>                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                      |
| <b>Description</b>              | Method of sampling the selected variable.                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>absolute-value</b>—Actual value of the selected variable is used when comparing against the thresholds.</p> <p><b>delta-value</b>—Difference between samples of the selected variable is used when comparing against the thresholds.</p> |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring the Sample Type</i></li> </ul>                                                                                                                                                           |

## security-level (Defining Access Privileges)

---

|                                 |                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>security-level (authentication   none   privacy) {<br/>    <a href="#">notify-view</a> <i>view-name</i>;<br/>    <a href="#">read-view</a> <i>view-name</i>;<br/>    <a href="#">write-view</a> <i>view-name</i>;<br/>}</code> |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i>) security-model (any   usm   v1   v2c)]</code>                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                       |
| <b>Description</b>              | Define the security level used for access privileges.                                                                                                                                                                                |
| <b>Default</b>                  | <code>none</code>                                                                                                                                                                                                                    |
| <b>Options</b>                  | <code>authentication</code> —Provide authentication but no encryption.<br><br><code>none</code> —No authentication and no encryption.<br><br><code>privacy</code> —Provide authentication and encryption.                            |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration.                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Security Level</i></li></ul>                                                                                                                                              |

## security-level (Generating SNMP Notifications)

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>security-level (authentication   none   privacy);</code>                                                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]</code>                                                                                                         |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the security level to use when generating SNMP notifications.                                                                                                                        |
| <b>Default</b>                  | <code>none</code>                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>authentication</b>—Provide authentication but no encryption.</p> <p><b>none</b>—No authentication and no encryption.</p> <p><b>privacy</b>—Provide authentication and encryption.</p>    |
| <b>Required Privilege Level</b> | <p><code>snmp</code>—To view this statement in the configuration.</p> <p><code>snmp-control</code>—To add this statement to the configuration.</p>                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Security Level</i></li> </ul>                                                                                                      |

## security-model (Access Privileges)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | security-model (usm   v1   v2c);                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i> )]                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the security model for an SNMPv3 group. The security model is used to determine access privileges for the group.                                                                     |
| <b>Options</b>                  | <b>usm</b> —SNMPv3 security model.<br><br><b>v1</b> —SNMPv1 security model.<br><br><b>v2c</b> —SNMPv2c security model.                                                                         |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Security Model</i></li></ul>                                                                                                        |



## security-model (Group)

|                                 |                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>security-model (usm   v1   v2c) {   security-name security-name {     group group-name;   } }</pre>                                                                                                      |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm <a href="#">security-to-group</a> ]                                                                                                                                                        |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> |
| <b>Description</b>              | Define a security model for an SNMPv3 group and associate the security name of a user with the group. All users in the group have the same access privileges.                                                 |
| <b>Options</b>                  | <p><b>usm</b>—SNMPv3 security model.</p> <p><b>v1</b>—SNMPv1 security model.</p> <p><b>v2c</b>—SNMPv2c security model.</p>                                                                                    |
| <b>Required Privilege Level</b> | <p><b>snmp</b>—To view this statement in the configuration.</p> <p><b>snmp-control</b>—To add this statement to the configuration.</p>                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring the Security Model</i></li> </ul>                                                                                                                       |

## security-model (SNMP Notifications)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | security-model (usm   v1   v2c);                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the security model for an SNMPv3 group. The security model is used for SNMP notifications.                                                                                           |
| <b>Options</b>                  | <b>usm</b> —SNMPv3 security model.<br><br><b>v1</b> —SNMPv1 security model.<br><br><b>v2c</b> —SNMPv2c security model.                                                                         |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Security Model</i></li></ul>                                                                                                        |

## security-name (Security Group)

|                                 |                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>security-name <i>security-name</i> {<br/>    group <i>group-name</i>;<br/>}</code>                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm security-to-group <b>security-model</b> (usm   v1   v2c)]                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                              |
| <b>Description</b>              | Associate the security name of a user (for SNMPv3 clients) or a community string (for SNMPv1 and SNMPv2c clients) with a configured security group.                                                                                                                                                                                         |
| <b>Options</b>                  | <b>security-name</b> —SNMPv3 secure username configured at the [edit snmp v3 usm local-engine <b>user <i>username</i></b> ] hierarchy level that is used for messaging security. For SNMPv1 and SNMPv2c, the security name is the community string configured at the [edit snmp v3 snmp-community <b>community-index</b> ] hierarchy level. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Assigning Security Names to Groups</i></li> <li>• <i>Assigning a Security Name to a Group</i></li> </ul>                                                                                                                                                                                        |

## security-name (Community String)

---


|                            |                                                                                                                                                                                                                                                                |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>security-name <i>security-name</i>;</code>                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | <code>[edit snmp v3 <i>snmp-community</i> <i>community-index</i>]</code>                                                                                                                                                                                       |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                 |
| <b>Description</b>         | Associate a community string with the security name of a user. The community string, which is used for SNMPv1 and SNMPv2c clients in an SNMPv3 system, is configured at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level. |
| <b>Options</b>             | <i>security-name</i> —Name that is used for messaging security and user access control.                                                                                                                                                                        |



**NOTE:** The security name must match the configured security name at the `[edit snmp v3 target-parameters target-parameters-name parameters]` hierarchy level when you configure traps or informs.

|                              |                                                                                         |
|------------------------------|-----------------------------------------------------------------------------------------|
| <b>Required Privilege</b>    | snmp—To view this statement in the configuration.                                       |
| <b>Level</b>                 | snmp-control—To add this statement to the configuration.                                |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Configuring the Security Names</i></li></ul> |

## security-name (SNMP Notifications)

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                               |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <code>security-name <i>security-name</i>;</code>                                                                                                                                                                                              |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <code>[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]</code>                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Configure the security name used when generating SNMP notifications.                                                                                                                                                                          |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b><i>security-name</i></b> —If the SNMPv3 USM security model is used, identify the user when generating the SNMP notification. If the v1 or v2c security models are used, identify the SNMP community used when generating the notification. |
| <div>  <p><b>NOTE:</b> The access privileges for the group associated with this security name must allow this notification to be sent.</p> <p>If you are using the v1 or v2 security models, the security name at the <code>[edit snmp v3 vacm security-to-group]</code> hierarchy level must match the security name at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level.</p> </div> |                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <ul style="list-style-type: none"> <li><i>Configuring the Security Name</i></li> </ul>                                                                                                                                                        |

## security-to-group

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>security-to-group {<br/>    security-model (usm   v1   v2c) {<br/>        group group-name;<br/>        security-name security-name;<br/>    }<br/>}</pre>                                |
| <b>Hierarchy Level</b>          | [edit snmp v3 vacm]                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the group to which a specific SNMPv3 security name belongs. The security name is used for messaging security.<br><br>The remaining statements are explained separately.              |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Assigning Security Model and Security Name to a Group</i></li></ul>                                                                                 |

## snmp

---

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>snmp { ... }</pre>                                                                                                   |
| <b>Hierarchy Level</b>          | [edit]                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches. |
| <b>Description</b>              | Configure SNMP.<br><br>SNMP modules cannot have the slash (/) character or the @ character in the name.                   |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring SNMP on a Device Running Junos OS</i></li></ul>                    |

## snmp

```
Syntax  snmp {
        rmon {
            history index {
                interface (SNMP RMON History) interface-name;
                bucket-size number;
                interval seconds;
                owner owner-name;
            }
        }
    }
```

Hierarchy Level [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure SNMP.

The statements are explained separately.

**Required Privilege Level** snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring SNMP \(J-Web Procedure\) on page 4040](#)

## snmp-community

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>snmp-community <i>community-index</i> {<br/>    <i>community-name</i> <i>community-name</i>;<br/>    <i>security-name</i> <i>security-name</i>;<br/>    tag <i>tag-name</i>;<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit snmp v3]                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the SNMP community which authorizes SNMPv1 or SNMPv2c clients in an SNMPv3 system.                                                                                                   |
| <b>Options</b>                  | <i>community-index</i> —(Optional) String that identifies an SNMP community.<br><br>The remaining statements are explained separately.                                                         |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the SNMPv3 Community</i></li></ul>                                                                                                      |



---

## source-address (SNMP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | source-address <i>address</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit snmp trap-options]                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Set the source address of every SNMP trap packet sent by this router to a single address regardless of the outgoing interface. If the source address is not specified, the default is to use the address of the outgoing interface as the source address.                                                                                                                                                                                          |
| <b>Options</b>                  | <b>address</b> —Source address of SNMP traps. You can configure the source address of trap packets two ways: <b>lo0</b> or a valid IPv4 address configured on one of the router interfaces. The value <b>lo0</b> indicates that the source address of all SNMP trap packets is set to the lowest loopback address configured at interface <b>lo0</b> .<br><b>Default:</b> Disabled. (The source address is the address of the outgoing interface.) |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Source Address for SNMP Traps</i></li></ul>                                                                                                                                                                                                                                                                                                                                             |

## startup-alarm

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | startup-alarm (falling-alarm   rising-alarm   rising-or-falling-alarm);                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit snmp rmon <a href="#">alarm</a> index]                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | The alarm that can be sent upon entry startup.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>falling-alarm</b>—Generated if the first sample after the alarm entry becomes active is less than or equal to the falling threshold.</p> <p><b>rising-alarm</b>—Generated if the first sample after the alarm entry becomes active is greater than or equal to the rising threshold.</p> <p><b>rising-or-falling-alarm</b>—Generated if the first sample after the alarm entry becomes active satisfies either of the corresponding thresholds.</p> <p><b>Default:</b> rising-or-falling-alarm</p> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Sample Type</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                     |

## syslog-subtag

---

|                                 |                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | syslog-subtag <i>syslog-subtag</i> ;                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit snmp rmon <a href="#">alarm</a> index]                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                     |
| <b>Description</b>              | Add a tag to the system log message.                                                                                                                      |
| <b>Options</b>                  | <p><b>syslog-subtag <i>syslog-subtag</i></b>—Tag of not more than 80 uppercase characters to be added to syslog messages.</p> <p><b>Default:</b> None</p> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the System Log Tag</i></li></ul>                                                                   |

## tag (SNMPv3)

|                                 |                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>tag tag-name;</code>                                                                                                                |
| <b>Hierarchy Level</b>          | [edit snmp v3 <b>notify</b> name],<br>[edit snmp v3 <b>snmp-community</b> community-index]                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                 |
| <b>Description</b>              | Configure a set of targets to receive traps or informs (for IPv4 packets only).                                                           |
| <b>Options</b>                  | <b>tag-name</b> —Identifies the address of managers that are allowed to use a community string.                                           |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Tag</i></li> <li>• <i>Configuring the SNMPv3 Trap Notification</i></li> </ul> |

## tag-list

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>tag-list tag-list;</code>                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit snmp v3 target-address <i>target-address-name</i> ]                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure an SNMP tag list used to select target addresses.                                                                                                                                    |
| <b>Options</b>                  | <b>tag-list</b> —Define sets of target addresses (tags). To specify more than one tag, specify the tag names as a space-separated list enclosed within double quotes.                          |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Trap Target Address</i></li> </ul>                                                                                                 |

## target-address

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>target-address <i>target-address-name</i> {<br/>    address <i>address</i>;<br/>    address-mask <i>address-mask</i>;<br/>    logical-system (SNMP) <i>logical-system</i>;<br/>    port (SNMP) <i>port-number</i>;<br/>    retry-count (SNMPv3) <i>number</i>;<br/>    routing-instance (SNMPv3) <i>instance</i>;<br/>    tag-list <i>tag-list</i>;<br/>    target-parameters <i>target-parameters-name</i>;<br/>    timeout <i>seconds</i>;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit snmp v3]                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure the address of an SNMP management application and the parameters to be used in sending notifications.                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b><i>target-address-name</i></b>—String that identifies the target address.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>Configuring the Trap Target Address</li></ul>                                                                                                                                                                                                                                                                                                                                                                           |

## target-parameters

**Syntax** At the `[edit snmp v3]` hierarchy level:

```
target-parameters target-parameters-name {
  profile-name;
  parameters {
    message-processing-model (v1 | v2c | V3);
    security-level (authentication | none | privacy);
    security-model (usm | v1 | v2c);
    security-name security-name;
  }
}
```

At the `[edit snmp v3 target-address target-address-name]` hierarchy level:

```
target-parameters target-parameters-name;
```

**Hierarchy Level** `[edit snmp v3]`  
`[edit snmp v3 target-address target-address-name]`

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the message processing and security parameters for sending notifications to a particular management target. The target parameters are configured at the `[edit snmp v3]` hierarchy level. The remaining statements at this level are explained separately.

Then apply the target parameters configured at the `[edit snmp v3 target-parameters target-parameters-name]` hierarchy level to the target address configuration at the `[edit snmp v3]` hierarchy level.

**Required Privilege Level** snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**

- *Defining and Configuring the Trap Target Parameters*
- *Applying Target Parameters*

## targets

---

|                                 |                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>targets {<br/>    address;<br/>}</code>                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit snmp trap-group group-name]</code>                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                 |
| <b>Description</b>              | Configure one or more systems to receive SNMP traps.                                                                                      |
| <b>Options</b>                  | <b>address</b> —IPv4 or IPv6 address of the system to receive traps. You must specify an address, not a hostname.                         |
| <b>Required Privilege Level</b> | <code>snmp</code> —To view this statement in the configuration.<br><code>snmp-control</code> —To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring SNMP Trap Groups</i></li></ul>                                                     |

## traceoptions (SNMP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i>;     memory-trace;     no-remote-trace;     no-default-memory-trace; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | [edit snmp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>file <i>filename</i></b> option added in Junos OS Release 8.1.</p> <p><b>world-readable   no-world-readable</b> option added in Junos OS Release 8.1.</p> <p><b>match <i>regular-expression</i></b> option added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>memory-trace</b> and <b>no-default-memory-trace</b> options introduced in Junos OS Release 13.3.</p>                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | <p>The output of the tracing operations is placed into log files in the <b>/var/log</b> directory. Each log file is named after the SNMP agent that generates it. Currently, the following logs are created in the <b>/var/log</b> directory when the <b>traceoptions</b> statement is used:</p> <ul style="list-style-type: none"> <li>• chassisd</li> <li>• craftd</li> <li>• ilmid</li> <li>• mib2d</li> <li>• rmopd</li> <li>• serviced</li> <li>• snmpd</li> </ul>                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—By default, the name of the log file that records trace output is the name of the process being traced (for example, <b>mib2d</b> or <b>snmpd</b>). Use this option to specify another name.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files per SNMP subagent. When a trace file (for example, <b>snmpd</b>) reaches its maximum size, it is archived by being renamed to <b>snmpd.0</b>. The previous <b>snmpd.1</b> is renamed to <b>snmpd.2</b>, and so on. The oldest archived file is deleted.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 10 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements:</p> |

- **all**—Log all SNMP events.
- **general**—Log general events.
- **interface-stats**—Log physical and logical interface statistics.
- **nonvolatile-sets**—Log nonvolatile SNMP set request handling.
- **pdu**—Log SNMP request and response packets.
- **protocol-timeouts**—Log SNMP response timeouts.
- **routing-socket**—Log routing socket calls.
- **subagent**—Log subagent restarts.
- **timer**—Log internally generated events.
- **varbind-error**—Log variable binding errors.

**memory-trace**—Enable tracing in memory. If **file** and **memory-trace** are configured, tracing will occur in file and memory.



**NOTE:** By default, memory tracing is enabled with default trace flags even after **snmp traceoptions** config statement is disabled.

---

**match *regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

**no-default-memory-trace**—(Optional) Disable default memory-tracing.

**size *size***—(Optional) Maximum size, in kilobytes (KB), of each trace file before it is closed and archived.

**Range:** 10 KB through 1 GB

**Default:** 1000 KB

**world-readable | no-world-readable**—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

**Required Privilege Level**    **snmp**—To view this statement in the configuration.  
                                      **snmp-control**—To add this statement to the configuration.

**Related Documentation**    • *Tracing SNMP Activity on a Device Running Junos OS*



## trap-group

|                                 |                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> trap-group <i>group-name</i> {   categories {     category;   }   destination-port <i>port-number</i>;   routing-instance <i>instance</i>;   targets {     address;   }   version (all   v1   v2); } </pre>                                                                        |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                     |
| <b>Description</b>              | Create a named group of hosts to receive the specified trap notifications. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name. At least one trap group must be configured for SNMP traps to be sent. |
| <b>Options</b>                  | <p><b><i>group-name</i></b>—Name of the trap group. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>                                                                                                      |
| <b>Required Privilege Level</b> | <p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Configuring SNMP Trap Groups</li> </ul>                                                                                                                                                                                                           |

## trap-options

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>trap-options {<br/>    agent-address outgoing-interface;<br/>    source-address address;<br/>}</pre>                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit snmp]                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router or switch to a single address, regardless of the outgoing interface. In addition, you can set the agent address of each SNMPv1 trap. For more information about the contents of SNMPv1 traps, see RFC 1157.</p> <p>The remaining statements are explained separately.</p> |
| <b>Default</b>                  | Disabled                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring SNMP Trap Options</i></li></ul>                                                                                                                                                                                                                                                                                            |

---

## type (SNMP RMON)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>type type;</code>                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit snmp rmon <a href="#">event index</a> ]                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                      |
| <b>Description</b>              | Type of notification generated when a threshold is crossed.                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>type</b>—Type of notification:</p> <ul style="list-style-type: none"><li>• <b>log</b>—Add an entry to <b>logTable</b>.</li><li>• <b>log-and-trap</b>—Send an SNMP trap and make a log entry.</li><li>• <b>none</b>—No notifications are sent.</li><li>• <b>snmptrap</b>—Send an SNMP trap.</li></ul> <p><b>Default:</b> log-and-trap</p> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring an RMON Event Entry and Its Attributes</i></li></ul>                                                                                                                                                                                                                                    |

## type (SNMPv3)

---

|                                 |                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | type (inform   trap);                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit snmp v3 notify <i>name</i> ]                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br><b>inform</b> option added in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the type of SNMP notification.                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>inform</b> —Defines the type of notification as an inform. SNMP informs are confirmed notifications.<br><br><b>trap</b> —Defines the type of notification as a trap. SNMP traps are unconfirmed notifications.                                     |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring SNMP Informs</i></li><li>• <i>Configuring the SNMPv3 Trap Notification</i></li></ul>                                                                                                           |

## v3

```

Syntax v3 {
    notify name {
        tag tag-name;
        type trap;
    }
    notify-filter profile-name {
        oid object-identifier (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        logical-system (SNMP) logical-system;
        port port-number;
        retry-count number;
        routing-instance instance;
        tag-list tag-list;
        target-parameters target-parameters-name;
        timeout seconds;
    }
    target-parameters target-parameters-name {
        notify-filter profile-name;
        parameters {
            message-processing-model (v1 | v2c | V3);
            security-level (authentication | none | privacy);
            security-model (usm | v1 | v2c);
            security-name security-name;
        }
    }
    usm {
        local-engine {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-sha {
                    authentication-password authentication-password;
                }
                authentication-none;
                privacy-aes128 {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
            }
        }
    }
}

```

```
        privacy-none;
    }
}
remote-engine engine-id {
    user username {
        authentication-md5 {
            authentication-password authentication-password;
        }
        authentication-sha {
            authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
            privacy-password privacy-password;
        }
        privacy-des {
            privacy-password privacy-password;
        }
        privacy-3des {
            privacy-password privacy-password;
        }
        privacy-none {
            privacy-password privacy-password;
        }
    }
}
}
}
vacm {
    access {
        group group-name {
            (default-context-prefix | context-prefix context-prefix) {
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
}
security-to-group {
    security-model (usm | v1 | v2c) {
        security-name security-name {
            group group-name;
        }
    }
}
}
```

Hierarchy Level    [edit snmp]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

|                                 |                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>              | Configure SNMPv3.<br><br>The remaining statements are explained separately.                                          |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Minimum SNMPv3 Configuration on a Device Running Junos OS</i></li> </ul> |

## vacm

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> vacm {   access {     group group-name {       (default-context-prefix   context-prefix context-prefix){         security-model (any   usm   v1   v2c) {           security-level (authentication   none   privacy) {             notify-view view-name;             read-view view-name;             write-view view-name;           }         }       }     }   }   security-to-group {     security-model (usm   v1   v2c);     security-name security-name {       group group-name;     }   } } </pre> |
| <b>Hierarchy Level</b>          | [edit snmp v3]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Configure view-based access control model (VACM) information, including access privileges such as security model and security level for a group of users.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Defining Access Privileges for an SNMP Group</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                           |

## variable

---

|                                 |                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>variable <i>oid-variable</i>;</code>                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit snmp rmon <a href="#">alarm</a> <i>index</i> ]                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                        |
| <b>Description</b>              | Object identifier (OID) of MIB variable to be monitored.                                                                                                                                                                         |
| <b>Options</b>                  | <i>oid-variable</i> —OID of the MIB variable that is being monitored. The OID can be a dotted decimal (for example, 1.3.6.1.2.1.2.1.2.1.10.1). Alternatively, use the MIB object name (for example, <code>ifInOctets.1</code> ). |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Variable</i></li></ul>                                                                                                                                                |

## version (SNMP)

---

|                                 |                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version (all   v1   v2);</code>                                                                                            |
| <b>Hierarchy Level</b>          | [edit snmp trap-group <i>group-name</i> ]                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.        |
| <b>Description</b>              | Specify the version number of SNMP traps.                                                                                        |
| <b>Default</b>                  | all—Send an SNMPv1 and SNMPv2 trap for every trap condition.                                                                     |
| <b>Options</b>                  | all—Send an SNMPv1 and SNMPv2 trap for every trap condition.<br><br>v1—Send SNMPv1 traps only.<br><br>v2—Send SNMPv2 traps only. |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring SNMP Trap Groups</i></li></ul>                                            |



---


## view (Associating a MIB View with a Community)

---

|                                 |                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>view view-name;</code>                                                                                                                            |
| <b>Hierarchy Level</b>          | <code>[edit snmp community community-name]</code>                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                               |
| <b>Description</b>              | Associate a view with a community. A view represents a group of MIB objects.                                                                            |
| <b>Options</b>                  | <b>view-name</b> —Name of the view. You must use a view name already configured in the <b>view</b> statement at the <b>[edit snmp]</b> hierarchy level. |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring SNMP Communities</i></li></ul>                                                                   |

## view (Configuring a MIB View)

---

|                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                                                                                                                                                                                                                              | <pre>view view-name {<br/>    oid object-identifier (include   exclude);<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                      |
| Hierarchy Level                                                                                                                                                                                                                     | [edit snmp]                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Release Information                                                                                                                                                                                                                 | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                 |
| Description                                                                                                                                                                                                                         | Define a MIB view. A MIB view identifies a group of MIB objects. Each MIB object in a view has a common OID prefix. Each object identifier represents a subtree of the MIB object hierarchy. The <b>view</b> statement uses a view to specify a group of MIB objects on which to define access. To enable a view, you must associate the view with a community by including the <b>view</b> statement at the <b>[edit snmp community community-name]</b> hierarchy level. |
| <div> <b>NOTE:</b> To remove an OID completely, use the <b>delete view all oid oid-number</b> command but omit the <b>include</b> parameter.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Options                                                                                                                                                                                                                             | <b>view-name</b> —Name of the view.<br><br>The remaining statement is explained separately.                                                                                                                                                                                                                                                                                                                                                                               |
| Required Privilege Level                                                                                                                                                                                                            | snmp—To view this statement in the configuration.<br>snmp-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                             |
| Related Documentation                                                                                                                                                                                                               | <ul style="list-style-type: none"><li>• <i>Configuring MIB Views</i></li><li>• <i>Associating MIB Views with an SNMP User Group</i></li><li>• <a href="#">community on page 4092</a></li></ul>                                                                                                                                                                                                                                                                            |

## write-view

---

|                                 |                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>write-view view-name;</code>                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit snmp v3 vacm access group group-name (default-context-prefix   context-prefix context-prefix) security-model (any   usm   v1   v2c) security-level (authentication   none   privacy)]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series switches.  |
| <b>Description</b>              | Associate the write view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).                                                                                          |
| <b>Options</b>                  | <b>view-name</b> —Name of the view for which the SNMP user group has write permission.                                                                                                                   |
| <b>Required Privilege Level</b> | <b>snmp</b> —To view this statement in the configuration.<br><b>snmp-control</b> —To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring MIB Views</a></li> <li>• <a href="#">Configuring the Write View</a></li> </ul>                                                          |

## Configuration Statements: Analyzers and Port Mirroring

---

- [\[edit forwarding-options\] Configuration Statement Hierarchy on EX Series Switches on page 4155](#)
- [\[edit forwarding-options port-mirroring\] Configuration Statement Hierarchy on page 4157](#)
- [\[edit forwarding-options analyzer\] Configuration Statement Hierarchy on page 4158](#)
- [egress on page 4158](#)
- [egress \(Analyzer\) on page 4159](#)
- [ingress \(vlans\) on page 4159](#)
- [ingress \(Analyzer\) on page 4160](#)
- [input \(Analyzer\) on page 4161](#)
- [interface on page 4162](#)
- [no-tag on page 4163](#)
- [output \(Mirroring\) on page 4164](#)
- [vlan \(Mirroring\) on page 4165](#)

### [\[edit forwarding-options\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration subhierarchies in the **[edit forwarding-options]** hierarchy level on EX Series switches.

- *Supported* subhierarchies are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* subhierarchies are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).
- [Supported Subhierarchies in the \[edit forwarding-options\] Hierarchy Level on page 4156](#)
- [Unsupported Subhierarchies in the \[edit forwarding-options\] Hierarchy Level on page 4156](#)

### Supported Subhierarchies in the [edit forwarding-options] Hierarchy Level

The following list shows the [edit forwarding-options] subhierarchies supported on EX Series switches:

Each of the following topics lists the statements at a subhierarchy of the [edit forwarding-options] hierarchy.

- [\[edit forwarding-options analyzer\] Configuration Statement Hierarchy on page 4158](#)
- [\[edit forwarding-options dhcp-relay\] Configuration Statement Hierarchy for EX Series Switches on page 1550](#)
- [\[edit forwarding-options enhanced-hash-key\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit forwarding-options port-mirroring\] Configuration Statement Hierarchy on page 4157](#)
- [\[edit forwarding-options storm-control-profiles\] Configuration Statement Hierarchy for EX Series Switches on page 2211](#)

### Unsupported Subhierarchies in the [edit forwarding-options] Hierarchy Level

All subhierarchies in the [edit forwarding-options] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 432: Unsupported [edit forwarding-options] Subhierarchies on EX Series Switches**

| Subhierarchy | Hierarchy Level           |
|--------------|---------------------------|
| accounting   | [edit forwarding-options] |
| helpers      | [edit forwarding-options] |
| sampling     | [edit forwarding-options] |

#### Related Documentation

- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)

## [edit forwarding-options port-mirroring] Configuration Statement Hierarchy

```

forwarding-options {
  port-mirroring {
    family {
      ethernet-switching
      output {
        interface interface-name {
        }
        no-filter-check;
      }
      vlan vlan-name {
        no-tag;
      }
    }
    inet
    output {
      ip-address address {
      }
      routing-instance instance-name {
        ip-address address {
        }
      }
    }
  }
  instance instance-name {
    family (Port Mirroring)
    ethernet-switching {
      output {
        interface interface-name {
        }
        no-filter-check;
      }
      vlan vlan-name {
        no-tag;
      }
    }
    inet
    output {
      ip-address address {
      }
      routing-instance instance-name {
        ip-address address {
        }
      }
    }
  }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
  - [\[edit forwarding-options\] Configuration Statement Hierarchy on EX Series Switches on page 2698](#)

## [edit forwarding-options analyzer] Configuration Statement Hierarchy

```

forwarding-options {
  analyzer analyzer-name {
    input {
      egress {
        interface (all | interface-name);
      }
      ingress {
        interface (all | interface-name);
        routing-instance routing-instance-name {
          vlan (vlan-name | vlan-id | vlan-list);
        }
        vlan (vlan-name | vlan-id | vlan-list);
      }
    }
    output {
      interface interface-name;
      routing-instance routing-instance-name {
        vlan (vlan-name | vlan-id);
        no-tag;
      }
      vlan (vlan-name | vlan-id) {
        no-tag;
      }
    }
  }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
  - [\[edit forwarding-options\] Configuration Statement Hierarchy on EX Series Switches on page 2698](#)

## egress

|                                 |                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | egress;                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit <b>vlan</b> vlan-name <b>vlan-id</b> number <b>interface</b> interface-name]                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                        |
| <b>Description</b>              | Specify that the member interface of the VLAN allows only egress traffic.                                                                                                                    |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches</i></li> </ul> |

## egress (Analyzer)

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>egress {   bridge-domain (bridge-domain   vlan-id   vlan);   interface (all   <i>interface-name</i>);   vlan (vlan-name   vlan-id   vlan-list); }</pre>                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <a href="#">[edit [edit forwarding-options] Configuration Statement Hierarchy on page 2698]</a><br><a href="#">[edit forwarding-options analyzer] Configuration Statement Hierarchy on page 4158</a><br><i>analyzer-name</i> input]                                                                                                                     |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level <a href="#">[edit forwarding-options]</a> introduced in Junos OS Release 13.2X50-D10 (ELS).</p> <p>Hierarchy level <a href="#">[edit forwarding-options analyzer <i>analyzer-name</i> input]</a> introduced in Junos OS Release 14.1 for MX routers.</p> |
| <b>Description</b>              | <p>Specify ports where traffic exiting the interface is to be mirrored in a mirroring configuration.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches</i></li> </ul>                                                                                                                                                            |

## ingress (vlans)

|                                 |                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ingress;                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <a href="#">[edit <i>vlans</i> <i>vlan-name</i> <i>vlan-id</i> number <i>interface</i> <i>interface-name</i>]</a>                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                        |
| <b>Description</b>              | Specify that the member interface of the VLAN allows only ingress traffic.                                                                                                                   |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches</i></li> </ul> |

## ingress (Analyzer)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ingress {<br/>  bridge-domain (<i>bridge-domain-name</i>   <i>vlan-id</i>   <i>vlan-name</i>)<br/>  interface (all   <i>interface-name</i>);<br/>  routing-instance <i>routing-instance-name</i> {<br/>    vlan (<i>vlan-name</i>   <i>vlan-id</i>   <i>vlan-list</i>);<br/>  }<br/>  vlan (<i>vlan-name</i>   <i>vlan-id</i>   <i>vlan-name</i>);<br/>}</pre> |
| <b>Hierarchy Level</b>          | [ <a href="#">edit [edit forwarding-options] Configuration Statement Hierarchy on page 2698</a> ]<br>[ <a href="#">edit forwarding-options analyzer Configuration Statement Hierarchy on page 4158</a> ]<br><i>analyzer-name</i> input]                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Hierarchy level [ <a href="#">edit forwarding-options</a> ] introduced in Junos OS Release 13.2X50-D10 (ELS).<br>Hierarchy level [ <a href="#">edit forwarding-options analyzer <i>analyzer-name</i> input</a> ] introduced in Junos OS Release 14.1 for MX routers.                        |
| <b>Description</b>              | Configure ports, routing instances, VLANs, or bridge domains for which the entering traffic is mirrored as part of a mirroring configuration.<br><br>The remaining statements are explained separately.                                                                                                                                                             |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX4300 Switches on page 4016</a></li></ul>                                                                                                                                                            |



## input (Analyzer)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> input {   egress {     bridge-domain (<i>bridge-domain</i>   <i>vlan-id</i>   <i>vlan-list</i>)     interface (all   <i>interface-name</i>);     vlan (<i>vlan-name</i>   <i>vlan-id</i>   <i>vlan-name</i>);   }   ingress {     bridge-domain (<i>bridge-domain</i>   <i>vlan-id</i>   <i>vlan-list</i>)     interface (all   <i>interface-name</i>);     routing-instance <i>routing-instance-name</i> {       vlan (<i>vlan-name</i>   <i>vlan-id</i>   <i>vlan-list</i>);     }     vlan (<i>vlan-name</i>   <i>vlan-id</i>   <i>vlan-name</i>);   } } </pre> |
| <b>Hierarchy Level</b>          | <p>[edit “[edit forwarding-options] Configuration Statement Hierarchy” on page 2698]</p> <p>[“[edit forwarding-options analyzer] Configuration Statement Hierarchy” on page 4158<i>analyzer-name</i>]</p>                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level <b>[edit forwarding-options]</b> introduced in Junos OS Release 13.2X50-D10 (ELS).</p> <p>Hierarchy level <b>[edit forwarding-options analyzer <i>analyzer-name</i>]</b> introduced in Junos OS Release 14.1 for MX Series routers.</p>                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Define the traffic to be mirrored in a mirroring configuration—the definition can be a combination of:</p> <ul style="list-style-type: none"> <li>• Packets entering or exiting a port</li> <li>• Packets entering or exiting a VLAN</li> <li>• Packets entering or exiting a bridge domain</li> </ul> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                      |
| <b>Default</b>                  | No default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Mirroring for Local Monitoring of Employee Resource Use on EX4300 Switches on page 4000</a></li> <li>• <a href="#">Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches on page 4006</a></li> <li>• <a href="#">Understanding Port Mirroring and Analyzers on EX4300 Switches on page 3964</a></li> </ul>                                                                                                                                          |

## interface

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | interface (all   <i>interface-name</i> );                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [ <a href="#">[edit forwarding-options analyzer] Configuration Statement Hierarchy on page 4158</a><br><i>analyzer-name</i> input egress],<br>[ <a href="#">[edit forwarding-options analyzer] Configuration Statement Hierarchy on page 4158</a><br><i>analyzer-name</i> input ingress],<br>[ <a href="#">[edit forwarding-options analyzer] Configuration Statement Hierarchy on page 4158</a><br><i>analyzer-name</i> output] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Hierarchy level <a href="#">[edit forwarding-options]</a> introduced in Junos OS Release 13.2X50-D10 (ELS).<br>Statement introduced in Junos OS Release 14.1 for MX Series routers.                                                                                                                                                                      |
| <b>Description</b>              | Configure the interfaces for which traffic is mirrored.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p>all—Apply mirroring to all interfaces on the network device. Mirroring a high volume of traffic can be performance intensive for the device. Therefore, you should generally select specific input interfaces in preference to using the all keyword, or use the all keyword in combination with setting a ratio for statistical sampling.</p> <p><i>interface-name</i>—Apply mirroring to the specified interface only.</p>  |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Port Mirroring Analyzers for Local Monitoring of Employee Resource Use</i></li><li>• <i>Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use</i></li><li>• <i>Understanding Port Mirroring Analyzers</i></li></ul>                                                                                                                 |

---

## no-tag

---

|                                 |                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-tag;                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">[edit forwarding-options analyzer]</a> Configuration Statement Hierarchy on page 4158 <i>analyzer-name</i> output vlan ( <i>vlan-id</i>   <i>vlan-name</i> )]                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for EX Series switches.<br>Hierarchy level <a href="#">[edit forwarding-options]</a> introduced in Junos OS Release 13.2X50-D10 (ELS).                                                                                                                                   |
| <b>Description</b>              | Specify that remote mirroring packets are not tagged.                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Mirroring for Local Monitoring of Employee Resource Use on EX4300 Switches on page 4000</a></li><li>• <a href="#">Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches on page 4006</a></li></ul> |

## output (Mirroring)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>output {<br/>  bridge-domain (<i>bridge-domain</i>   <i>vlan-id</i>   <i>vlan-list</i>);<br/>  interface <i>interface-name</i>;<br/>  routing-instance <i>routing-instance-name</i> {<br/>    vlan (<i>vlan-name</i>   <i>vlan-id</i>);<br/>    no-tag<br/>  }<br/>  vlan (<i>vlan-name</i>   <i>vlan-id</i>   <i>vlan-list</i>) {<br/>    no-tag;<br/>  }<br/>}</pre>              |
| <b>Hierarchy Level</b>          | [ <a href="#">[edit forwarding-options analyzer] Configuration Statement Hierarchy on page 4158</a> <i>analyzer-name</i> ]                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level <a href="#">[edit forwarding-options]</a> introduced in Junos OS Release 13.2X50-D10 (ELS).</p> <p>Hierarchy level [ "<a href="#">[edit forwarding-options analyzer] Configuration Statement Hierarchy" on page 4158</a> ] introduced in Junos OS Release 14.1 for MX Series routers.</p> |
| <b>Description</b>              | <p>Configure the destination for mirrored traffic, either an interface on the network device, for local monitoring, or a VLAN or bridge domain, for remote monitoring. You can optionally configure the <b>no-tag</b> statement so that remote port mirroring packets are not tagged.</p> <p>The remaining statements are explained separately.</p>                                      |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Mirroring for Local Monitoring of Employee Resource Use on EX4300 Switches on page 4000</a></li><li>• <a href="#">Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches on page 4006</a></li></ul>                                                                   |

## vlan (Mirroring)

|                                 |                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vlan (vlan-id   vlan-name) {<br/>    no-tag;<br/>}</code>                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit “[edit forwarding-options] Configuration Statement Hierarchy” on page 2698]<br><br>[“[edit forwarding-options analyzer] Configuration Statement Hierarchy” on page 4158<br>analyzer-name output]                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Hierarchy level <b>[edit forwarding-options]</b> introduced in Junos OS Release 13.2X50-D10 (ELS).                                                                                                                                                                          |
| <b>Description</b>              | Configure mirrored traffic to be sent to a VLAN for remote monitoring. On a destination ( <b>output</b> ) VLAN, you can also configure the <b>no-tag</b> statement.                                                                                                                                                                                 |
| <b>Options</b>                  | <i>vlan-id</i> —Numeric VLAN identifier.<br><br><i>vlan-name</i> —Name of the VLAN.<br><br>The remaining statement is explained separately.                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <b>system</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX4300 Switches on page 4016</a></li> <li>• <a href="#">Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches on page 4006</a></li> </ul> |

## Configuration Statements for Network Analytics

- [address \(Analytics Collector\) on page 4166](#)
- [analytics on page 4167](#)
- [collector \(Analytics\) on page 4171](#)
- [depth-threshold on page 4172](#)
- [export-profiles on page 4173](#)
- [file \(Analytics\) on page 4175](#)
- [interface \(Export Profiles\) on page 4176](#)
- [interfaces \(Analytics Resource\) on page 4177](#)
- [latency-threshold on page 4178](#)
- [local \(Analytics Collector\) on page 4179](#)
- [resource \(Analytics\) on page 4180](#)

- [resource-profiles \(Analytics\)](#) on page 4181
- [system \(Analytics Resource\)](#) on page 4182
- [system \(Export Profiles\)](#) on page 4183
- [traceoptions \(Analytics\)](#) on page 4184

## address (Analytics Collector)

---

**Syntax**     `address ip-address {  
                  port number {  
                    transport protocol {  
                      export-profile profile-name;  
                    }  
                  }  
          }`

**Hierarchy Level**     [edit services analytics collector]

**Release Information**     Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description**     Configure the address of a remote server to receive streamed analytics (queue and traffic statistics) data.



**NOTE:** The `address` statement is not available in Junos OS Releases prior to 13.2X51-D15.

**Options**     `ip-address`—IP address of the remote server receiving the streamed data.

`port number`—Port number of the remote server receiving the streaming data.

`export-profile profile-name`—Name of the export profile containing the parameters for the analytics data being streamed.

`transport protocol`—A transport protocol used to stream data to the port.

**Values:**

- `tcp`—Transmission Control Protocol (TCP)
- `udp`—User Datagram Protocol (UDP)

**Required Privilege Level**     interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview](#) on page 3978
- [analytics](#) on page 4167
- [show analytics collector](#) on page 4346

## analytics

**Syntax** *Junos OS Release 13.2X51-D15 and later:*

```

analytics {
  collector {
    local {
      file filename {
        size size;
        files number;
      }
    }
    address ip-address {
      port number {
        transport protocol {
          export-profile profile-name;
        }
      }
    }
  }
  export-profiles {
    profile-name {
      interface {
        information;
        statistics {
          queue;
          traffic;
        }
        status {
          link;
          queue;
          traffic;
        }
      }
    }
    stream-format format;
    system {
      information;
      status {
        queue;
        traffic;
      }
    }
  }
  resource {
    interfaces {
      interface-name {
        resource-profile name;
      }
    }
    system {
      polling-interval {
        queue-monitoring interval;
        traffic-monitoring interval;
      }
    }
  }
}

```

```
        resource-profile name;  
    }  
}  
resource-profiles {  
    profile-name {  
        depth-threshold {  
            high number;  
            low number;  
        }  
        latency-threshold {  
            high number;  
            low number;  
        }  
        no-queue-monitoring;  
        no-traffic-monitoring;  
        queue-monitoring;  
        traffic-monitoring;  
    }  
}  
traceoptions {  
    file filename {  
        files number;  
        size size;  
    }  
}
```



*Junos OS Release 13.2X50-D15 and 13.2X51-D10 only:*

```

analytics {
  interfaces {
    all {
      depth-threshold high number low number;
      latency-threshold high number low number;
      queue-statistics;
      no-queue-statistics;
      traffic-statistics;
      no-traffic-statistics;
    }
    interface-name {
      depth-threshold high number low number;
      latency-threshold high number low number;
      queue-statistics;
      no-queue-statistics;
      traffic-statistics;
      no-traffic-statistics;
    }
  }
  queue-statistics {
    file filename {
      files number-of-files;
      size size;
    }
    interval interval;
  }
  streaming-servers {
    address ip-address {
      port number {
        stream-format format;
        stream-type type
      }
    }
  }
  traceoptions {
    file filename {
      files number;
      size size;
    }
  }
  traffic-statistics {
    file filename {
      files number-of-files;
      size size;
    }
    interval interval;
  }
}

```

**Hierarchy Level** [edit services]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Configure the network analytics feature that includes monitoring for traffic and queue statistics. The network analytics processes running on the Packet Forwarding Engine and Routing Engine collect and analyze the data, and generate reports that may be saved in log files or sent as streaming data to remote servers.

The remaining statements are explained separately.

**Required Privilege** interface—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview on page 3978](#)
- [show analytics traffic-statistics on page 4358](#)
- [show analytics collector on page 4346](#)
- [show analytics status on page 4354](#)
- [show analytics queue-statistics on page 4352](#)
- [show analytics configuration on page 4348](#)

## collector (Analytics)

```
Syntax collector {
    local {
        file filename {
            size size;
            files number;
        }
    }
    address ip-address {
        port number {
            transport protocol {
                export-profile profile-name;
            }
        }
    }
}
```

**Hierarchy Level** [edit services analytics]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Configure a local file for storing network analytics statistics and/or a remote server for receiving streamed statistics data.



**NOTE:** The `collector` statement is not available in Junos OS Releases prior to 13.2X51-D15.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview on page 3978](#)

## depth-threshold

---

|                            |                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>depth-threshold {<br/>    high <i>number</i>;<br/>    low <i>number</i>;<br/>}</pre>                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>     | [edit services analytics interfaces]<br>[edit services analytics resource-profiles]                                                                                                                                                                                                                                                                               |
| <b>Release Information</b> | Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement in the <b>[edit services analytics resource-profiles]</b> hierarchy level introduced in Junos OS Release 13.2X51-D15.<br>Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.                                                                              |
| <b>Description</b>         | If network analytics queue statistics monitoring is enabled, specify the high and low values (in bytes) of the queue depth (buffer) threshold. If you configure a depth threshold, you cannot configure the latency threshold. You can configure the depth threshold for one interface or all interfaces. Specify the high and low queue depth threshold numbers: |



**NOTE:** The configuration for a specific interface supersedes the global configuration for all interfaces.

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <p><b>high <i>number</i></b>—Specify the maximum value for the depth threshold.</p> <p><b>Range:</b> 1 to 1,250,000,000 bytes</p> <p><b>Default:</b></p> <ul style="list-style-type: none"><li>• Junos OS Release 13.2X51-D10 or later—0 bytes</li><li>• Junos OS Release 13.2X50-D15—14,680,064 bytes (14 MB)</li></ul> <p><b>low <i>number</i></b>—Specify the minimum value for the depth threshold.</p> <p><b>Range:</b> 1 to 1,250,000,000 bytes</p> <p><b>Default:</b></p> <ul style="list-style-type: none"><li>• Junos OS Release 13.2X51-D10 or later—0 bytes</li><li>• Junos OS Release 13.2X50-D15—1024 bytes (1 KB)</li></ul> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Network Analytics Overview on page 3978</a></li><li>• <a href="#">analytics on page 4167</a></li><li>• <a href="#">latency-threshold on page 4178</a></li><li>• <a href="#">resource-profiles (Analytics) on page 4181</a></li></ul>                                                                                                                                                                                                                                                                                                                                                  |

## export-profiles

```
Syntax  export-profiles {
        profile-name {
            interface {
                information;
                statistics {
                    queue;
                    traffic;
                }
            }
            status {
                link;
                queue;
                traffic;
            }
        }
    }
    stream-format format;
    system {
        information;
        status {
            queue;
            traffic;
        }
    }
}
```

**Hierarchy Level** [edit services analytics]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Configure an profile to specify the network analytics data being streamed to remote servers. Each profile is a template that defines the type of data being streamed.



**NOTE:** The `export-profile` statement is not available in Junos OS Releases prior to 13.2X51-D15.

**Options** *profile-name*—Name of the export profile containing the configuration of the data being streamed.

*stream-format format*—Format of the streaming data being sent to a server. Only one format can be sent to each port on a server.

**Values:**

- **csv**—Comma-separated Values (CSV). Data sent in this format is newline separated, and each record contains one stream type (queue or traffic data) per interface. Each record contains either a “q” for a queue statistics, or a “t” for a traffic statistics.

- **gpb**—Google Protocol Buffer (GPB). Data sent in this format has a hierarchical format, and is categorized by resource type (system or interfaces), which is specified in the message header. You can generate data formatted in other formats (CSV, TSV, and JSON) from GPB-encoded data.

Each message includes a 8-byte header containing the following information:

- Bytes 0 to 3—Length of the message.
- Byte 4—Message version.
- Bytes 5 to 7—Reserved for future use.



**NOTE:** A schema file called `analytics.proto` containing the definitions of the GPB messages is available for downloading from the following location:

[http://www.juniper.net/techpubs/en\\_US/junos132/topics/reference/proto-files/analytics-prototxt](http://www.juniper.net/techpubs/en_US/junos132/topics/reference/proto-files/analytics-prototxt)

---

- **json**—JavaScript Object Notation (JSON). Data sent in this format is newline separated, and each record contains one stream type (queue or traffic data) per interface. Each record contains either “queue-statistics” or “traffic-statistics” in the “record type” field.
- **tsv**—Tab-separated Values (TSV). Data sent in this format is newline separated, and each record contains one stream type (queue or traffic data) per interface. Each record contains a “q” for a queue statistics, or a “t” for a traffic statistics.

The remaining statements are explained separately.

|                                 |                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration. |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                                                                            |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Network Analytics Overview on page 3978</a></li><li>• <a href="#">analytics on page 4167</a></li></ul> |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|

## file (Analytics)

|                            |                                                                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | file <i>filename</i> {<br>files <i>number-of-files</i> ;<br>size <i>size</i> ;<br>}                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>     | [edit services analytics collector local]<br>[edit services analytics queue-statistics]<br>[edit services analytics traffic-statistics]                                                                                                                                                    |
| <b>Release Information</b> | Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.                                                                                                                                          |
| <b>Description</b>         | Enable the logging of queue or traffic monitoring statistics in a local file. This statement does not enable monitoring.                                                                                                                                                                   |
| <b>Default</b>             | This feature is disabled by default.                                                                                                                                                                                                                                                       |
| <b>Options</b>             | <b><i>filename</i></b> —Specify a filename for storing queue and traffic monitoring statistics in the Comma-separated Values (CSV) format. The file is stored in the <b>/var/log/</b> directory of your device.<br><br>If you do not specify a filename, the data is not stored in a file. |



**NOTE:** In Junos OS Release 13.2X51-D15 or later, you configure a single filename to store both queue and traffic monitoring statistics. In Junos OS Release 13.2X51-D10 and earlier, you configure separate files for storing monitoring data, one for queue statistics, and another for traffic statistics.

**files *number-of-files***—Specify the number of files to store locally. After the number of files with the maximum file size is reached, the system starts over and writes the data to the first file.

**Range:** 2 to 1,000 files.

**size *size***—Configure the file size in megabytes (MB).

**Syntax:** *xm* to specify MB.

**Range:** 10 to 4095 MB

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview on page 3978](#)
- [analytics on page 4167](#)

## interface (Export Profiles)

---

**Syntax**

```
interface {  
    information;  
    statistics {  
        queue;  
        traffic;  
    }  
    status {  
        link;  
        queue;  
        traffic;  
    }  
}
```

**Hierarchy Level** [edit services analytics export-profiles]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Configure an export profile for streaming network analytics data for a specific interface to remote servers. Each profile is a template that defines the type of data being streamed for that interface.



**NOTE:** The `interface` statement is not available in Junos OS Releases prior to 13.2X51-D15.

---

**Options** **information**—Information about the specified interface, including SNMP index, interface index, slot, port number, media type, capability, and port type.

**statistics**—Type of monitoring statistics to be streamed.

**Values:**

- `queue`
- `traffic`

**status**—Status information about the interface to be streamed.

**Values:**

- `link`
- `queue`
- `traffic`

**Required Privilege Level** `interface`—To view this statement in the configuration.  
`interface-control`—To add this statement to the configuration.



- Related Documentation**
- [Network Analytics Overview on page 3978](#)
  - [analytics on page 4167](#)

## interfaces (Analytics Resource)

**Syntax**

```
interfaces {
  interface-name {
    resource-profile profile-name;
  }
}
```

**Hierarchy Level** [edit services analytics resource]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Apply the network analytics resource profile to an interface for which you wish to enable queue or traffic statistics monitoring. The resource profile is a template that specifies the parameters for queue and traffic monitoring, as well as for the depth and latency thresholds.



**NOTE:** The `interfaces` statement in the [edit services analytics resource] hierarchy is not available in Junos OS Releases prior to 13.2X51-D15.

**Options** *interface-name*—Name of the interface for which a resource profile has been configured.

*resource-profile profile-name*—Name of a resource profile containing the analytics parameters that have been specified for interfaces. Information contained in a resource profile includes the configuration of queue and traffic monitoring (whether enabled or disabled), and values for the depth and latency thresholds (if applicable).

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

- Related Documentation**
- [Network Analytics Overview on page 3978](#)
  - [analytics on page 4167](#)

## latency-threshold

---

|                            |                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | latency-threshold {<br>high <i>number</i> ;<br>low <i>number</i> ;<br>}                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>     | [edit services analytics interfaces]<br>[edit services analytics resource-profiles]                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b> | Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement in the <b>[edit services analytics resource-profiles]</b> hierarchy level introduced in Junos OS Release 13.2X51-D15.<br>Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.                                                                                   |
| <b>Description</b>         | If network analytics queue statistics monitoring is enabled, specify the high and low values (in microseconds) of the latency threshold of the queue. If you configure a latency threshold, you cannot configure the depth threshold. You can configure the latency threshold for one interface or all interfaces. Specify the high and low latency threshold numbers: |



**NOTE:** The configuration for a specific interface supersedes the global configuration for all interfaces.

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b> | <b>high <i>number</i></b> —Specify the maximum value for the latency threshold.<br><b>Range:</b> <ul style="list-style-type: none"><li>Junos OS Release 13.2X51-D15 or later—1 to 100,000,000 nanoseconds (0.001 to 100,000 microseconds)</li><li>Junos OS Release 13.2X51-D10 or earlier—1 to 100,000 microseconds</li></ul> <b>Default:</b> <ul style="list-style-type: none"><li>Junos OS Release 13.2X51-D15 or later—1,000,000 nanoseconds (1000 microseconds or 1 millisecond)</li><li>Junos OS Release 13.2X51-D10—1000 microseconds</li><li>Junos OS Release 13.2X50-D15—900 microseconds</li></ul> <b>low <i>number</i></b> —Specify the minimum value for the latency threshold.<br><b>Range:</b> <ul style="list-style-type: none"><li>Junos OS Release 13.2X51-D15 or later—1 to 100,000,000 nanoseconds</li><li>Junos OS Release 13.2X51-D10 or earlier—1 to 100,000 microseconds</li></ul> <b>Default:</b> <ul style="list-style-type: none"><li>Junos OS Release 13.2X51-D15 or later—100 nanoseconds (0.1 microseconds)</li></ul> |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

- Junos OS Release 13.2X51-D10—50 microseconds
- Junos OS Release 13.2X50-D15—300 microseconds

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview on page 3978](#)
- [analytics on page 4167](#)
- [depth-threshold on page 4172](#)

## local (Analytics Collector)

**Syntax**

```
local {
  file filename {
    size size;
    files number;
  }
}
```

**Hierarchy Level** [edit services analytics collector]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Configure a local file for logging network analytics (queue and traffic) statistics.



**NOTE:** The `local` statement is not available in Junos OS Releases prior to 13.2X51-D15.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview on page 3978](#)
- [collector \(Analytics\) on page 4171](#)

## resource (Analytics)

---

**Syntax**

```
resource {  
  interfaces {  
    interface-name {  
      resource-profile profile-name;  
    }  
  }  
  system {  
    polling-interval {  
      queue-monitoring interval;  
      traffic-monitoring interval;  
    }  
    resource-profile profile-name;  
  }  
}
```

**Hierarchy Level** [edit services analytics]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Configure network analytics resources such as resource profiles (for interfaces and system), and polling intervals (for queue and traffic monitoring).



**NOTE:** The **resource** statement is not available in Junos OS Releases prior to 13.2X51-D15.

---

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview on page 3978](#)
- [analytics on page 4167](#)

## resource-profiles (Analytics)

**Syntax**

```
resource-profiles {
  profile-name {
    depth-threshold {
      high number;
      low number;
    }
    latency-threshold {
      high number;
      low number;
    }
    no-queue-monitoring;
    no-traffic-monitoring;
    queue-monitoring;
    traffic-monitoring;
  }
}
```

**Hierarchy Level** [edit services analytics]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Configure resource profiles that are used as templates for specifying network analytics parameters. You use resource profiles to enable and disable queue and traffic monitoring, and specify depth and latency thresholds as applicable. Once you have defined a resource profile, you can apply it specifically to individual interfaces, or globally to a system.



**NOTE:** The `resource-profiles` statement is not available in Junos OS Releases prior to 13.2X51-D15.

The remaining statements are explained separately.

**Options** *profile-name*—Specify a name for the resource profile.

*no-queue-monitoring*—Disable queue monitoring.

*no-traffic-monitoring*—Disable traffic monitoring.

*queue-monitoring*—Enable queue monitoring.

*traffic-monitoring*—Enable traffic monitoring.


**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**


- [Network Analytics Overview on page 3978](#)
- [analytics on page 4167](#)

## system (Analytics Resource)

---

|                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                                                                                                                                                                                                                                                      | <pre>system {<br/>  polling-interval {<br/>    queue-monitoring <i>interval</i>;<br/>    traffic-monitoring <i>interval</i>;<br/>  }<br/>  resource-profile <i>profile-name</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Hierarchy Level                                                                                                                                                                                                                                             | [edit services analytics resource]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Release Information                                                                                                                                                                                                                                         | Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Description                                                                                                                                                                                                                                                 | Apply a network analytics resource profile to a system for which you wish to enable queue or traffic monitoring. The resource profile is a template that specifies the parameters for queue and traffic monitoring, as well as for the depth and latency thresholds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <div> <b>NOTE:</b> The <code>system</code> statement in the [edit services analytics resource] hierarchy is not available in Junos OS Releases prior to 13.2X51-D15.</div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Options                                                                                                                                                                                                                                                     | <p><b>polling-interval</b>—Configure the polling interval for queue and traffic monitoring:</p> <p><b>queue-monitoring <i>polling-interval</i></b>—Configure the queue monitoring interval in milliseconds.</p> <p><b>Range:</b> 1 to 1000 milliseconds (1 millisecond to 1 second) on devices other than EX4300 switches. 8 to 1000 milliseconds (8 milliseconds to 1 second) on EX4300 switches.</p> <p><b>traffic-monitoring <i>polling-interval</i></b>—Configure the traffic monitoring interval in seconds.</p> <p><b>Range:</b> 1 to 300 seconds (1 second to 5 minutes) on devices other than EX4300 switches. 5 to 300 seconds (5 seconds to 5 minutes) on EX 4300 switches.</p> <p><b>resource-profile <i>profile-name</i></b>—Name of a resource profile containing the global analytics parameters that have been configured for the system. Information contained in a resource profile includes the configuration of queue and traffic monitoring (whether enabled or disabled), and values for the depth and latency thresholds (if applicable).</p> |
| Required Privilege Level                                                                                                                                                                                                                                    | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Related Documentation                                                                                                                                                                                                                                       | <ul style="list-style-type: none"><li>• <a href="#">Network Analytics Overview on page 3978</a></li><li>• <a href="#">analytics on page 4167</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## system (Export Profiles)

|                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                              | <pre> system {   information;   status {     queue;     traffic;   } } </pre>                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>                                                                                                                                                                                     | [edit services analytics export-profiles]                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>                                                                                                                                                                                 | <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.</p>                                                                                                                                                                                           |
| <b>Description</b>                                                                                                                                                                                         | Configure a system-wide export profile for streaming network analytics data to remote servers. Each profile is a template that defines the type of data being streamed for that system.                                                                                                                                                                |
| <div>  <p><b>NOTE:</b> The <b>system</b> statement is not available in Junos OS Releases prior to 13.2X51-D15.</p> </div> |                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                                                                                                                                                                             | <p><b>information</b>—Information about the system, including boot time, model, serial number, maximum number of ports, collector information, and interface list.</p> <p><b>status</b>—System status information to be streamed.</p> <p><b>Values:</b></p> <ul style="list-style-type: none"> <li>• <b>queue</b></li> <li>• <b>traffic</b></li> </ul> |
| <b>Required Privilege Level</b>                                                                                                                                                                            | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                                               | <ul style="list-style-type: none"> <li>• <a href="#">Network Analytics Overview on page 3978</a></li> <li>• <a href="#">analytics on page 4167</a></li> </ul>                                                                                                                                                                                          |

## traceoptions (Analytics)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>traceoptions {<br/>    file <i>filename</i>;<br/>    files <i>number-of-files</i>;<br/>    size <i>size</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit services analytics]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure traceoptions for the network analytics daemon (analyticsd) running on the Routing Engine.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>file <i>filename</i></b>—Specify a filename for storing the traceoptions data. The file is stored in the <code>/var/log/</code> directory of your device.</p> <p>If you do not specify a filename, the data is not stored in a file.</p> <p><b>files <i>number-of-files</i></b>—Specify the number of files to store locally. After the number files with the maximum file size is reached, the system starts over and writes the data to the first file.</p> <p><b>Range:</b> 2 to 1,000 files.</p> <p><b>size <i>size</i></b>—Configure the file size in megabytes (MB).</p> <p><b>Syntax:</b> <i>xm</i> to specify MB.</p> <p><b>Range:</b> 10 to 4095 MB</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Network Analytics Overview on page 3978</a></li><li>• <a href="#">analytics on page 4167</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## Configuration Statements: sFlow Technology

---

- [\[edit protocols\] Configuration Statement Hierarchy on EX4300 Switches on page 4185](#)
- [\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches on page 4186](#)
- [agent-id on page 4188](#)
- [collector on page 4188](#)
- [disable \(sFlow Monitoring Technology\) on page 4189](#)
- [interfaces \(sFlow Monitoring Technology\) on page 4190](#)
- [polling-interval on page 4191](#)



- [sample-rate on page 4192](#)
- [sflow on page 4193](#)
- [source-ip on page 4194](#)
- [udp-port on page 4194](#)

## **[edit protocols]** Configuration Statement Hierarchy on EX4300 Switches

Each of the following topics lists the statements at a subhierarchy of the **[edit protocols]** hierarchy:

- [\[edit protocols bfd\] Configuration Statement Hierarchy on EX Series Switches on page 391](#)
- [\[edit protocols bgp\] Configuration Statement Hierarchy on EX Series Switches on page 392](#)
- [\[edit protocols dot1x\] Configuration Statement Hierarchy on EX Series Switches on page 404](#)
- [\[edit protocols igmp\] Configuration Statement Hierarchy on EX Series Switches on page 406](#)
- [\[edit protocols igmp-snooping\] Configuration Statement Hierarchy on page 3743](#)
- [\[edit protocols isis\] Configuration Statement Hierarchy on EX Series Switches on page 408](#)
- [\[edit protocols lacp\] Configuration Statement Hierarchy on EX Series Switches on page 411](#)
- [\[edit protocols l2-learning\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols layer2-control\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols lldp\] Configuration Statement Hierarchy on EX Series Switches on page 413](#)
- [\[edit protocols lldp-med\] Configuration Statement Hierarchy on EX Series Switches on page 415](#)
- [\[edit protocols msdp\] Configuration Statement Hierarchy on EX Series Switches on page 429](#)
- [\[edit protocols mstp\] Configuration Statement Hierarchy on EX Series Switches on page 431](#)
- [\[edit protocols mvrp\] Configuration Statement Hierarchy on EX Series Switches on page 433](#)
- [\[edit protocols neighbor-discovery\] Configuration Statement Hierarchy on EX Series Switches on page 434](#)
- [\[edit protocols oam\] Configuration Statement Hierarchy on EX Series Switches on page 435](#)
- [\[edit protocols ospf\] Configuration Statement Hierarchy on EX Series Switches on page 438](#)

- [\[edit protocols ospf3\] Configuration Statement Hierarchy on EX Series Switches on page 441](#)
- [\[edit protocols pim\] Configuration Statement Hierarchy on EX Series Switches on page 444](#)
- [\[edit protocols rip\] Configuration Statement Hierarchy on EX Series Switches on page 447](#)
- [\[edit protocols ripng\] Configuration Statement Hierarchy on EX Series Switches on page 450](#)
- [\[edit protocols router-advertisement\] Configuration Statement Hierarchy on EX Series Switches on page 451](#)
- [\[edit protocols router-discovery\] Configuration Statement Hierarchy on EX Series Switches on page 452](#)
- [\[edit protocols rstp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches on page 459](#)
- [\[edit protocols uplink-failure-detection\] Configuration Statement Hierarchy on EX Series Switches on page 461](#)
- [\[edit protocols vrrp\] Configuration Statement Hierarchy on EX Series Switches on page 462](#)
- [\[edit protocols vstp\] Configuration Statement Hierarchy on EX Series Switches](#)

**Related  
Documentation**

- [EX Series Switch Software Features Overview](#)

## [\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the [\[edit protocols sflow\]](#) hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols sflow\] Hierarchy Level on page 4187](#)
- [Unsupported Statements in the \[edit sflow\] Hierarchy Level on page 4187](#)

### Supported Statements in the [edit protocols sflow] Hierarchy Level

The following hierarchy shows the [edit protocols sflow] configuration statements supported on EX Series switches:

```
sflow {
  agent-id;
  collector {
    ip-address;
    udp-port port-number;
  }
  interfaces interface-name {
    polling-interval seconds;
    sample-rate {
      egress number;
      ingress number;
    }
  }
  polling-interval seconds;
  sample-rate {
    egress number;
    ingress number;
  }
  source-ip;
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
  no-world-readable>;
  flag (all | client-server | configuration | interface | rtsock);
}
```

### Unsupported Statements in the [edit sflow] Hierarchy Level

All statements in the [edit protocols sflow] hierarchy level that are displayed in the command-line interface (CLI) on the EX Series switch are supported on the switch and operate as documented.

- Related Documentation**
- [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\) on page 4049](#)
  - [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 389](#)

## agent-id

---

|                                 |                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>agent-id ip-address;</code>                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">sflow</a> ]                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                   |
| <b>Description</b>              | Configure the IP address to be assigned as the agent ID for the sFlow agent. By assigning an IP address, you ensure that the IP address is not dynamic. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li></ul>      |

## collector

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>collector {<br/>    ip-address;<br/>    <a href="#">udp-port</a> port-number;<br/>}</pre>                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">sflow</a> ]                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Configure a remote collector for sFlow network traffic monitoring. The switch sends sFlow UDP datagrams to this collector for analysis. You can configure up to four collectors on the switch. You configure a collector by specifying its IP address and a UDP port.</p> <p>The remaining statements are explained separately.</p>                                                          |
| <b>Default</b>                  | The default port is 6343.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">[edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 389</a></li><li>• <a href="#">Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995</a></li><li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li></ul> |

---

## disable (sFlow Monitoring Technology)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">sflow</a> ],<br>[edit protocols <a href="#">sflow interfaces</a> <i>interface-name</i> ]                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Disable the sFlow monitoring protocol on all interfaces on the switch or on the specified interface.                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">[edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 389</a></li><li>• <a href="#">Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995</a></li><li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li></ul> |

## interfaces (sFlow Monitoring Technology)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interfaces <i>interface-name</i> {<br/>    <b>polling-interval</b> <i>seconds</i>;<br/>    <b>sample-rate</b> {<br/>        egress <i>number</i>;<br/>        ingress <i>number</i>;<br/>    }<br/>}</pre>                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols <b>sflow</b> ]                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Configure sFlow network traffic monitoring on the specified interface on the switch. You can configure sFlow parameters such as polling interval and sampling rate with different values on different interfaces, and you can also disable sFlow monitoring on individual interfaces.</p> <p>The remaining statements are explained separately.</p>                                          |
| <b>Options</b>                  | <b><i>interface-name</i></b> —Name of the interface on which to configure sFlow parameters.                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">[edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 389</a></li><li>• <a href="#">Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995</a></li><li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li></ul> |

## polling-interval

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>polling-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">sflow</a> ],<br>[edit protocols <a href="#">sflow interfaces</a> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Configure the interval (in seconds) that the switch waits between port statistics update messages. <i>Polling</i> refers to the switch gathering various statistics for the network interfaces configured for sFlow monitoring and exporting the statistics to the configured sFlow collector.                                                                                                      |
| <b>Default</b>                  | If no polling interval is configured for a particular interface, the switch waits the number of seconds that is configured as the polling interval in the global sFlow configuration. If no polling interval is specified in the global configuration, the switch waits 20 seconds between messages.                                                                                                |
| <b>Options</b>                  | <i>seconds</i> —Number of seconds between port statistics update messages. A value of <b>0</b> (zero) specifies that polling is disabled.<br><b>Range:</b> 0 through 3600 seconds<br><b>Default:</b> 20 seconds                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">[edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 389</a></li> <li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li> <li>• <a href="#">Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995</a></li> </ul> |

## sample-rate

---

|                     |                                                                                                                                                                                                                                                                                                 |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax              | <pre>sample-rate {<br/>    egress <i>number</i>;<br/>    ingress <i>number</i>;<br/>}</pre>                                                                                                                                                                                                     |
| Hierarchy Level     | [edit protocols <a href="#">sflow</a> ],<br>[edit protocols <a href="#">sflow interfaces</a> <i>interface-name</i> ]                                                                                                                                                                            |
| Release Information | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>The option <i>number</i> (which immediately followed <b>sample-rate</b> ) is deprecated and options <b>egress <i>number</i></b> and <b>ingress <i>number</i></b> added in Junos OS Release 10.4 for EX Series switches. |
| Description         | Specify the number of egress or ingress packets out of which one packet is sampled. If no interface sampling rates are configured, the global sampling rates take effect. If neither is configured, by default both ingress and egress packet sampling is disabled                              |
| Default             | By default, both egress and ingress sampling rates are disabled.                                                                                                                                                                                                                                |



**NOTE:** The **sample-rate *number*** (the global sampling rate) statement has been deprecated and might be removed from future product releases. We strongly recommend that you phase out its use.

|                          |                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Options                  | <b>egress <i>number</i></b> —Value for egress sampling rate.<br><b>Range:</b> 100 through 1,073,741,823<br><br><b>ingress <i>number</i></b> —Value for ingress sampling rate.<br><b>Range:</b> 100 through 1,073,741,823                                                                                                                                                                        |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                             |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">[edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 389</a></li><li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li><li>• <a href="#">Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995</a></li></ul> |



## sflow

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> sflow {   agent-id ip-address;   collector {     ip-address;     udp-port port-number;   }   interfaces interface-name {     disable;     polling-interval seconds;     sample-rate {       egress number;       ingress number;     }   }   polling-interval seconds;   sample-rate {     egress number;     ingress number;   }   source-ip ip-address; } </pre>                            |
| <b>Hierarchy Level</b>          | [edit protocols]                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Options <b>agent-id</b> and <b>source-ip</b> added in Junos OS Release 10.2 for EX Series switches.</p>                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure sFlow technology to continuously monitor traffic at wire speed on specified interfaces simultaneously. sFlow data can be used to provide network traffic visibility information.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                         |
| <b>Default</b>                  | sFlow technology is disabled by default.                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">[edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 389</a></li> <li>• <a href="#">Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995</a></li> <li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li> </ul> |

## source-ip

---

|                                 |                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-ip <i>ip-address</i>;</code>                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">sflow</a> ]                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                              |
| <b>Description</b>              | Configure the IP address to be used for the sFlow datagrams. By configuring an IP address, you ensure that the IP address is not dynamic.          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li></ul> |

## udp-port

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>udp-port <i>port-number</i>;</code>                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">sflow collector</a> ]                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure the UDP port for a remote collector for sFlow network traffic monitoring. The switch sends sFlow UDP datagrams to the collector for analysis.                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <i>port-number</i> —UDP port number for this collector.<br><b>Default:</b> 6343                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">[edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 389</a></li><li>• <a href="#">Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995</a></li><li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li></ul> |

## Configuration Statements: Ethernet OAM Connectivity Fault Management

---

- [action-profile \(Applying to OAM CFM, for EX Series Switch Only\) on page 4196](#)
- [age \(EX Series Switch Only\) on page 4197](#)
- [auto-discovery \(EX Series Switch Only\) on page 4197](#)
- [calculation-weight on page 4198](#)
- [connectivity-fault-management \(EX Series Switch Only\) on page 4199](#)

- [continuity-check \(EX Series Switch Only\) on page 4200](#)
- [cycle-time on page 4201](#)
- [delay on page 4202](#)
- [delay-variation on page 4203](#)
- [direction \(EX Series Switch Only\) on page 4204](#)
- [hold-interval \(OAM CFM, for EX Series Switch Only\) on page 4204](#)
- [interface \(OAM CFM, for EX Series Switch Only\) on page 4205](#)
- [interval \(EX Series Switch Only\) on page 4206](#)
- [iteration-period on page 4207](#)
- [level \(EX Series Switch Only\) on page 4207](#)
- [linktrace \(EX Series Switch Only\) on page 4208](#)
- [loss-threshold \(EX Series Switch Only\) on page 4208](#)
- [maintenance-association \(EX Series Switch Only\) on page 4209](#)
- [maintenance-domain \(EX Series Switch Only\) on page 4210](#)
- [measurement-type \(OAM LFM\) on page 4211](#)
- [mep \(EX Series Switch Only\) on page 4212](#)
- [mip-half-function \(EX Series Switch Only\) on page 4213](#)
- [name-format \(EX Series Switch Only\) on page 4214](#)
- [path-database-size \(EX Series Switch Only\) on page 4214](#)
- [performance-monitoring \(OAM LFM\) on page 4215](#)
- [remote-mep \(EX Series Switch Only\) on page 4216](#)
- [sla-iterator-profiles \(OAM LFM\) on page 4217](#)

## action-profile (Applying to OAM CFM, for EX Series Switch Only)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>action-profile <i>profile-name</i> {<br/>    action {<br/>        interface-down;<br/>    }<br/>    default-actions {<br/>        interface-down;<br/>    }<br/>    event {<br/>        adjacency-loss;<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet connectivity-fault-management</a> ]<br>[edit protocols <a href="#">oam ethernet connectivity-fault-management maintenance-domain</a> <i>domain-name</i> <a href="#">maintenance-association</a> <i>ma-name</i> <a href="#">mep</a> <i>mep-id</i> <a href="#">remote-mep</a> <i>mep-id</i> ]                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure a name and default action for an action profile.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><i>profile-name</i>—Name of the action profile.</p> <p><i>action</i>—Defines the action to be taken when connectivity to the remote MEP is lost.</p> <p><i>interface-down</i>—Brings the interface down when a remote MEP connectivity failure is detected.</p> <p><i>default-actions</i>—Defines the default action to be taken when connectivity to the remote MEP is lost.</p> <p><i>interface-down</i>—Brings the interface down when a remote MEP connectivity failure is detected.</p> <p><i>event</i>—Defines the event to be monitored when a remote MEP connectivity failure is detected.</p> <p><i>adjacency-loss</i>—Defines the connectivity loss to the remote MEP.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)</a> on page 4058</li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## age (EX Series Switch Only)

---

|                                 |                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | age (30m   10m   1m   30s   10s);                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet connectivity-fault-management linktrace</a> ]                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                              |
| <b>Description</b>              | Configure the time to wait (in minutes or seconds) for a response. If no response is received, the request and response entry is deleted from the linktrace database.                                                              |
| <b>Default</b>                  | 10 minutes                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li> <li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li> </ul> |

## auto-discovery (EX Series Switch Only)

---

|                                 |                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | auto-discovery;                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id</a> ]                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                              |
| <b>Description</b>              | Enable the MEP to accept continuity check messages from all remote MEPs.                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li> <li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li> </ul> |

## calculation-weight

---

|                            |                                                                                                                                                           |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>calculation-weight {<br/>    <code>delay</code> <i>delay-value</i>;<br/>    <code>delay-variation</code> <i>delay-variation-value</i>;<br/>}</code> |
| <b>Hierarchy Level</b>     | [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles <i>profile-name</i> ]                             |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1.<br>Statement introduced in Junos OS Release 11.4 for EX Series switches.                                   |
| <b>Description</b>         | Configure the calculation weight for delay and delay variation.                                                                                           |



**NOTE:** This option is applicable only for two-way delay measurement.

---

The remaining statements are explained separately.

|                           |                                        |
|---------------------------|----------------------------------------|
| <b>Required Privilege</b> | Configure—To enter configuration mode. |
| <b>Level</b>              | Control—To modify any configuration.   |

|                              |                                                                                                                                                                                                                                                                                                          |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Configuring an Iterator Profile</a></li><li>• <a href="#">Configuring an Iterator Profile on a Switch (CLI Procedure) on page 4067</a></li><li>• <a href="#">delay on page 4202</a></li><li>• <a href="#">delay-variation on page 4203</a></li></ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## connectivity-fault-management (EX Series Switch Only)

```

Syntax  connectivity-fault-management {
        action-profile profile-name {
            action {
                interface-down;
            }
            default-actions {
                interface-down;
            }
            event {
                adjacency-loss;
            }
        }
        linktrace {
            age (30m | 10m | 1m | 30s | 10s);
            path-database-size path-database-size;
        }
        maintenance-domain domain-name {
            level number;
            mip-half-function (none | default | explicit);
            name-format (character-string | none | dns | mac+2oct);
            maintenance-association ma-name {
                continuity-check {
                    hold-interval minutes;
                    interface-status-tlv;
                    interval (10m | 10s | 1m | 1s | 100ms);
                    loss-threshold number;
                    port-status-tlv;
                }
                mep mep-id {
                    auto-discovery;
                    direction down;
                    interface interface-name;
                    remote-mep mep-id {
                        action-profile profile-name;
                    }
                }
            }
        }
        performance-monitoring {
            sla-iterator-profiles {
                profile-name {
                    calculation-weight {
                        delay delay-value;
                        delay-variation delay-variation-value;
                    }
                    cycle-time cycle-time-value;
                    iteration-period iteration-period-value;
                    measurement-type two-way-delay;
                    passive;
                }
            }
        }
    }

```

```
}
```

|                                 |                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet</a> ]                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.<br><b>performance-monitoring</b> introduced in Junos OS Release 11.4 for EX Series switches.                                                              |
| <b>Description</b>              | Configure connectivity fault management for IEEE 802.1ag Operation, Administration, and Management (OAM) support.<br><br>The remaining statements are explained separately.                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul> |

---

## continuity-check (EX Series Switch Only)

---

|                                 |                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>continuity-check {<br/>  <a href="#">hold-interval</a> <i>minutes</i>;<br/>  interface-status-tlv;<br/>  <a href="#">interval</a> (10m   10s   1m   1s   100ms);<br/>  <a href="#">loss-threshold</a> <i>number</i>;<br/>  port-status-tlv;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet connectivity-fault-management maintenance-domain</a> <i>domain-name</i> <a href="#">maintenance-association</a> <i>ma-name</i> ]                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                        |
| <b>Description</b>              | Specify continuity check protocol options.<br><br>The remaining statements are explained separately.                                                                                                                                                         |
| <b>Options</b>                  | <b>interface-status-tlv</b> —Includes interface status TLV in CCM.<br><br><b>port-status-tlv</b> —Includes port status TLV in CCM.<br><br>The remaining statements are explained separately.                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul>                              |



---

## cycle-time

---

|                                 |                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>cycle-time <i>cycle-time-value</i>;</code>                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles <i>profile-name</i> ]                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1.<br>Statement introduced in Junos OS Release 11.4 for EX Series switches.                                                                              |
| <b>Description</b>              | Configure the time (in milliseconds) taken between back-to-back transmissions of SLA frames for a single connection.                                                                                 |
| <b>Options</b>                  | <b><i>cycle-time-value</i></b> —Cycle time value in milliseconds.<br><b>Range:</b> 10 through 3,600,000<br><b>Default:</b> 1000                                                                      |
| <b>Required Privilege Level</b> | Configure—To enter configuration mode.<br>Control—To modify any configuration.                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring an Iterator Profile</a></li><li>• <a href="#">Configuring an Iterator Profile on a Switch (CLI Procedure) on page 4067</a></li></ul> |

## delay

---

|                            |                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>delay <i>delay-value</i>;</code>                                                                                                                  |
| <b>Hierarchy Level</b>     | [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles <i>profile-name</i> <b>calculation-weight</b> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1.<br>Statement introduced in Junos OS Release 11.4 for EX Series switches.                                 |
| <b>Description</b>         | Configure the calculation weight for delay.                                                                                                             |
| <b>Options</b>             | <b><i>delay-value</i></b> —Calculation weight for delay.                                                                                                |



**NOTE:** This option is applicable only for two-way delay measurement.

---

**Range:** 1 through 65,535

**Default:** 1

|                                 |                                                                                |
|---------------------------------|--------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | Configure—To enter configuration mode.<br>Control—To modify any configuration. |
|---------------------------------|--------------------------------------------------------------------------------|

|                              |                                                                                                                                                                                                                                                                |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Configuring an Iterator Profile</a></li><li>• <a href="#">Configuring an Iterator Profile on a Switch (CLI Procedure) on page 4067</a></li><li>• <a href="#">calculation-weight on page 4198</a></li></ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## delay-variation

|                            |                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>delay-variation <i>delay-variation-value</i>;</code>                                                                                              |
| <b>Hierarchy Level</b>     | [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles <i>profile-name</i> <b>calculation-weight</b> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1.<br>Statement introduced in Junos OS Release 11.4 for EX Series switches.                                 |
| <b>Description</b>         | Configure the calculation weight for delay variation.                                                                                                   |
| <b>Options</b>             | <b><i>delay-variation-value</i></b> —Calculation weight for delay variation.                                                                            |



**NOTE:** This option is applicable only for two-way delay measurement.

**Range:** 1 through 65,535

**Default:** 1

|                                 |                                                                                |
|---------------------------------|--------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | Configure—To enter configuration mode.<br>Control—To modify any configuration. |
|---------------------------------|--------------------------------------------------------------------------------|

|                              |                                                                                                                                                                                                                                                                    |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an Iterator Profile</a></li> <li>• <a href="#">Configuring an Iterator Profile on a Switch (CLI Procedure) on page 4067</a></li> <li>• <a href="#">calculation-weight on page 4198</a></li> </ul> |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## direction (EX Series Switch Only)

---

|                                 |                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | direction down;                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id</a> ]                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                    |
| <b>Description</b>              | Specify that connectivity fault management (CFM) packets (CCMs) be transmitted only in one direction for the MEP, that is, the direction be set as <b>down</b> so that CCMs are transmitted only out of (not into) the interface configured on this MEP. |
| <b>Options</b>                  | <b>down</b> —Down MEP CCMs are transmitted only out (not into) of the interface configured on this MEP.                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul>                          |

## hold-interval (OAM CFM, for EX Series Switch Only)

---

|                                 |                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | hold-interval <i>minutes</i> ;                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name continuity-check</a> ]                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                           |
| <b>Description</b>              | Configure the time to wait before flushing the maintenance association end point (MEP) database, if no updates occur.                                                                                                           |
| <b>Options</b>                  | <i>minutes</i> —Time to wait, in minutes.                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul> |

## interface (OAM CFM, for EX Series Switch Only)

|                                 |                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interface (interface-name   ((ge-   xe-) (fpc/pic/port   fpc/pic/port.unit-number   fpc/pic/port.unit-number vlan vlan-id)));</code>                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id</a> ]                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                              |
| <b>Description</b>              | Configure IEEE 802.1ag Operation, Administration, and Management (OAM) Connectivity Fault Management (CFM) support for the specified interface.                                                                                    |
| <b>Options</b>                  | <b>interface-name</b> —Interface to which the MEP is attached. It can be a physical Ethernet interface or a logical interface. If the interface is a trunk interface, the VLAN associated with the interface must have a VLAN ID.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li> <li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li> </ul> |

## interval (EX Series Switch Only)

---

|                                 |                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | interval (10m   10s   1m   1s   100ms   10ms);                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name continuity-check</a> ]                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                           |
| <b>Description</b>              | Configure the time between continuity check messages.                                                                                                                                                                           |
| <b>Options</b>                  | <p>10m—10 minutes.</p> <p>10s—10 seconds.</p> <p>1m—1 minute.</p> <p>1s—1 second.</p> <p>100ms—100 milliseconds.</p> <p>10ms—10 milliseconds.</p>                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul> |

## iteration-period

|                                 |                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>iteration-period <i>iteration-period-value</i>;</code>                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles <i>profile-name</i> ]                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1.<br>Statement introduced in Junos OS Release 11.4 for EX Series switches.                                                                                 |
| <b>Description</b>              | Configure the iteration period, which is the maximum number of cycles per iteration (that is, the number of connections registered to an iterator cannot exceed this value).                            |
| <b>Options</b>                  | <i>iteration-period-value</i> —Maximum number of cycles per iteration.<br><b>Range:</b> 1 through 2000<br><b>Default:</b> 2000                                                                          |
| <b>Required Privilege Level</b> | Configure—To enter configuration mode.<br>Control—To modify any configuration.                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an Iterator Profile</a></li> <li>• <a href="#">Configuring an Iterator Profile on a Switch (CLI Procedure) on page 4067</a></li> </ul> |

## level (EX Series Switch Only)

|                                 |                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>level <i>number</i>;</code>                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> ]                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                              |
| <b>Description</b>              | Configure a number to be used in CFM messages to identify the maintenance association.                                                                                                                                             |
| <b>Options</b>                  | <i>number</i> —Number used to identify the maintenance domain to which the CFM message belongs.<br><b>Range:</b> 0 through 7                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li> <li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li> </ul> |

## linktrace (EX Series Switch Only)

---

|                                 |                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | linktrace {<br><b>age</b> (30m   10m   1m   30s   10s);<br><b>path-database-size</b> <i>path-database-size</i> ;<br>}                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <b>oam ethernet connectivity-fault-management</b> ]                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                           |
| <b>Description</b>              | Configure connectivity fault management linktrace parameters.<br><br>The remaining statements are explained separately.                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul> |

## loss-threshold (EX Series Switch Only)

---

|                                 |                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | loss-threshold <i>number</i> ;                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols <b>oam ethernet connectivity-fault-management maintenance-domain</b> <i>domain-name</i> <b>maintenance-association</b> <i>ma-name</i> <b>continuity-check</b> ]                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                           |
| <b>Description</b>              | Configure the number of continuity check messages that can be lost before the remote MEP is marked as down.                                                                                                                     |
| <b>Options</b>                  | <i>number</i> —Number of continuity check messages that can be lost before the remote MEP is marked down.                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul> |



## maintenance-association (EX Series Switch Only)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> maintenance-association <i>ma-name</i> {   continuity-check {     hold-interval <i>minutes</i>;     interface-status-tlv;     interval (10m   10s   1m   1s   100ms);     loss-threshold <i>number</i>;     port-status-tlv;   }   mep <i>mep-id</i> {     auto-discovery;     direction down;     interface <i>interface-name</i>;     remote-mep <i>mep-id</i> {       action-profile <i>profile-name</i>;     }   } } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet connectivity-fault-management maintenance-domain domain-name</a> ]                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure the name of the maintenance association in IEEE-compliant format.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b><i>ma-name</i></b>—The name of the maintenance association within the maintenance domain.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li> <li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li> </ul>                                                                                                                                                                                                    |

## **maintenance-domain (EX Series Switch Only)**

---

**Syntax**    `maintenance-domain domain-name {  
                  level number;  
                  mip-half-function (none | default | explicit);  
                  name-format (character-string | none | dns | mac+2oct);  
                  maintenance-association ma-name {  
                      continuity-check {  
                          hold-interval minutes;  
                          interface-status-tlv;  
                          interval (10m | 10s | 1m | 1s | 100ms);  
                          loss-threshold number;  
                          port-status-tlv;  
                        }  
                      mep mep-id {  
                          auto-discovery;  
                          direction down;  
                          interface interface-name;  
                          remote-mep mep-id {  
                              action-profile profile-name;  
                          }  
                        }  
                  }  
                }`

**Hierarchy Level**    [edit protocols [oam](#) [ethernet](#) [connectivity-fault-management](#) ]

**Release Information**    Statement introduced in Junos OS Release 10.2 for EX Series switches.

**Description**    Configure the name of the maintenance domain in IEEE-compliant format.

**Options**    *domain-name*—The name for the maintenance domain.

The remaining statements are explained separately.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Ethernet OAM Connectivity Fault Management \(CLI Procedure\) on page 4058](#)
- [Junos OS Network Interfaces Configuration Guide](#)

## measurement-type (OAM LFM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | measurement-type two-way-delay;                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">connectivity-fault-management</a> <a href="#">performance-monitoring</a> <a href="#">sla-iterator-profiles</a> <i>profile-name</i> ]                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure the measurement type for the service-level agreement (SLA) frames. An SLA frame is a type of packet used to measure frame loss in Ethernet connections.                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <b>two-way-delay</b> —Use Y.1731-compliant two-way ETH-DM frames to measure frame loss.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | Configure—To enter configuration mode.<br>Control—To modify any configuration.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements (CLI Procedure) on page 4065</a></li> <li>• <a href="#">Configuring One-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure) on page 4066</a></li> <li>• <a href="#">Configuring Two-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure) on page 4069</a></li> </ul> |

## mep (EX Series Switch Only)

---

|                                 |                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>mep mep-id {<br/>  auto-discovery;<br/>  direction down;<br/>  interface interface-name;<br/>  remote-mep mep-id {<br/>    action-profile profile-name;<br/>  }<br/>}</pre>                                                |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name</a> ]                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                           |
| <b>Description</b>              | Configure the numeric identifier of the maintenance association end point (MEP) within the maintenance association.                                                                                                             |
| <b>Options</b>                  | <p><b>mep-id</b>—Numeric identifier of the MEP.</p> <p><b>Range:</b> 1 through 8191</p> <p>The remaining statements are explained separately.</p>                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul> |

## mip-half-function (EX Series Switch Only)

|                            |                                                                                                             |
|----------------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | mip-half-function (none   default   explicit);                                                              |
| <b>Hierarchy Level</b>     | [edit protocols <a href="#">oam ethernet connectivity-fault-management maintenance-domain domain-name</a> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                       |
| <b>Description</b>         | Specify the OAM Ethernet CFM maintenance domain MIP half functions.                                         |



**NOTE:** Whenever a MIP is configured, the MIP half function value for all maintenance domains and maintenance associations must be the same.

|                                 |                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <p><b>none</b>—Specify to not use the mip-half-function.</p> <p><b>default</b>—Specify to use the default mip-half-function.</p> <p><b>explicit</b>—Specify an explicit mip-half-function.</p>                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li> <li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li> </ul> |

## name-format (EX Series Switch Only)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | name-format (character-string   none   dns   mac+2oct);                                                                                                                                                                                                                                                                                                                                                                                       |
| Hierarchy Level          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">connectivity-fault-management</a> <a href="#">maintenance-domain</a> <i>domain-name</i> ]                                                                                                                                                                                                                                                                            |
| Release Information      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                         |
| Description              | Specify the format of the maintenance domain name.                                                                                                                                                                                                                                                                                                                                                                                            |
| Options                  | <p><b>character-string</b>—The name is an ASCII character string.</p> <p><b>none</b>—Name format <b>none</b> means that maintenance domain name is not used.</p> <p><b>dns</b>—Name is in domain name service (DNS) format. For example: www.juniper.net.</p> <p><b>mac+2oct</b>—Name is the MAC address plus a two-octet maintenance association identifier. For example: 08:00:22:33:44:55.100.</p> <p><b>Default:</b> character-string</p> |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                           |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul>                                                                                                                                                                                                               |

## path-database-size (EX Series Switch Only)

---

|                          |                                                                                                                                                                                                                                 |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | path-database-size <i>path-database-size</i> ;                                                                                                                                                                                  |
| Hierarchy Level          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">connectivity-fault-management</a> <a href="#">linktrace</a> ]                                                                                          |
| Release Information      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                           |
| Description              | Specify the number of linktrace reply entries to be stored per linktrace request.                                                                                                                                               |
| Options                  | <p><b>path-database-size</b>—Database size (number of entries stored per request).</p> <p><b>Range:</b> 1 through 500</p> <p><b>Default:</b> 100</p>                                                                            |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                             |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul> |

## performance-monitoring (OAM LFM)

```
Syntax performance-monitoring {
    sla-iterator-profiles {
        profile-name {
            calculation-weight {
                delay delay-value;
                delay-variation delay-variation-value;
            }
            cycle-time cycle-time-value;
            iteration-period iteration-period-value;
            measurement-type two-way-delay;
            passive;
        }
    }
}
```

**Hierarchy Level** [edit protocols [oam ethernet connectivity-fault-management](#)]

**Release Information** Statement introduced in Junos OS Release 11.4 for EX Series switches.

**Description** Specify performance monitoring support for Ethernet frame delay measurement.

The remaining statements are explained separately.

**Required Privilege Level** Configure—To enter configuration mode.  
Control—To modify any configuration.

**Related Documentation**

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 4065](#)
- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4066](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 4069](#)

## remote-mep (EX Series Switch Only)

---

|                                 |                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>remote-mep mep-id {<br/>    action-profile profile-name;<br/>}</code>                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">connectivity-fault-management</a> <a href="#">maintenance-domain</a> <a href="#">domain-name</a> <a href="#">maintenance-association</a> <a href="#">ma-name</a> <a href="#">mep</a> <a href="#">mep-id</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                  |
| <b>Description</b>              | Specify the numeric identifier of the remote maintenance association end point (MEP) within the maintenance association.                                                                                                                                                               |
| <b>Options</b>                  | <p><b>mep-id</b>—Specify the numeric identifier of the MEP.</p> <p><b>Range:</b> 1 through 8191</p> <p>The remaining statement is explained separately.</p>                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li></ul>                                                        |



## sla-iterator-profiles (OAM LFM)

|                                 |                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>sla-iterator-profiles {   profile-name {     calculation-weight {       delay delay-value;       delay-variation delay-variation-value;     }     cycle-time cycle-time-value;     iteration-period iteration-period-value;     measurement-type two-way-delay;     passive;   } }</pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">connectivity-fault-management</a> <a href="#">performance-monitoring</a> ]                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4 for EX Series switches.                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure an iterator application and specify the iterator profile options.                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><i>profile-name</i>—Name of the iterator profile.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>Configure—To enter configuration mode.</p> <p>Control—To modify any configuration.</p>                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an Iterator Profile on a Switch (CLI Procedure) on page 4067</a></li> </ul>                                                                                                                                                  |

## Configuration Statements: Ethernet OAM Link Fault Management

- [action \(OAM LFM\) on page 4218](#)
- [action-profile on page 4219](#)
- [allow-remote-loopback on page 4220](#)
- [ethernet \(OAM LFM\) on page 4221](#)
- [event \(OAM LFM\) on page 4224](#)
- [event-thresholds on page 4224](#)
- [frame-error on page 4225](#)
- [frame-period on page 4225](#)
- [frame-period-summary on page 4226](#)
- [interface \(OAM LFM\) on page 4227](#)
- [link-adjacency-loss on page 4228](#)
- [link-discovery on page 4228](#)
- [link-down on page 4229](#)

- [link-event-rate on page 4229](#)
- [link-fault-management on page 4230](#)
- [negotiation-options on page 4231](#)
- [no-allow-link-events on page 4231](#)
- [oam on page 4232](#)
- [pdu-interval on page 4234](#)
- [pdu-threshold on page 4235](#)
- [remote-loopback on page 4235](#)
- [symbol-period on page 4236](#)
- [syslog \(OAM LFM\) on page 4236](#)

---

## action (OAM LFM)

---

|                                 |                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>action {<br/>  syslog;<br/>  link-down;<br/>}</pre>                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management</a> ]                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                   |
| <b>Description</b>              | <p>Define the action or actions to be taken when the OAM link fault management (LFM) fault event occurs.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li></ul>                          |

## action-profile

|                                 |                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> action-profile <i>profile-name</i>;   action {     syslog;     link-down;   }   event {     link-adjacency-loss;     link-event-rate {       frame-error <i>count</i>;       frame-period <i>count</i>;       frame-period-summary <i>count</i>;       symbol-period <i>count</i>;     }   } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management</a> ]                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure an Ethernet OAM link fault management (LFM) action profile by specifying a profile name.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                       |
| <b>Options</b>                  | <i>profile-name</i> —Name of the action profile.                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure)</a> on page 4062</li> </ul>                                                                                                                                                           |

## allow-remote-loopback

---

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | allow-remote-loopback;                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet link-fault-management interface</a> <i>interface-name</i> ]                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                               |
| <b>Description</b>              | Advertise that the interface is capable of getting into loopback mode. Enable remote loopback in Ethernet OAM link fault management (LFM) on all Ethernet interfaces or the specified interface on the EX Series switch.                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027</a></li><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li></ul> |

## ethernet (OAM LFM)

```
Syntax ethernet {
    connectivity-fault-management {
        action-profile profile-name {
            action {
                interface-down;
            }
            default-actions {
                interface-down;
            }
            event {
                adjacency-loss;
            }
        }
    }
    esp-traceoptions {
        file filename <files number> <no-stamp> <replace> <size size> <world-readable |
            no-world-readable>;
        flag (all | error | esp | interface | krt | lib | normal | task | timer);
    }
    linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
    }
    maintenance-domain domain-name {
        level number;
        mip-half-function (none | default | explicit);
        name-format (character-string | none | dns | mac+2oct);
        maintenance-association ma-name {
            continuity-check {
                hold-interval minutes;
                interface-status-tlv;
                interval (10m | 10s | 1m | 1s | 100ms);
                loss-threshold number;
                port-status-tlv;
            }
            mep mep-id {
                auto-discovery;
                direction down;
                interface interface-name;
                priority
                remote-mep mep-id {
                    action-profile profile-name;
                    sla-iterator-profile profile-name {
                        data-tlv-size size;
                        iteration-count count-value;
                        priority priority-value;
                    }
                }
            }
        }
        short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
    }
}
performance-monitoring {
```

```
sla-iterator-profiles {
  profile-name {
    calculation-weight {
      delay delay-value;
      delay-variation delay-variation-value;
    }
    cycle-time cycle-time-value;
    iteration-period iteration-period-value;
    measurement-type two-way-delay;
    passive;
  }
}
}
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
}
}
link-fault-management {
  action-profile profile-name;
  action {
    syslog;
    link-down;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
  }
}
interface interface-name {
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
}
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
```

```

    }
  }
}

```

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.<br><b>connectivity-fault-management</b> introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) support for Ethernet interfaces on EX Series switches or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support on the switches.<br><br>The remaining statements are explained separately.                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027</a></li> <li>• <a href="#">Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches on page 4023</a></li> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li> <li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 4058</a></li> </ul> |

## event (OAM LFM)

---

|                                 |                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>event {<br/>  link-adjacency-loss;<br/>  link-event-rate {<br/>    frame-error count;<br/>    frame-period count;<br/>    frame-period-summary count;<br/>    symbol-period count;<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile</a> <i>profile-name</i> ]                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                        |
| <b>Description</b>              | Configure link events in an action profile for Ethernet OAM link fault management (LFM).<br><br>The remaining statements are explained separately.                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li></ul>                                                               |

## event-thresholds

---

|                                 |                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>event-thresholds {<br/>  frame-error count;<br/>  frame-period count;<br/>  frame-period-summary count;<br/>  symbol-period count;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management interface</a> <i>interface-name</i> ]                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                |
| <b>Description</b>              | Configure threshold limit values for link events in periodic OAM PDUs.<br><br>The remaining statements are explained separately.                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li></ul>       |



## frame-error

---

|                                 |                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>frame-error count;</code>                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management event link-event-rate</a> ],<br>[edit protocols <a href="#">oam ethernet link-fault-management interface interface-name event-thresholds</a> ]                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the threshold value for sending frame error events or taking the action specified in the action profile.<br><br>Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value. |
| <b>Options</b>                  | <i>count</i> —Threshold count in seconds for frame error events.<br><b>Range:</b> 1 through 100 seconds<br><b>Default:</b> 1 second                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li> </ul>                                                                                                                       |

## frame-period

---

|                                 |                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>frame-period count;</code>                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management event link-event-rate</a> ],<br>[edit protocols <a href="#">oam ethernet link-fault-management interface interface-name event-thresholds</a> ]                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                            |
| <b>Description</b>              | Configure the number of frame errors within the last N frames that has exceeded a threshold.<br><br>Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value. |
| <b>Options</b>                  | <i>count</i> —Threshold count in seconds for frame error events.<br><b>Range:</b> 1 through 100 seconds                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li> </ul>                                                                                                 |

## frame-period-summary

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>frame-period-summary count;</code>                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management event link-event-rate</a> ],<br>[edit protocols <a href="#">oam ethernet link-fault-management interface interface-name event-thresholds</a> ]                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Configure the threshold value for sending frame period summary error events or taking the action specified in the action profile.</p> <p>An errored frame second is any 1-second period that has at least one errored frame. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period.</p> |
| <b>Options</b>                  | <p><i>count</i>—Threshold count in seconds for frame period summary error events.</p> <p><b>Range:</b> 1 through 100 seconds</p>                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li></ul>                                                                                                                                                                                                                           |

## interface (OAM LFM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interface <i>interface-name</i> {   link-discovery (active   passive);   pdu-interval <i>interval</i>;   pdu-threshold <i>threshold-value</i>;   remote-loopback;   event-thresholds {     frame-error <i>count</i>;     frame-period <i>count</i>;     frame-period-summary <i>count</i>;     symbol-period <i>count</i>;   }   negotiation-options {     allow-remote-loopback;     no-allow-link-events;   } } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">link-fault-management</a> ]                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface to be enabled for IEEE 802.3ah OAM link fault management (LFM) support.                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027</a></li> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li> </ul>                                                                                                                                                          |

## link-adjacency-loss

---

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-adjacency-loss;                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile event</a> ]                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                               |
| <b>Description</b>              | Configure <b>loss of adjacency</b> event with the IEEE 802.3ah link fault management (LFM) peer. When included, the loss of adjacency event triggers the action specified under the <a href="#">action</a> statement.                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027</a></li><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li></ul> |

## link-discovery

---

|                                 |                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-discovery (active   passive);                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management interface interface-name</a> ]                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on an interface. Link monitoring is done when the interface sends periodic OAM PDUs.         |
| <b>Options</b>                  | <p><i>active</i>—In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality.</p> <p><i>passive</i>—In passive mode, the peer initiates the discovery process.</p> <p>Once the discovery process is initiated, both sides participate in discovery.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li></ul>                                                                                                                                                                                  |

## link-down

---

|                                 |                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-down;                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile action</a> ]                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                            |
| <b>Description</b>              | Mark the interface as down for transit traffic.                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li> </ul> |

## link-event-rate

---

|                                 |                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>link-event-rate {   frame-error count;   frame-period count;   frame-period-summary count;   symbol-period count; }</pre>                  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile event</a> ]                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                            |
| <b>Description</b>              | <p>Configure the number of link fault management (LFM) events per second.</p> <p>The remaining statements are explained separately.</p>         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li> </ul> |

## link-fault-management

---

**Syntax**    link-fault-management {  
              [action-profile](#) *profile-name*;  
              action {  
                  syslog;  
                  link-down;  
              }  
              event {  
                  link-adjacency-loss;  
                  link-event-rate {  
                      frame-error *count*;  
                      frame-period *count*;  
                      frame-period-summary *count*;  
                      symbol-period *count*;  
                  }  
              }  
              interface *interface-name* {  
                  link-discovery (active | passive);  
                  pdu-interval *interval*;  
                  pdu-threshold *threshold-value*;  
                  remote-loopback;  
                  event-thresholds {  
                      frame-error *count*;  
                      frame-period *count*;  
                      frame-period-summary *count*;  
                      symbol-period *count*;  
                  }  
                  negotiation-options {  
                      allow-remote-loopback;  
                      no-allow-link-events;  
                  }  
              }  
              }  
              }

**Hierarchy Level**    [edit protocols [oam](#) [ethernet](#)]

**Release Information**    Statement introduced in Junos OS Release 9.4 for EX Series switches.

**Description**    Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

The remaining statements are explained separately.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 4062](#)

## negotiation-options

---

|                                 |                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | negotiation-options {<br>allow-remote-loopback;<br>no-allow-link-events;<br>}                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols oam ethernet link-fault-management interface <i>interface-name</i> ]                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                     |
| <b>Description</b>              | Enable and disable IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) features for Ethernet interfaces.<br><br>The remaining statements are explained separately. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li> </ul>                                                          |

## no-allow-link-events

---

|                                 |                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-allow-link-events;                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols oam ethernet link-fault-management interface <i>interface-name</i> negotiation-options]                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                            |
| <b>Description</b>              | Disable the sending of link event TLVs.                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li> </ul> |

## oam

---

```
Syntax  oam {
    ethernet {
        connectivity-fault-management {
            action-profile profile-name {
                action {
                    interface-down;
                }
                default-actions {
                    interface-down;
                }
                event {
                    adjacency-loss;
                }
            }
        }
        linktrace {
            age (30m | 10m | 1m | 30s | 10s);
            path-database-size path-database-size;
        }
        maintenance-domain domain-name {
            level number;
            mip-half-function (none | default | explicit);
            name-format (character-string | none | dns | mac+2oct);
            maintenance-association ma-name {
                continuity-check {
                    hold-interval minutes;
                    interface-status-tlv;
                    interval (10m | 10s | 1m | 1s | 100ms);
                    loss-threshold number;
                    port-status-tlv;
                }
                mep mep-id {
                    auto-discovery;
                    direction down;
                    interface interface-name;
                    remote-mep mep-id {
                        action-profile profile-name;
                    }
                }
            }
        }
    }
    performance-monitoring {
        sla-iterator-profiles {
            profile-name {
                calculation-weight {
                    delay delay-value;
                    delay-variation delay-variation-value;
                }
                cycle-time cycle-time-value;
                iteration-period iteration-period-value;
                measurement-type two-way-delay;
                passive;
            }
        }
    }
}
```



```

    }
  }
}
link-fault-management {
  action-profile profile-name;
  action {
    syslog;
    link-down;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
  }
}
interface interface-name {
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
}
}
}

```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 9.4 for EX Series switches.  
**connectivity-fault-management** introduced in Junos OS Release 10.2 for EX Series switches.

**Description** Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support for Ethernet interfaces on EX Series switches or configure connectivity fault management (CFM) for IEEE 802.lag Operation, Administration, and Management (OAM) support on the switches.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
 interface-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027](#)
  - [Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches on page 4023](#)
  - [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 4062](#)
  - [Configuring Ethernet OAM Connectivity Fault Management \(CLI Procedure\) on page 4058](#)

---

## pdu-interval

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>pdu-interval <i>interval</i>;</code>                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <b>oam</b> <b>ethernet link-fault-management interface</b> <i>interface-name</i> ]                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                               |
| <b>Description</b>              | Specify the periodic OAM PDU sending interval for fault detection. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support.                                                                               |
| <b>Options</b>                  | <i>interval</i> —Periodic OAM PDU sending interval.<br><b>Range:</b> 400 through 1000 milliseconds<br><b>Default:</b> 1000 milliseconds                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027</a></li><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li></ul> |

## pdu-threshold

---

|                                 |                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>pdu-threshold <i>threshold-value</i>;</code>                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management interface <i>interface-name</i></a> ]                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                               |
| <b>Description</b>              | Configure how many protocol data units (PDUs) are missed before declaring the peer lost in Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces. |
| <b>Options</b>                  | <p><b><i>threshold-value</i></b> —Number of PDUs missed before declaring the peer lost.</p> <p><b>Range:</b> 3 through 10 PDUs</p> <p><b>Default:</b> 3 PDUs</p>                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li> </ul>                                    |

## remote-loopback

---

|                                 |                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>remote-loopback;</code>                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management interface <i>interface-name</i></a> ]                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                  |
| <b>Description</b>              | Set the data terminal equipment (DTE) in loopback mode. Remove the statement from the configuration to take the DTE out of loopback mode. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support.           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027</a></li> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li> </ul> |

## symbol-period

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>symbol-period count;</code>                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile</a> <i>profile-name</i> ; <a href="#">event link-event-rate</a> ] ,<br>[edit protocols <a href="#">oam ethernet link-fault-management interface</a> <i>interface-name</i> <a href="#">event-thresholds</a> ]                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Configure the threshold for sending symbol period events or taking the action specified in the action profile.</p> <p>Symbol code errors occur on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period. You cannot configure the default value to a different value.</p> |
| <b>Options</b>                  | <p><i>count</i>—Threshold count in seconds for symbol period events.</p> <p><b>Range:</b> 1 through 100 seconds</p>                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li></ul>                                                                                                                                                                                                                                |

## syslog (OAM LFM)

---

|                                 |                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>syslog;</code>                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile</a> <i>profile-name</i> ; <a href="#">action</a> ]              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.                                                                          |
| <b>Description</b>              | Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event.            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li></ul> |

## Configuration Statements: RPM

---

- [data-fill on page 4238](#)
- [data-size on page 4239](#)
- [destination-port on page 4240](#)

- [dscp-code-point on page 4241](#)
- [hardware-timestamp on page 4242](#)
- [history-size on page 4243](#)
- [moving-average-size on page 4243](#)
- [one-way-hardware-timestamp on page 4244](#)
- [port \(RPM\) on page 4244](#)
- [probe on page 4245](#)
- [probe-count on page 4246](#)
- [probe-interval on page 4246](#)
- [probe-limit on page 4247](#)
- [probe-server on page 4248](#)
- [probe-type on page 4249](#)
- [routing-instance on page 4250](#)
- [routing-instances on page 4250](#)
- [rpm \(Interfaces\) on page 4251](#)
- [source-address \(Services\) on page 4251](#)
- [tcp on page 4252](#)
- [test on page 4253](#)
- [test-interval on page 4254](#)
- [thresholds on page 4255](#)
- [traps on page 4256](#)
- [udp on page 4257](#)

## data-fill

---

|                                 |                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>data-fill data;</code>                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit services rpm bgp],<br>[edit services rpm <b>probe</b> owner <b>test</b> test-name]                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 9.3. for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes. The <b>data-fill</b> statement is not valid with the <b>http-get</b> or <b>http-metadata-get</b> probe types.          |
| <b>Options</b>                  | <b>data</b> —A hexadecimal value; for example, <b>0-9</b> , <b>A-F</b> .                                                                                                                                            |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring BGP Neighbor Discovery Through RPM</i></li><li>• <i>Configuring RPM Probes</i></li></ul>                                                                     |

## data-size

|                            |                                                                                                                                                                                                                     |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>data-size size;</code>                                                                                                                                                                                        |
| <b>Hierarchy Level</b>     | [edit services rpm bgp],<br>[edit services rpm <b>probe</b> owner <b>test</b> test-name]                                                                                                                            |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>         | Specify the size of the data portion of ICMP probes. The <b>data-size</b> statement is not valid with the <b>http-get</b> or <b>http-metadata-get</b> probe type.                                                   |
| <b>Options</b>             | <b>data</b> —The size can be from 0 through 65400<br><b>Default:</b> 0                                                                                                                                              |



**NOTE:** If you configure the hardware timestamp feature (see *Configuring RPM Timestamping*):

- The **data-size** default value is 32 bytes and 32 is the minimum value for explicit configuration. The UDP timestamp probe type is an exception; it requires a minimum data size of 52 bytes.
- The **data-size** must be at least 100 bytes smaller than the default MTU of the interface of the RPM client interface.

|                                 |                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring BGP Neighbor Discovery Through RPM</i></li> </ul>            |

## destination-port

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>destination-port port;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit services rpm bgp],<br>[edit services rpm <b>probe</b> owner <b>test</b> test-name]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Specify the User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) port to which a probe is sent. This statement is used only for TCP or UDP probe types.</p> <p>The value for the <b>destination-port</b> can be only 7 when you configure along with hardware timestamping. A constraint check prevents you for configuring any other value for the destination port in this case.</p> <p>This constraint does not apply when you are using one-way hardware timestamping along with <b>destination-port</b> and either <b>probe-type udp-ping</b> or <b>probe-type udp-ping-timestamp</b>.</p> |
| <b>Options</b>                  | <b>port</b> —The port number can be 7 or from 49,160 to 65,535.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>Configuring BGP Neighbor Discovery Through RPM</li><li>Configuring RPM Probes</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |



## dscp-code-point

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>dscp-code-point <i>dscp-bits</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>     | [edit services rpm <b>probe</b> <i>owner</i> <b>test</b> <i>test-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release for PTX Series Packet Transport Routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>         | Specify the value of the Differentiated Services (DiffServ) field within the IP header. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>             | <p><b><i>dscp-bits</i></b>—A valid 6-bit pattern; for example, 001111, or one of the following configured DSCP aliases:</p> <ul style="list-style-type: none"> <li>• <b>af11</b>—Default: 001010</li> <li>• <b>af12</b>—Default: 001100</li> <li>• <b>af13</b>—Default: 001110</li> <li>• <b>af21</b>—Default: 010010</li> <li>• <b>af22</b>—Default: 010100</li> <li>• <b>af23</b> —Default: 010110</li> <li>• <b>af31</b> —Default: 011010</li> <li>• <b>af32</b> —Default: 011100</li> <li>• <b>af33</b> —Default: 011110</li> <li>• <b>af41</b> —Default: 100010</li> <li>• <b>af42</b> —Default:100100</li> <li>• <b>af43</b> —Default:100110</li> <li>• <b>be</b>—Default: 000000</li> <li>• <b>cs1</b>—Default: 001000</li> <li>• <b>cs2</b>—Default: 010000</li> <li>• <b>cs3</b>—Default: 011000</li> <li>• <b>cs4</b>—Default: 100000</li> <li>• <b>cs5</b>—Default: 101000</li> <li>• <b>cs6</b>—Default: 110000</li> <li>• <b>cs7</b>—Default: 111000</li> <li>• <b>ef</b>—Default: 101110</li> <li>• <b>nc1</b>—Default: 110000</li> </ul> |

- **nc2**—Default: 111000

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring RPM Probes*

---

## hardware-timestamp

---

**Syntax** hardware-timestamp;

**Hierarchy Level** [edit services rpm **probe** owner **test** test-name]

**Release Information** Statement introduced in Junos OS Release 8.1.  
Statement applied to MX Series routers in Junos OS Release 10.0.  
Statement introduced in Junos OS Release 10.3 for EX Series switches.

**Description** On MX Series routers, on M-320 routers using the Enhanced Queuing MPC, and on EX Series switches only, enable timestamping of RPM probe messages in the Packet Forwarding Engine host processor. This feature is supported only with **icmp-ping**, **icmp-ping-timestamp**, **udp-ping**, and **udp-ping-timestamp** probe types.

When you configure either **probe-type udp-ping** or **probe-type udp-ping-timestamp** along with the **hardware-timestamp** command, the value for the **destination-port** can be only 7. A constraint check prevents you for configuring any other value for the destination port in this case.

This constraint does not apply when you are configuring **one-way-hardware-timestamp**.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring RPM Timestamping*

## history-size

|                                 |                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>history-size size;</code>                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit services rpm bgp],<br>[edit services rpm <b>probe</b> owner <b>test</b> test-name]                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Specify the number of stored history entries.                                                                                                                                                                       |
| <b>Options</b>                  | <b>size</b> —A value from 0 to 512.<br><b>Default:</b> 50                                                                                                                                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring BGP Neighbor Discovery Through RPM</i></li> <li>• <i>Configuring RPM Probes</i></li> </ul>                                                                  |

## moving-average-size

|                                 |                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>moving-average-size number;</code>                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit services rpm bgp],<br>[edit services rpm <b>probe</b> owner <b>test</b> test-name]                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement Introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Enable statistical calculation operations to be performed across a configurable number of the most recent samples.                                                                                              |
| <b>Options</b>                  | <b>number</b> —Number of samples to be used in calculations.<br><b>Range:</b> 0 through 255                                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring RPM Probes</i></li> </ul>                                                                                                                               |

## one-way-hardware-timestamp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>one-way-hardware-timestamp;</code>                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit services rpm <a href="#">probe owner test test-name</a> ]                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.                                                                                                                                                                                                                                 |
| <b>Description</b>              | Enable timestamping of RPM probe messages for one-way delay and jitter measurements. You must configure this statement along with the <b>destination-interface</b> statement to invoke timestamping. This feature is supported only with <b>icmp-ping</b> , <b>icmp-ping-timestamp</b> , <b>udp-ping</b> , and <b>udp-ping-timestamp</b> probe types. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring RPM Timestamping</i></li><li>• <i>destination-interface</i></li><li>• <a href="#">hardware-timestamp on page 4242</a></li></ul>                                                                                                                                                                |

## port (RPM)

---

|                                 |                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>port number;</code>                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit services rpm <a href="#">probe-server (tcp   udp)</a> ]                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Specify the port number for the probe server.                                                                                                                                                                       |
| <b>Options</b>                  | <b>number</b> —Port number for the probe server. The value can be 7 or 49,160 through 65,535.                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring RPM Receiver Servers</i></li></ul>                                                                                                                           |

## probe

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> probe owner {   test test-name {     data-fill data;     data-size size;     destination-interface interface-name;     destination-port port;     dscp-code-point dscp-bits;     hardware-timestamp;     history-size size;     moving-average-size number;     one-way-hardware-timestamp;     probe-count count;     probe-interval seconds;     probe-type type;     routing-instance instance-name;     source-address address;     target (url   address);     test-interval interval;     thresholds thresholds;     traps traps;   } } </pre> |
| <b>Hierarchy Level</b>          | [edit services rpm]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Specify an owner name. The owner name combined with the test name represent a single RPM configuration instance.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>owner</b>—Specify an owner name up to 32 characters in length.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p><b>system</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Configuring RPM Probes</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## probe-count

---

|                                 |                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>probe-count count;</code>                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit services rpm bgp],<br>[edit services rpm <b>probe</b> owner <b>test</b> test-name]                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Specify the number of probes within a test.                                                                                                                                                                         |
| <b>Options</b>                  | <b>count</b> —A value from 1 through 15.                                                                                                                                                                            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring BGP Neighbor Discovery Through RPM</i></li><li>• <i>Configuring RPM Probes</i></li></ul>                                                                     |

## probe-interval

---

|                                 |                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>probe-interval interval;</code>                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit services rpm bgp],<br>[edit services rpm <b>probe</b> owner <b>test</b> test-name]                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Specify the time to wait between sending packets, in seconds.                                                                                                                                                       |
| <b>Options</b>                  | <b>interval</b> —Number of seconds, from 1 through 255.                                                                                                                                                             |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring BGP Neighbor Discovery Through RPM</i></li><li>• <i>Configuring RPM Probes</i></li></ul>                                                                     |

---

## probe-limit

---

|                                 |                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>probe-limit <i>limit</i>;</code>                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit services rpm]                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Configure the maximum number of concurrent probes allowed.                                                                                                                                                          |
| <b>Options</b>                  | <i>limit</i> —Maximum number of concurrent probes allowed.<br><b>Range:</b> 1 through 500 (PTX Series Packet Transport Routers only) 1 through 200<br><b>Default:</b> 100                                           |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Limiting the Number of Concurrent RPM Probes</i></li></ul>                                                                                                               |

## probe-server

---

**Syntax**    `probe-server {  
              tcp {  
                  destination-interface interface-name;  
                  port number;  
              }  
              udp {  
                  destination-interface interface-name;  
                  port number;  
              }  
          }`

**Hierarchy Level**    [edit services rpm]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.3 for EX Series switches.  
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.

**Description**    Specify the server to act as a receiver for the probes.  
  
The remaining statements are explained separately.



**NOTE:** The `destination-interface` statement is not supported on PTX Series routers.

---

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**    • *Configuring RPM Receiver Servers*



## probe-type

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>probe-type type;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit services rpm bgp],<br>[edit services rpm <b>probe</b> owner <b>test</b> test-name]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Specify the packet and protocol contents of a probe.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>type</b>—Specify one of the following probe type values:</p> <ul style="list-style-type: none"> <li>• <b>http-get</b>—(Not available at the [edit services rpm bgp] hierarchy level.) Sends a Hypertext Transfer Protocol (HTTP) get request to a target URL.</li> <li>• <b>http-metadata-get</b>—(Not available at the [edit services rpm bgp] hierarchy level.) Sends an HTTP get request for metadata to a target URL.</li> <li>• <b>icmp-ping</b>—Sends ICMP echo requests to a target address.</li> <li>• <b>icmp-ping-timestamp</b>—Sends ICMP timestamp requests to a target address.</li> <li>• <b>tcp-ping</b>—Sends TCP packets to a target.</li> <li>• <b>udp-ping</b>—Sends UDP packets to a target.</li> <li>• <b>udp-ping-timestamp</b>—Sends UDP timestamp requests to a target address.</li> </ul> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring BGP Neighbor Discovery Through RPM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## routing-instance

---

|                                 |                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>routing-instance <i>instance-name</i>;</code>                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit services rpm <b>probe</b> owner <b>test</b> <i>test-name</i>]</code>                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Specify the routing instance used by the probes.                                                                                                                                                                    |
| <b>Options</b>                  | <b><i>instance-name</i></b> —A routing instance configured at the <code>[edit routing-instance]</code> hierarchy level.<br><b>Default:</b> Internet routing table <code>inet.0</code> .                             |
| <b>Required Privilege Level</b> | <code>interface</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring RPM Probes</i></li></ul>                                                                                                                                     |

## routing-instances

---

|                                 |                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>routing-instances <i>instance-name</i>;</code>                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <code>[edit services rpm <b>bgp</b>],</code><br><code>[edit services rpm <b>bgp</b> logical-system <i>logical-system-name</i>]</code>                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Specify the routing instance used by the probes.                                                                                                                                                                |
| <b>Options</b>                  | <b><i>instance-name</i></b> —A routing instance configured at the <code>[edit routing-instances]</code> hierarchy level.<br><b>Default:</b> Internet routing table <code>inet.0</code> .                        |
| <b>Required Privilege Level</b> | <code>interface</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring BGP Neighbor Discovery Through RPM</i></li></ul>                                                                                                         |

## rpm (Interfaces)

|                                 |                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rpm (client   server);</code>                                                                                     |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ]                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.   |
| <b>Description</b>              | Associate an RPM client (router or switch that originates RPM probes) or RPM server with a specified interface.         |
| <b>Options</b>                  | <i>client</i> —Identifier for RPM client router or switch.<br><i>server</i> —Identifier for RPM server.                 |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring RPM Timestamping</i></li> </ul>                                 |

## source-address (Services)

|                                 |                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-address <i>address</i>;</code>                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit services rpm <i>probe owner test test-name</i> ]                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Specify the source IP address used for probes. If the source IP address is not one of the router's or switch's assigned addresses, the packet will use the outgoing interface's address as its source.              |
| <b>Options</b>                  | <i>address</i> —Valid IP address.                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring RPM Probes</i></li> </ul>                                                                                                                                   |

## tcp

---

|                                 |                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>tcp {<br/>    destination-interface <i>interface-name</i>;<br/>    port <i>port</i>;<br/>}</pre>                                                                                                               |
| <b>Hierarchy Level</b>          | [edit services rpm <a href="#">probe-server</a> ]                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Specify the port information for the TCP server.<br><br>The remaining statements are explained separately.                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring RPM Receiver Servers</i></li></ul>                                                                                                                           |

## test

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>test test-name {   data-fill data;   data-size size;   destination-interface interface-name;   destination-port port;   dscp-code-point dscp-bits;   hardware-timestamp;   history-size size;   moving-average-size number;   one-way-hardware-timestamp;   probe-count count;   probe-interval seconds;   probe-type type;   routing-instance instance-name;   source-address address;   target (url url   address address);   test-interval interval;   thresholds thresholds;   traps traps; }</pre> |
| <b>Hierarchy Level</b>          | [edit services rpm <b>probe</b> owner]                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.</p>                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify the range of probes over which the standard deviation, average, and jitter are calculated. The test name combined with the owner name represent a single RPM configuration instance.                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>test-name</b>—Specify a test name. The name can be up to 32 characters in length.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Configuring RPM Probes</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## test-interval

---

|                                 |                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>test-interval frequency;</code>                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit services rpm bgp],<br>[edit services rpm <b>probe</b> owner <b>test</b> test-name]                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Specify the time to wait between tests, in seconds.                                                                                                                                                                 |
| <b>Options</b>                  | <b>frequency</b> —Number of seconds, from 0 through 86400.                                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring BGP Neighbor Discovery Through RPM</i></li><li>• <i>Configuring RPM Probes</i></li></ul>                                                                     |

## thresholds

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>thresholds thresholds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit services rpm <b>probe</b> owner <b>test</b> test-name]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Packet Series Transport Routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify thresholds used for the probes. A system log message is generated when the configured threshold is exceeded. Likewise, an SNMP trap (if configured) is generated when a threshold is exceeded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>thresholds</b>—Specify one or more threshold measurements. The following options are supported:</p> <ul style="list-style-type: none"> <li>• <b>egress-time</b>—Measures maximum source-to-destination time per probe.</li> <li>• <b>ingress-time</b>—Measures maximum destination-to-source time per probe.</li> <li>• <b>jitter-egress</b>—Measures maximum source-to-destination jitter per test.</li> <li>• <b>jitter-ingress</b>—Measures maximum destination-to- source jitter per test.</li> <li>• <b>jitter-rtt</b>—Measures maximum jitter per test, from 0 through 60,000,000 microseconds.</li> <li>• <b>rtt</b>—Measures maximum round-trip time per probe, in microseconds.</li> <li>• <b>std-dev-egress</b>—Measures maximum source-to-destination standard deviation per test.</li> <li>• <b>std-dev-ingress</b>—Measures maximum destination-to-source standard deviation per test.</li> <li>• <b>std-dev-rtt</b>—Measures maximum standard deviation per test, in microseconds.</li> <li>• <b>successive-loss</b>—Measures successive probe loss count, indicating probe failure.</li> <li>• <b>total-loss</b>—Measures total probe loss count indicating test failure, from 0 through 15.</li> </ul> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring RPM Probes</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## traps

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>traps traps;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | [edit services rpm <b>probe</b> owner <b>test</b> test-name]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | Set the trap bit to generate traps for probes. Traps are sent if the configured threshold is met or exceeded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>             | <p><b>traps</b>—Specify one or more traps. The following options are supported:</p> <ul style="list-style-type: none"><li>• <b>egress-jitter-exceeded</b>—Generates traps when the jitter in egress time threshold is met or exceeded.</li><li>• <b>egress-std-dev-exceeded</b>—Generates traps when the egress time standard deviation threshold is met or exceeded.</li><li>• <b>egress-time-exceeded</b>—Generates traps when the maximum egress time threshold is met or exceeded.</li><li>• <b>ingress-jitter-exceeded</b>—Generates traps when the jitter in ingress time threshold is met or exceeded.</li><li>• <b>ingress-std-dev-exceeded</b>—Generates traps when the ingress time standard deviation threshold is met or exceeded.</li><li>• <b>ingress-time-exceeded</b>—Generates traps when the maximum ingress time threshold is met or exceeded.</li><li>• <b>jitter-exceeded</b>—Generates traps when the jitter in round-trip time threshold is met or exceeded.</li><li>• <b>probe-failure</b>—Generates traps for successive probe loss thresholds crossed.</li><li>• <b>rtt-exceeded</b>—Generates traps when the maximum round-trip time threshold is met or exceeded.</li><li>• <b>std-dev-exceeded</b>—Generates traps when the round-trip time standard deviation threshold is met or exceeded.</li><li>• <b>test-completion</b>—Generates traps when a test is completed.</li><li>• <b>test-failure</b>—Generates traps when the total probe loss threshold is met or exceeded.</li></ul> |



**NOTE:** For RPM traps to be generated, you must configure the remote-operations SNMP trap category by including the **categories** statement at the [edit snmp trap-group trap-group-name hierarchy level.

---



**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring RPM Probes*
- [categories on page 4089](#)

## udp

**Syntax**

```
udp {
    destination-interface interface-name;
    port port;
}
```

**Hierarchy Level** [edit services rpm [probe-server](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.3 for EX Series switches.  
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.

**Description** Specify the port information for the UDP server.  
  
The remaining statements are explained separately.



**NOTE:** The `destination-interface` statement is not supported on PTX Series routers.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring RPM Receiver Servers*

## Configuration Statements: Uplink Failure Detection

- [action \(Uplink Failure Detection\) on page 4258](#)
- [group \(Uplink Failure Detection\) on page 4258](#)
- [link-to-disable on page 4259](#)
- [link-to-monitor on page 4259](#)
- [traceoptions \(Uplink Failure Detection\) on page 4260](#)
- [uplink-failure-detection on page 4261](#)

## action (Uplink Failure Detection)

---

|                                 |                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>action {<br/>  log;<br/>}</pre>                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">uplink-failure-detection</a> ]                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.                                                                              |
| <b>Description</b>              | Define an action on uplink-failure-detection group state change.                                                                                   |
| <b>Options</b>                  | <b>log</b> —Generate a system log message.                                                                                                         |
| <b>Required Privilege Level</b> | <b>admin</b> —To view this statement in the configuration.<br><b>admin-control</b> —To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Interfaces for Uplink Failure Detection (CLI Procedure)</a> on page 4064</li></ul> |

## group (Uplink Failure Detection)

---

|                                 |                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>group <i>group-name</i> {<br/>  <a href="#">link-to-monitor</a> <i>interface-name</i>;<br/>  <a href="#">link-to-disable</a> <i>interface-name</i>;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols uplink-failure-detection]                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for EX Series switches.                                                                                                |
| <b>Description</b>              | Configure a group of uplink and downlink interfaces for uplink failure detection.                                                                                    |
| <b>Options</b>                  | <b><i>group-name</i></b> —Name of the uplink-failure-detection group.<br><br>The remaining statements are explained separately.                                      |
| <b>Required Privilege Level</b> | <b>admin</b> —To view this statement in the configuration.<br><b>admin-control</b> —To add this statement to the configuration.                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Interfaces for Uplink Failure Detection (CLI Procedure)</a> on page 4064</li></ul>                   |

## link-to-disable

---

|                                 |                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>link-to-disable <i>interface-name</i>;</code>                                                                                                                |
| <b>Hierarchy Level</b>          | <code>[edit protocols uplink-failure-detection group <i>group-name</i>]</code>                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for EX Series switches.                                                                                              |
| <b>Description</b>              | Configure the downlink interfaces to be disabled when the switch detects an uplink failure. The switch can monitor a maximum of 48 downlink interfaces in a group. |
| <b>Options</b>                  | <i>interface-name</i> —Name of the downlink interface or interface range in the group. The interface can be a physical interface or a logical interface.           |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 4064</a></li> </ul>               |

## link-to-monitor

---

|                                 |                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>link-to-monitor <i>interface-name</i>;</code>                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit protocols uplink-failure-detection group <i>group-name</i>]</code>                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for EX Series switches.                                                                                                                                                           |
| <b>Description</b>              | Configure the uplink interfaces to be monitored for uplink failure detection. The switch can monitor a maximum of 48 uplink interfaces in a group.<br><br>An interface can be configured as link-to-monitor in multiple groups. |
| <b>Options</b>                  | <i>interface-name</i> —Name of the uplink interface or interface range in the group. The interface can be a physical interface or a logical interface.                                                                          |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 4064</a></li> </ul>                                                                            |

## traceoptions (Uplink Failure Detection)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable  <br/>    no-world-readable&gt;;<br/>    flag <i>flag</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>     | [edit protocols <a href="#">uplink-failure-detection</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | Define tracing operations for uplink failure detection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>             | The <b>traceoptions</b> feature is disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>             | <p><b>file <i>filename</i></b> —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks.</p> <p><b>files <i>number</i></b> —(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached (<b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b> —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—Trace everything.</li><li>• <b>dcd</b>—Trace ufdi interaction with dcd.</li><li>• <b>groups</b>—Trace uplink-failure-detection group handling.</li><li>• <b>interface</b>—Trace interface notification handlers of ufdi.</li><li>• <b>parse</b>—Trace configuration parsing.</li></ul> <p><b>no-stamp</b>—(Optional) Do not place a timestamp on any trace file.</p> <p><b>no-world-readable</b>—(Optional) Restricted file access to the user who created the file.</p> <p><b>size <i>size</i></b> —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>files</b> option.</p> <p><b>Syntax:</b> <b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify gigabytes</p> |

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation** • [Configuring Interfaces for Uplink Failure Detection \(CLI Procedure\) on page 4064](#)

## uplink-failure-detection

**Syntax**

```
uplink-failure-detection {
  action {
    log;
  }
  group group-name {
    link-to-monitor interface-name;
    link-to-disable interface-name;
  }
  traceoptions {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable |
    no-world-readable>;
    flag flag);
  }
}
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 11.1 for EX Series switches.

**Description** Configure uplink and downlink interfaces in a group to monitor uplink failures and to propagate uplink failures to the downlink interfaces.

The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation** • [Configuring Interfaces for Uplink Failure Detection \(CLI Procedure\) on page 4064](#)



## CHAPTER 68

# Administration

- Routine Monitoring on page 4263
- Operational Commands: General on page 4269
- Operational Commands: RPM on page 4292
- Operational Commands: SNMP on page 4302
- Operational Commands: Analyzers and Port Mirroring on page 4338
- Commands for Network Analytics on page 4342
- Operational Commands: sFlow on page 4360
- Operational Commands: Ethernet OAM Connectivity Fault Management on page 4366
- Operational Commands: Ethernet OAM Link Fault Management on page 4404
- Operational Commands: Uplink Failure Detection on page 4409

### Routine Monitoring

---

- Monitoring Hosts Using the J-Web Ping Host Tool on page 4263
- Monitoring Network Traffic Using Traceroute on page 4265
- Verifying Input and Output for Port Mirroring Analyzers on EX Series Switches on page 4267
- Viewing Real-Time Performance Monitoring Information on page 4268
- Verifying That Uplink Failure Detection Is Working Correctly on page 4268

### Monitoring Hosts Using the J-Web Ping Host Tool

#### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the J-Web ping host tool to verify that the host can be reached over the network. The output is useful for diagnosing host and network connectivity problems. The switch sends a series of ICMP echo (ping) requests to a specified host and receives ICMP echo responses.

**Action** To use the J-Web ping host tool:

1. Select **Troubleshoot>Ping Host**.
2. Next to Advanced options, click the expand icon.
3. Enter information into the Ping Host page, as described in [Table 433 on page 4264](#).

The Remote Host field is the only required field.

4. Click **Start**.

The results of the ping operation are displayed in the main pane. If no options are specified, each ping response is in the following format:

*bytes bytes from ip-address: icmp\_seq=number ttl=number time=time*

5. To stop the ping operation before it is complete, click **OK**.

**Meaning** [Table 433 on page 4264](#) lists the fields.

**Table 433: J-Web Ping Host Field Summary**

| Field                   | Function                                                                                                                                                                    | Your Action                                                                                                                                                                                                              |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Remote Host             | Identifies the host to ping.                                                                                                                                                | Type the hostname or IP address of the host to ping.                                                                                                                                                                     |
| <b>Advanced Options</b> |                                                                                                                                                                             |                                                                                                                                                                                                                          |
| Don't Resolve Addresses | Determines whether to display hostnames of the hops along the path.                                                                                                         | <ul style="list-style-type: none"> <li>• To suppress the display of the hop hostnames, select the check box.</li> <li>• To display the hop hostnames, clear the check box.</li> </ul>                                    |
| Interface               | Specifies the interface on which the ping requests are sent.                                                                                                                | Select the interface on which ping requests are sent from the list. If you select <b>any</b> , the ping requests are sent on all interfaces.                                                                             |
| Count                   | Specifies the number of ping requests to send.                                                                                                                              | Select the number of ping requests to send from the list.                                                                                                                                                                |
| Don't Fragment          | Specifies the Don't Fragment (DF) bit in the IP header of the ping request packet.                                                                                          | <ul style="list-style-type: none"> <li>• To set the DF bit, select the check box.</li> <li>• To clear the DF bit, clear the check box.</li> </ul>                                                                        |
| Record Route            | Sets the record route option in the IP header of the ping request packet. The path of the ping request packet is recorded within the packet and displayed in the main pane. | <ul style="list-style-type: none"> <li>• To record and display the path of the packet, select the check box.</li> <li>• To suppress the recording and display of the path of the packet, clear the check box.</li> </ul> |
| Type-of-Service         | Specifies the type-of-service (TOS) value in the IP header of the ping request packet.                                                                                      | Select the decimal value of the TOS field from the list.                                                                                                                                                                 |
| Routing Instance        | Name of the routing instance for the ping attempt.                                                                                                                          | Select the routing instance name from the list.                                                                                                                                                                          |
| Interval                | Specifies the interval, in seconds, between transmissions of individual ping requests.                                                                                      | Select the interval from the list.                                                                                                                                                                                       |
| Packet Size             | Specifies the size of the ping request packet.                                                                                                                              | Type the size, in bytes, of the packet. The size can be from 0 through 65468. The switch adds 8 bytes of ICMP header to the size.                                                                                        |



Table 433: J-Web Ping Host Field Summary (*continued*)

| Field          | Function                                                                                                                                                                                                                                                                            | Your Action                                                                                                                                                                                                                                                |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Source Address | Specifies the source address of the ping request packet.                                                                                                                                                                                                                            | Type the source IP address.                                                                                                                                                                                                                                |
| Time-to-Live   | Specifies the time-to-live (TTL) hop count for the ping request packet.                                                                                                                                                                                                             | Select the TTL value from the list.                                                                                                                                                                                                                        |
| Bypass Routing | <p>Determines whether ping requests are routed by means of the routing table.</p> <p>If the routing table is not used, ping requests are sent only to hosts on the interface specified in the Interface box. If the host is not on that interface, ping responses are not sent.</p> | <ul style="list-style-type: none"> <li>To bypass the routing table and send the ping requests to hosts on the specified interface only, select the check box.</li> <li>To route the ping requests using the routing table, clear the check box.</li> </ul> |

**Related Documentation** • [Monitoring Interface Status and Traffic on page 2835](#)

## Monitoring Network Traffic Using Traceroute

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the Traceroute page in the J-Web interface to trace a route between the switch and a remote host. You can use a traceroute task to display a list of waypoints between the switch and a specified destination host. The output is useful for diagnosing a point of failure in the path from the switch platform to the destination host and addressing network traffic latency and throughput problems.

### Action

To use the traceroute tool:

1. Select **Troubleshoot > Traceroute**.
2. Next to **Advanced options**, click the expand icon.
3. Enter information into the Traceroute page.  
The **Remote Host** field is the only required field.
4. Click **Start**.
5. To stop the traceroute operation before it is complete, click **OK** while the results of the traceroute operation are being displayed.

### Meaning

The switch generates the list of waypoints by sending a series of ICMP traceroute packets in which the time-to-live (TTL) value in the messages sent to each successive waypoint is incremented by 1. (The TTL value of the first traceroute packet is set to 1.) In this manner, each waypoint along the path to the destination host replies with a Time Exceeded packet from which the source IP address can be obtained.

The results of the traceroute operation are displayed in the main pane. If no options are specified, each line of the traceroute display is in the following format:

**hop-number host (ip-address) [as-number] time1 time2 time3**

The switch sends a total of three traceroute packets to each waypoint along the path and displays the round-trip time for each traceroute operation. If the switch times out before receiving a **Time Exceeded** message, an asterisk (\*) is displayed for that round-trip time.

**Table 434: Traceroute field summary**

| Field                   | Function                                                                                                                                                                                                                                                                              | Your Action                                                                                                                                      |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| Remote Host             | Identifies the destination host of the traceroute.                                                                                                                                                                                                                                    | Type the hostname or IP address of the destination host.                                                                                         |
| Advanced Options        |                                                                                                                                                                                                                                                                                       |                                                                                                                                                  |
| Don't Resolve Addresses | Determines whether hostnames of the hops along the path are displayed, in addition to IP addresses.                                                                                                                                                                                   | To suppress the display of the hop hostnames, select the check box.                                                                              |
| Gateway                 | Specifies the IP address of the gateway to route through.                                                                                                                                                                                                                             | Type the gateway IP address.                                                                                                                     |
| Source Address          | Specifies the source address of the outgoing traceroute packets.                                                                                                                                                                                                                      | Type the source IP address.                                                                                                                      |
| Bypass Routing          | Determines whether traceroute packets are routed by means of the routing table. If the routing table is not used, traceroute packets are sent only to hosts on the interface specified in the Interface box. If the host is not on that interface, traceroute responses are not sent. | To bypass the routing table and send the traceroute packets to hosts on the specified interface only, select the check box.                      |
| Interface               | Specifies the interface on which the traceroute packets are sent.                                                                                                                                                                                                                     | From the list, select the interface on which traceroute packets are sent. If you select any, the traceroute requests are sent on all interfaces. |
| Time-to-live            | Specifies the maximum time-to-live (TTL) hop count for the traceroute request packet.                                                                                                                                                                                                 | From the list, select the TTL.                                                                                                                   |
| Type-of-Service         | Specifies the type-of-service (TOS) value to include in the IP header of the traceroute request packet.                                                                                                                                                                               | From the list, select the decimal value of the TOS field.                                                                                        |
| Resolve AS Numbers      | Determines whether the autonomous system (AS) number of each intermediate hop between the router and the destination host is displayed.                                                                                                                                               | To display the AS numbers, select the check box.                                                                                                 |

**Related Documentation**

- [Connecting and Configuring an EX Series Switch \(CLI Procedure\)](#)
- [Connecting and Configuring an EX Series Switch \(J-Web Procedure\)](#)
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 2619](#)

- [Monitoring Interface Status and Traffic on page 2835](#)

## Verifying Input and Output for Port Mirroring Analyzers on EX Series Switches

**Purpose** Verify that an analyzer has been created on the switch and has the appropriate output interfaces, and appropriate output interface.

**Action** You can verify the port mirror analyzer is configured as expected using the **show analyzer** command.

```
[edit]
user@switch> show analyzer
Analyzer name           : employee-monitor
Output VLAN             : remote-analyzer
Mirror ratio            : 1
Loss priority           : High
Ingress monitored interfaces : ge-0/0/0.0
Ingress monitored interfaces : ge-0/0/1.0
```

You can view all of the port mirror analyzers configured on the switch, including any that are disabled, using the **show ethernet-switching-options** command in configuration mode.

```
user@switch# show ethernet-switching-options
inactive: analyzer employee-web-monitor {
    loss-priority high;
    output {

analyzer employee-monitor {
    loss-priority high;
    input {
        ingress {
            interface ge-0/0/0.0;
            interface ge-0/0/1.0;
        }
    }
    output {
        vlan {
            remote-analyzer;
        }
    }
}
```

**Meaning** This output shows that the employee-monitor analyzer has a ratio of 1 (mirroring every packet, the default), a loss priority of high (set this option to high whenever the analyzer output is to a VLAN), is mirroring the traffic entering **ge-0/0/0** and **ge-0/0/1**, and sending the mirrored traffic to the analyzer called remote-analyzer.

**Related Documentation**

- [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 4047](#)
- [Configuring Port Mirroring to Analyze Traffic \(CLI Procedure\)](#)
- [Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches](#)

- *Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches*
- *Understanding Port Mirroring on EX Series Switches*

## Viewing Real-Time Performance Monitoring Information



**NOTE:** This topic applies only to the J-Web Application package.

Real-time performance monitoring (RPM) on EX Series switches enables you to configure and send probes to a specified target and monitor the analyzed results to determine packet loss, round-trip time, and jitter. The J-Web interface provides a graphical view of RPM information for EX Series switches.

To view the RPM information using the J-Web interface:

1. Select **Troubleshoot > RPM > View RPM**.
2. Select the **Round Trip Time** check box to display the graph with round-trip time included. Clear the check-box to view the graph without the round-trip time.
3. From the **Refresh Time** list, select a refresh time interval for the graph.

### Related Documentation

- [Configuring Real-Time Performance Monitoring \(J-Web Procedure\) on page 4050](#)

## Verifying That Uplink Failure Detection Is Working Correctly

**Purpose** Verify that the switch disables the downlink interface when it detects an uplink failure.

- Action** 1. View the current uplink-failure-detection status:

```
user@switch> show uplink-failure-detection
Group                : group1
Uplink               : ge-0/0/0*
Downlink             : ge-0/0/1*
Failure Action       : Inactive
```



**NOTE:** The asterisk (\*) indicates that the link is up.

2. Disable the uplink interface:

```
[edit]
user@switch# set interface ge-0/0/0 disable
```

3. Save the configuration on the switch.

4. View the current uplink-failure-detection status:

```
user@switch> show uplink-failure-detection
Group                : group1
Uplink               : ge-0/0/0
Downlink             : ge-0/0/1
Failure Action       : Active
```

**Meaning** The output in Step 1 shows that the uplink interface is up, and hence that the downlink interface is also up, and that the status of **Failure Action** is **Inactive**.

The output in Step 4 shows that both the uplink and downlink interfaces are down and that the status of **Failure Action** is changed to **Active**. This output shows that uplink failure detection is working.

- Related Documentation**
- [Configuring Interfaces for Uplink Failure Detection \(CLI Procedure\) on page 4064](#)
  - [Understanding Uplink Failure Detection on page 3977](#)

## Operational Commands: General

- [monitor traffic](#)
- [ping](#)
- [show pfe statistics bridge](#)
- [traceroute](#)

## monitor traffic

---

**Syntax**    monitor traffic  
              <brief | detail | extensive>  
              <absolute-sequence>  
              <count *count*>  
              <interface *interface-name*>  
              <layer2-headers>  
              <matching *matching*>  
              <no-domain-names>  
              <no-promiscuous>  
              <no-resolve>  
              <no-timestamp>  
              <print-ascii>  
              <print-hex>  
              <resolve-timeout>  
              <size *size*>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              Command introduced in Junos OS Release 9.0 for EX Series switches.  
                              Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Display packet headers or packets received and sent from the Routing Engine.



**NOTE:**

- Using the **monitor-traffic** command can degrade router or switch performance.
  - Delays from DNS resolution can be eliminated by using the **no-resolve** option.
- 
- 



**NOTE:** This command is not supported on the QFabric system.

---

**Options**    **none**—(Optional) Display packet headers transmitted through **fxp0**. On a TX Matrix Plus router, display packet headers transmitted through **em0**.

**brief | detail | extensive**—(Optional) Display the specified level of output.

**absolute-sequence**—(Optional) Display absolute TCP sequence numbers.

**count *count***—(Optional) Specify the number of packet headers to display (0 through 1,000,000). The monitor traffic command quits automatically after displaying the number of packets specified.

**interface *interface-name***—(Optional) Specify the interface on which the **monitor traffic** command displays packet data. If no interface is specified, the **monitor traffic** command displays packet data arriving on the lowest-numbered interface.

**layer2-headers**—(Optional) Display the link-level header on each line.

**matching *matching***—(Optional) Display packet headers that match a regular expression. Use matching expressions to define the level of detail with which the **monitor traffic** command filters and displays packet data.

**no-domain-names**—(Optional) Suppress the display of the domain portion of hostnames. With the **no-domain-names** option enabled, the **monitor traffic** command displays only **team** for the hostname **team.company.net**.

**no-promiscuous**—(Optional) Do not put the interface into promiscuous mode.

**no-resolve**—(Optional) Suppress reverse lookup of the IP addresses.

**no-timestamp**—(Optional) Suppress timestamps on displayed packets.

**print-ascii**—(Optional) Display each packet in ASCII format.

**print-hex**—(Optional) Display each packet, except the link-level header, in hexadecimal format.

**resolve-timeout *timeout***—(Optional) Amount of time the router or switch waits for each reverse lookup before timing out. You can set the timeout for 1 through 4,294,967,295 seconds. The default is 4 seconds. To display each packet, use the **print-ascii**, **print-hex**, or **extensive** option.

**size *size***—(Optional) Read but do not display up to the specified number of bytes for each packet. When set to **brief** output, the default packet size is 96 bytes and is adequate for capturing IP, ICMP, UDP, and TCP packet data. When set to **detail** and **extensive** output, the default packet size is 1514. The **monitor traffic** command truncates displayed packets if the matched data exceeds the configured size.

**Additional Information** In the **monitor traffic** command, you can specify an expression to match by using the **matching** option and including the expression in quotation marks:

```
monitor traffic matching "expression"
```

Replace ***expression*** with one or more of the match conditions listed in [Table 435 on page 4272](#).

Table 435: Match Conditions for the monitor traffic Command

| Match Type    | Condition                                             | Description                                                                                                                                                                                                                                                                            |
|---------------|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Entity        | <b>host</b> [ <i>address</i>   <i>hostname</i> ]      | Matches packets that contain the specified address or hostname.<br><br>The protocol match conditions <b>arp</b> , <b>ip</b> , or <b>rarp</b> , or any of the directional match conditions can be prepended to the <b>host</b> match condition.                                         |
|               | <b>net</b> <i>address</i>                             | Matches packets with source or destination addresses containing the specified network address.                                                                                                                                                                                         |
|               | <b>net</b> <i>address mask mask</i>                   | Matches packets containing the specified network address and subnet mask.                                                                                                                                                                                                              |
|               | <b>port</b> ( <i>port-number</i>   <i>port-name</i> ) | Matches packets containing the specified source or destination TCP or UDP port number or port name.<br><br>In place of the numeric port address, you can specify a text synonym, such as <b>bgp</b> (179), <b>dhcp</b> (67), or <b>domain</b> (53) (the port numbers are also listed). |
| Directional   | <b>dst</b>                                            | Matches packets going to the specified destination. This match condition can be prepended to any of the entity type match conditions.                                                                                                                                                  |
|               | <b>src</b>                                            | Matches packets from a specified source. This match condition can be prepended to any of the entity type match conditions.                                                                                                                                                             |
|               | <b>src and dst</b>                                    | Matches packets that contain the specified source and destination addresses. This match condition can be prepended to any of the entity type match conditions.                                                                                                                         |
|               | <b>src or dst</b>                                     | Matches packets containing either of the specified addresses. This match condition can be prepended to any of the entity type match conditions.                                                                                                                                        |
| Packet Length | <b>less</b> <i>value</i>                              | Matches packets shorter than or equal to the specified value, in bytes.                                                                                                                                                                                                                |
|               | <b>greater</b> <i>value</i>                           | Matches packets longer than or equal to the specified value, in bytes.                                                                                                                                                                                                                 |



Table 435: Match Conditions for the monitor traffic Command (*continued*)

| Match Type | Condition                                                | Description                                                                                                                                                                                                                                                                                                                     |
|------------|----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Protocol   | <b>amt</b>                                               | Matches all AMT packets. Use the extensive level of output to decode the inner IGMP packets in addition to the AMT outer packet.                                                                                                                                                                                                |
|            | <b>arp</b>                                               | Matches all ARP packets.                                                                                                                                                                                                                                                                                                        |
|            | <b>ether</b>                                             | Matches all Ethernet packets.                                                                                                                                                                                                                                                                                                   |
|            | <b>ether (broadcast   multicast)</b>                     | Matches broadcast or multicast Ethernet frames. This match condition can be prepended with <b>src</b> and <b>dst</b> .                                                                                                                                                                                                          |
|            | <b>ether protocol (address   (arp   ip   rarp))</b>      | Matches packets with the specified Ethernet address or Ethernet packets of the specified protocol type. The <b>ether protocol</b> arguments <b>arp</b> , <b>ip</b> , and <b>rarp</b> are also independent match conditions, so they must be preceded by a backslash (\) when used in the <b>ether protocol</b> match condition. |
|            | <b>icmp</b>                                              | Matches all ICMP packets.                                                                                                                                                                                                                                                                                                       |
|            | <b>ip</b>                                                | Matches all IP packets.                                                                                                                                                                                                                                                                                                         |
|            | <b>ip (broadcast   multicast)</b>                        | Matches broadcast or multicast IP packets.                                                                                                                                                                                                                                                                                      |
|            | <b>ip protocol (address   (icmp   igmp   tcp   udp))</b> | Matches packets with the specified address or protocol type. The <b>ip protocol</b> arguments <b>icmp</b> , <b>tcp</b> , and <b>udp</b> are also independent match conditions, so they must be preceded by a backslash (\) when used in the <b>ip protocol</b> match condition.                                                 |
|            | <b>isis</b>                                              | Matches all IS-IS routing messages.                                                                                                                                                                                                                                                                                             |
|            | <b>rarp</b>                                              | Matches all RARP packets.                                                                                                                                                                                                                                                                                                       |
|            | <b>tcp</b>                                               | Matches all TCP datagrams.                                                                                                                                                                                                                                                                                                      |
|            | <b>udp</b>                                               | Matches all UDP datagrams.                                                                                                                                                                                                                                                                                                      |

To combine expressions, use the logical operators listed in [Table 436 on page 4273](#).

Table 436: Logical Operators for the monitor traffic Command

| Logical Operator (Highest to Lowest Precedence) | Description                                                                          |
|-------------------------------------------------|--------------------------------------------------------------------------------------|
| <b>!</b>                                        | Logical NOT. If the first condition does not match, the next condition is evaluated. |

Table 436: Logical Operators for the monitor traffic Command (*continued*)

| Logical Operator (Highest to Lowest Precedence) | Description                                                                                                                                         |
|-------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| &&                                              | Logical AND. If the first condition matches, the next condition is evaluated. If the first condition does not match, the next condition is skipped. |
|                                                 | Logical OR. If the first condition matches, the next condition is skipped. If the first condition does not match, the next condition is evaluated.  |
| ()                                              | Group operators to override default precedence order. Parentheses are special characters, each of which must be preceded by a backslash (\).        |

You can use relational operators to compare arithmetic expressions composed of integer constants, binary operators, a length operator, and special packet data accessors. The arithmetic expression matching condition uses the following syntax:

```
monitor traffic matching "ether[0] & 1 != 0"arithmetic_expression relational_operator arithmetic_expression
```

The packet data accessor uses the following syntax:

```
protocol [byte-offset <size>]
```

The optional *size* field represents the number of bytes examined in the packet header. The available values are 1, 2, or 4 bytes. The following sample command captures all multicast traffic:

```
user@host> monitor traffic matching "ether[0] & 1 != 0"
```

To specify match conditions that have a numeric value, use the arithmetic and relational operators listed in [Table 437 on page 4275](#).



**NOTE:** Because the Packet Forwarding Engine removes Layer 2 header information before sending packets to the Routing Engine:

- The **monitor traffic** command cannot apply match conditions to inbound traffic.
- The **monitor traffic interface** command also cannot apply match conditions for Layer 3 and Layer 4 packet data, resulting in the match pipe option (`| match`) for this command for Layer 3 and Layer 4 packets not working either. Therefore, ensure that you specify match conditions as described in this command summary. For more information about match conditions, see [Table 435 on page 4272](#).
- The 802.1Q VLAN tag information included in the Layer 2 header is removed from all inbound traffic packets. Because the **monitor traffic interface ae[x]** command for aggregated Ethernet interfaces (such as ae0) only shows inbound traffic data, the command does not show VLAN tag information in the output.

**Table 437: Arithmetic and Relational Operators for the monitor traffic Command**

| Arithmetic or Relational Operator                         | Description                                                                         |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------|
| <b>Arithmetic Operator</b>                                |                                                                                     |
| +                                                         | Addition operator.                                                                  |
| -                                                         | Subtraction operator.                                                               |
| /                                                         | Division operator.                                                                  |
| &                                                         | Bitwise AND.                                                                        |
| *                                                         | Bitwise exclusive OR.                                                               |
|                                                           | Bitwise inclusive OR.                                                               |
| <b>Relational Operator (Highest to Lowest Precedence)</b> |                                                                                     |
| <=                                                        | If the first expression is less than or equal to the second, the packet matches.    |
| >=                                                        | If the first expression is greater than or equal to the second, the packet matches. |
| <                                                         | If the first expression is less than the second, the packet matches.                |
| >                                                         | If the first expression is greater than the second, the packet matches.             |
| =                                                         | If the compared expressions are equal, the packet matches.                          |
| !=                                                        | If the compared expressions are unequal, the packet matches.                        |

**Required Privilege Level** trace  
maintenance

**List of Sample Output** [monitor traffic count on page 4276](#)  
[monitor traffic detail count on page 4276](#)  
[monitor traffic extensive \(Absolute Sequence\) on page 4276](#)  
[monitor traffic extensive \(Relative Sequence\) on page 4276](#)  
[monitor traffic extensive count on page 4276](#)  
[monitor traffic interface on page 4277](#)  
[monitor traffic matching on page 4277](#)  
[monitor traffic \(TX Matrix Plus Router\) on page 4277](#)  
[monitor traffic \(QFX3500 Switch\) on page 4278](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### monitor traffic count

```
user@host> monitor traffic count 2
listening on fxp0
04:35:49.814125 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529478 win 16798 (DF)
04:35:49.814185
Out my-server.work.net.telnet > my-server.home.net.1295: P
1:38(37) ack 0 win 17680 (DF) [tos 0x10]
```

### monitor traffic detail count

```
user@host> monitor traffic detail count 2
listening on fxp0
04:38:16.265864 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529971 win 17678 (DF) (ttl 121, id 6812)
04:38:16.265926
Out my-server.work.net.telnet.telnet > my-server.home.net.1295: P 1:38(37) ack 0
win 17680 (DF) [tos 0x10] (ttl 6)
```

### monitor traffic extensive (Absolute Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp" absolute-sequence
listening on fxp0
In 207.17.136.193.179 > 192.168.4.227.1024: . 4042780859:4042780859(0)
ack 1845421797 win 16384 <nop,nop,timestamp 4935628 965951> [tos 0xc0] (ttl )
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042780859:4042780912(53)
ack 1845421797 win 16384
<nop,nop,timestamp 4935628 965951>:
BGP [|BGP UPDAT)
In 192.168.4.227.1024 > 207.17.136.193.179:
P 1845421797:1845421852(55) ack 4042780912 win 16384 <nop,nop,timestamp 965951
4935628>: BGP [|BGP UPDAT)
...
```

### monitor traffic extensive (Relative Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp"
listening on fxp0
In 172.24.248.221.1680 > 192.168.4.210.23: . 396159737:396159737(0)
ack 1664980689 win 17574 (DF) (ttl 121, id 50003)
Out 192.168.4.210.23 > 172.24.248.221.1680: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10] (ttl 64, id 5394)
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042775817:4042775874(57)
ack 1845416593 win 16384 <nop,nop,timestamp 4935379 965690>: BGP [|BGP UPDAT)
...
```

### monitor traffic extensive count

```
user@host> monitor traffic extensive count 5 no-domain-names no-resolve
listening on fxp013:18:17.406933
In 192.168.4.206.2723610880 > 172.17.28.8.2049:
40 null (ttl 64, id 38367)13:18:17.407577
In 172.17.28.8.2049 > 192.168.4.206.2723610880:
```

```

reply ok 28 null (ttl 61, id 35495)13:18:17.541140
In 0:e0:1e:42:9c:e0 0:e0:1e:42:9c:e0 9000 60:
0000 0100 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 000013:18:17.591513
In 172.24.248.156.4139 > 192.168.4.210.23:
3556964918:3556964918(0)
ack 295526518 win 17601 (DF)
(ttl 121, id 14)13:18:17.591568
Out 192.168.4.210.23 >
172.24.248.156.4139: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10]
(ttl 64, id 52376)

```

### monitor traffic interface

```

user@host> monitor traffic interface fxp0
listening on fxp0.0
18:17:28.800650 In server.home.net.723 > host1-0.lab.home.net.log
18:17:28.800733 Out host2-0.lab.home.net.login > server.home.net.7
18:17:28.817813 In host30.lab.home.net.syslog > host40.home0
18:17:28.817846 In host30.lab.home.net.syslog > host40.home0
...

```

### monitor traffic matching

```

user@host> monitor traffic matching "net 192.168.1.0/24"
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on fxp0, capture size 96 bytes

Reverse lookup for 192.168.1.255 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use no-resolve to avoid reverse lookups on IP addresses.

21:55:54.003511 In IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003585 Out IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003864 In arp who-has 192.168.1.17 tell 192.168.1.9
...

```

### monitor traffic (TX Matrix Plus Router)

```

user@host> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on em0, capture size 96 bytes
04:11:59.862121 Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog > sv-log-01.englab.juniper.net.syslog:
SYSLOG kernel.info, length: 57
04:11:59.862303
Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog >
sv-log-02.englab.juniper.net.syslog: SYSLOG kernel.info, length: 57
04:11:59.923948
In IP aj-em0.englab.juniper.net.65235 >

```

```

summit-em0.englab.juniper.net.telnet: .
ack 1087492766 win 33304 <nop,nop,timestamp 42366734 993490>
04:11:59.923983 Out IP truncated-ip - 232 bytes missing!
summit-em0.englab.juniper.net.telnet > aj-em0.englab.juniper.net.65235: P
1:241(240) ack 0 win 33304
<nop,nop,timestamp 993590 42366734>
04:12:00.022900
In IP aj-em0.englab.juniper.net.65235 >
summit-em0.englab.juniper.net.telnet: . ack 241 win 33304 <nop,nop,timestamp
42366834 993590>
04:12:00.141204
In IP truncated-ip - 40 bytes missing!
ipg-lnx-shell11.juniper.net.46182 > summit-em0.englab.juniper.net.telnet: P
2950530356:2950530404(48) ack 485494987 win 63712
<nop,nop,timestamp 1308555294 987086>
04:12:00.141345
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 1:6(5)
ack 48 win 33304
<nop,nop,timestamp 993809 1308555294>
04:12:00.141572
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 6 win 63712
<nop,nop,timestamp 1308555294 993809>
04:12:00.141597
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 6:10(4) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.141821
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 10 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.141837 Out IP truncated-ip - 2 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 10:20(10) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.142072
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: . ack 20 win 63712
<nop,nop,timestamp 1308555294 993810>
04:12:00.142089 Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 20:28(8) ack 48 win 33304 <nop,nop,timestamp
993810 1308555294>
04:12:00.142321
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 28 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.142337
Out IP truncated-ip - 1 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 28:37(9) ack 48 win 33304 <nop,nop,timestamp
993810 1308555294>
...

```

### monitor traffic (QFX3500 Switch)

```

user@switch> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.

```

```
Listening on me4, capture size 96 bytes
Reverse lookup for 172.22.16.246 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use <no-resolve> to avoid reverse lookups on IP addresses.
16:35:32.240873 Out IP truncated-ip - 112 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 4200727624:4200727756(132) ack 2889954831 win 65535
16:35:32.240900 Out IP truncated-ip - 176 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 132:328(196) ack 1 win 65535
...
```

## ping

---


**List of Syntax**   [Syntax on page 4280](#)  
                          [Syntax \(QFX Series\) on page 4280](#)

**Syntax**   `ping host`  
              `<bypass-routing>`  
              `<count requests>`  
              `<detail>`  
              `<do-not-fragment>`  
              `<inet | inet6>`  
              `<interface source-interface>`  
              `<interval seconds>`  
              `<logical-system logical-system-name>`  
              `<loose-source value>`  
              `<mac-address mac-address>`  
              `<no-resolve>`  
              `<pattern string>`  
              `<rapid>`  
              `<record-route>`  
              `<routing-instance routing-instance-name>`  
              `<size bytes>`  
              `<source source-address>`  
              `<strict >`  
              `<strict-source value.>`  
              `<tos type-of-service>`  
              `<ttl value>`  
              `<verbose>`  
              `<vpls instance-name>`  
              `<wait seconds>`

**Syntax (QFX Series)**   `ping host`  
                          `<bypass-routing>`  
                          `<count requests>`  
                          `<detail>`  
                          `<do-not-fragment>`  
                          `<inet>`  
                          `<interface source-interface>`  
                          `<interval seconds>`  
                          `<logical-system logical-system-name>`  
                          `<loose-source value>`  
                          `<mac-address mac-address>`  
                          `<no-resolve>`  
                          `<pattern string>`  
                          `<rapid>`  
                          `<record-route>`  
                          `<routing-instance routing-instance-name>`  
                          `<size bytes>`  
                          `<source source-address>`  
                          `<strict>`  
                          `< strict-source value>`  
                          `<tos type-of-service>`  
                          `<ttl value>`  
                          `<verbose>`



<wait *seconds*>

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b> | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>         | <p>Check host reachability and network connectivity. The <b>ping</b> command sends Internet Control Message Protocol (ICMP) ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host. Press Ctrl+c to interrupt a ping command.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>             | <p><b>host</b>—IP address or hostname of the remote system to ping.</p> <p><b>bypass-routing</b>—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.</p> <p><b>count requests</b>—(Optional) Number of ping requests to send. The range of values is 1 through 2,000,000,000. The default value is an unlimited number of requests.</p> <p><b>detail</b>—(Optional) Include in the output the interface on which the ping reply was received.</p> <p><b>do-not-fragment</b>—(Optional) Set the do-not-fragment (DF) flag in the IP header of the ping packets. For IPv6 packets, this option disables fragmentation.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> <b>NOTE:</b> In Junos OS Release 11.1 and later, when issuing the <b>ping</b> command for an IPv6 route with the <b>do-not-fragment</b> option, the maximum ping packet size is calculated by subtracting 48 bytes (40 bytes for the IPV6 header and 8 bytes for the ICMP header) from the MTU. Therefore, if the ping packet size (including the 48-byte header) is greater than the MTU, the ping operation might fail.</p> </div> <p><b>inet</b>—(Optional) Ping Packet Forwarding Engine IPv4 routes.</p> <p><b>inet6</b>—(Optional) Ping Packet Forwarding Engine IPv6 routes.</p> <p><b>interface source-interface</b>—(Optional) Interface to use to send the ping requests.</p> <p><b>interval seconds</b>—(Optional) How often to send ping requests. The range of values, in seconds, is 1 through infinity. The default value is 1.</p> <p><b>logical-system logical-system-name</b>—(Optional) Name of logical system from which to send the ping requests.</p> <p>Alternatively, enter the <b>set cli logical-system logical-system-name</b> command and then run the <b>ping</b> command. To return to the main router or switch, enter the <b>clear cli logical-system</b> command.</p> |

**loose-source value**—(Optional) Intermediate loose source route entry (IPv4). Open a set of values.

**mac-address mac-address**—(Optional) Ping the physical or hardware address of the remote system you are trying to reach.

**no-resolve**—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

**pattern string**—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.

**rapid**—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the **count** option.

**record-route**—(Optional) Record and report the packet's path (IPv4).

**routing-instance routing-instance-name**—(Optional) Name of the routing instance for the ping attempt.

**size bytes**—(Optional) Size of ping request packets. The range of values, in bytes, is 0 through 65,468. The default value is 56, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

**source source-address**—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

**strict**—(Optional) Use the strict source route option (IPv4).

**strict-source value**—(Optional) Intermediate strict source route entry (IPv4). Open a set of values.

**tos type-of-service**—(Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is 0 through 255.

If the device configuration includes the **dscp-code-point value** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level, the configured DSCP value overrides the value specified in this command option. In this case, the ToS field of ICMP echo request packets sent on behalf of this command carries the DSCP value specified in the **dscp-code-point** configuration statement instead of the value you specify in this command option.

**ttl value**—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is 0 through 255.

**verbose**—(Optional) Display detailed output.

**vpls instance-name**—(Optional) Ping the instance to which this VPLS belongs.

**wait seconds**—(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is 10 seconds. If this option is used without the count option, a default count of 5 packets is used.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | network                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages</i></li> </ul>                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">ping hostname on page 4283</a><br><a href="#">ping hostname rapid on page 4283</a><br><a href="#">ping hostname size count on page 4283</a>                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | <p>When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. These packets are not counted in the received packets count. They are accounted for separately.</p> |

## Sample Output

### ping hostname

```
user@host> ping skye
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

### ping hostname rapid

```
user@host> ping skye rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!
--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

### ping hostname size count

```
user@host> ping skye size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

## show pfe statistics bridge

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show pfe statistics bridge</b><br><b>&lt;fpc slot&gt;</b>                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Display information about the number of packets discarded in the ingress pipeline of the Packet Forwarding Engine, packets discarded because of egress filtering or congestion filtering, number of control packets, and general counters for dropped packets. You can use this information to inform troubleshooting investigations.                                                                          |
| <b>Options</b>                  | <b>none</b> —Display bridge counter statistics for all Flexible PIC Concentrator (FPC) slots.<br><b>fpc slot</b> —(Optional) Display bridge counter statistics for a specific FPC slot.                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring System Process Information on page 850</a></li> <li>• <a href="#">Monitoring Switch Control Traffic on page 843</a></li> </ul>                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <a href="#">show pfe statistics bridge (EX3200 and EX4200 Switches) on page 4285</a><br><a href="#">show pfe statistics bridge (EX8200 Switches and EX8200 Virtual Chassis) on page 4286</a><br><a href="#">show pfe statistics bridge fpc (EX8200 Switches and EX8200 Virtual Chassis) on page 4287</a><br><a href="#">show pfe statistics bridge fpc (EX8200-40XS (40-port SFP+) Line Card) on page 4287</a> |
| <b>Output Fields</b>            | Table 438 on page 4284 lists the output fields for the <b>show pfe statistics bridge</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                      |

**Table 438: show pfe statistics bridge Output Fields**

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Ingress Counters</b> | Information about ingress counters: <ul style="list-style-type: none"> <li>• <b>Received</b>—Number of packets received by the bridge.</li> <li>• <b>VLAN Filtered</b>—Number of packets discarded because of VLAN filtering.</li> <li>• <b>Security Filtered</b>—Number of packets discarded because of security filtering.</li> <li>• <b>Other Discards</b>—Number of packets dropped by the bridge for reasons other than VLAN or security filtering.</li> </ul> |

Table 438: show pfe statistics bridge Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Egress Counters</b>  | <p>Information about egress counters:</p> <ul style="list-style-type: none"> <li>• <b>Unicast</b>—Number of unicast packets transmitted.</li> <li>• <b>Multicast</b>—Number of multicast packets transmitted.</li> <li>• <b>Broadcast</b>—Number of broadcast packets transmitted.</li> <li>• <b>Egress Filtered</b>—Number of egress-filtered packets (regardless of port, priority, or mode).</li> <li>• <b>TailDrop</b>—Number of packets filtered because of egress queue congestion.</li> <li>• <b>Forward Restrict</b>—Number of packets filtered because of egress forward restrictions.</li> <li>• <b>Congestion Filtered</b>—Number of packets filtered because of transmit queue (TxQ) congestion.</li> <li>• <b>Control Packets</b>—Number of control packets (sent to CPU, received from CPU, and sent to analyzer).</li> </ul> |
| <b>Drop Counters</b>    | <p>Information about drop counters:</p> <ul style="list-style-type: none"> <li>• <b>Drop Mode</b>—Count mode of the counter.</li> <li>• <b>Drop Counter</b>—Counter value.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>General Counters</b> | <p>Information about general counters:</p> <ul style="list-style-type: none"> <li>• <b>Drop Mode</b>—Count mode of the counter.</li> <li>• <b>Drop Counter</b>—Counter value.</li> <li>• <b>Source Not Learnt</b>—Number of source addresses that were not learnt because of internal congestion.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>MUX PFE</b>          | <p>Information about multiplexer PFE for oversubscribed cards:</p> <ul style="list-style-type: none"> <li>• <b>Drop Mode</b>—Count mode of the counter.</li> <li>• <b>Drop Count</b>—Counter value.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Sample Output

### show pfe statistics bridge (EX3200 and EX4200 Switches)

```

user@switch> show pfe statistics bridge
Slot 0

PFE:                0          1          2
-----
---- Ingress Counters ----
Received:            0          52          0
VLAN Filtered:       0          0          0
Security Filtered:   0          0          0
Other Discards:      0          0          0
---- Egress Counters ----
Unicast:             0         104         52
Multicast:           0          0          0
Broadcast:           0          0          0
Egress Filtered:     0          0          0
Congestion Filtered: 0          0          0
Control Packets:     5          0          0
---- General Counters ----
Drop Mode:           0          0          0

```

```

Drop Counter:          34217      36080      6367
Source Not Learnt:      0          0          0

```

### show pfe statistics bridge (EX8200 Switches and EX8200 Virtual Chassis)

```
user@switch> show pfe statistics bridge
```

```
Slot 0
```

```
PFE:                0          1
```

```
----- Ingress Counters -----
```

```

Received:              946          48
VLAN Filtered:         0           0
Security Filtered:     0           0
Other Discards:        0           0

```

```
----- Egress Counters -----
```

```

Unicast:               0           0
Multicast:             0           0
Broadcast:             0           0
Egress Filtered:       0           0
TailDrop:              0           0
Forward Restrict:      0           0
Congestion Filtered:   0           0
Control Packets:       4103        896

```

```
----- Drop Counters -----
```

```

Drop Mode:             0           0
Drop Counter:          12528        2

```

```
Slot 1
```

```
PFE:                0          1
```

```
----- Ingress Counters -----
```

```

Received:              0           0
VLAN Filtered:         0           0
Security Filtered:     0           0
Other Discards:        0           0

```

```
----- Egress Counters -----
```

```

Unicast:               0           0
Multicast:             0           0
Broadcast:             0           0
Egress Filtered:       0           0
TailDrop:              0           0
Forward Restrict:      0           0
Congestion Filtered:   0           0
Control Packets:       0           0

```

```
----- Drop Counters -----
```

```

Drop Mode:             0           0
Drop Counter:          0           0

```

```
Slot 2
```

```
PFE:                0          1
```

```
----- Ingress Counters -----
```

```

Received:              0           0
VLAN Filtered:         0           0
Security Filtered:     0           0
Other Filtered:        0           0

```

```
----- Egress Counters -----
```

```

Unicast:               0           0
Multicast:             0           0

```

```

Broadcast:                0          0
Egress Filtered:          0          0
TailDrop:                 0          0
Forward Restrict:         0          0
Congestion Filtered:      0          0
Control Packets:          0          0
---- Drop Counters ----
Drop Mode:                0          0
Drop Counter:             0          0

```

#### show pfe statistics bridge fpc (EX8200 Switches and EX8200 Virtual Chassis)

```

user@switch> show pfe statistics bridge fpc 2
Slot 2

```

```

PFE:                0          1
-----
---- Ingress Counters ----
Received:            0          0
VLAN Filtered:       0          0
Security Filtered:   0          0
Other Discards:      0          0
---- Egress Counters ----
Unicast:             0          0
Multicast:           0          0
Broadcast:           0          0
Egress Filtered:     0          0
TailDrop:            0          0
Forward Restrict:    0          0
Congestion Filtered: 0          0
Control Packets:     0          0
---- Drop Counters ----
Drop Mode:           0          0
Drop Counter:        0          0

```

#### show pfe statistics bridge fpc (EX8200-40XS (40-port SFP+) Line Card)

```

user@switch> show pfe statistics bridge fpc 8
Slot 8

```

```

PFE:                0          1          2          3
-----
---- Ingress Counters ----
Received:            0          3          0          0
VLAN Filtered:       0          0          0          0
Security Filtered:   0          0          0          0
Other Discards:      0          1          0          0
---- Egress Counters ----
Unicast:             0          0          0          0
Multicast:           0          0          0          0
Broadcast:           0          0          0          0
Egress Filtered:     0          0          0          0
TailDrop:            0          0          0          0
Forward Restrict:    0          0          0          0
Congestion Filtered: 0          2          0          0
Control Packets:     4          0          0          0
---- Drop Counters ----
Drop Mode:           0          0          0          0
Drop Counter:        0          1          0          0

MUX PFE:            4          5

```

-----

|             |   |   |
|-------------|---|---|
| Drop Mode:  | 0 | 0 |
| Drop Count: | 0 | 0 |



## traceroute

**List of Syntax**    [Syntax on page 4289](#)  
                           [Syntax \(QFX Series\) on page 4289](#)

**Syntax**    `traceroute host`  
                   `<as-number-lookup>`  
                   `<bypass-routing>`  
                   `<clns>`  
                   `<gateway address>`  
                   `<inet | inet6>`  
                   `<interface interface-name>`  
                   `<logical system logical-system-name>`  
                   `<monitor host>`  
                   `<mpls (ldp FEC address | rsvp label-switched-path-name)>`  
                   `<no-resolve>`  
                   `<propagate-ttl>`  
                   `<routing-instance routing-instance-name>`  
                   `<source source-address>`  
                   `<tos value>`  
                   `<ttl value>`  
                   `<wait seconds>`

**Syntax (QFX Series)**    `traceroute host`  
                           `<as-number-lookup>`  
                           `<bypass-routing>`  
                           `<gateway address>`  
                           `<inet>`  
                           `<interface interface-name>`  
                           `<monitor host>`  
                           `<no-resolve>`  
                           `<routing-instance routing-instance-name>`  
                           `<source source-address>`  
                           `<tos value>`  
                           `<ttl value>`  
                           `<wait seconds>`

**Release Information**    Command introduced before Junos OS Release 7.4.  
                           Command introduced in Junos OS Release 9.0 for EX Series switches.  
                           **mpls** option introduced in Junos OS Release 9.2.  
                           Command introduced in Junos OS Release 11.1 for the QFX Series.  
                           **propagate-ttl** option introduced in Junos OS Release 12.1.

**Description**    Display the route that packets take to a specified network host. Use **traceroute** as a debugging tool to locate points of failure in a network.

**Options**    **host**—IP address or name of remote host.

**as-number-lookup**—(Optional) Display the autonomous system (AS) number of each intermediate hop on the path from the host to the destination.

**bypass-routing**—(Optional) Bypass the normal routing tables and send requests directly to a system on an attached network. If the system is not on a directly attached

network, an error is returned. Use this option to display a route to a local system through an interface that has no route through it.

**clns**—(Optional) Trace the route belonging to the Connectionless Network Service (CLNS).

**gateway address**—(Optional) Address of a router or switch through which the route transits.

**inet | inet6**—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.

**interface interface-name**—(Optional) Name of the interface over which to send packets.

**logical-system logical-system-name**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**monitor host**—(Optional) Display real-time monitoring information for the specified host.

**mpls (ldp FEC address | rsvp label-switched-path name)**—(Optional) See *traceroute mpls ldp* and *traceroute mpls rsvp*.

**no-resolve**—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

**propagate-ttl**—(Optional) On the PE routing device, use this option to view locally generated Routing Engine transit traffic. This is applicable for MPLS L3VPN traffic only.

Use for troubleshooting, when you want to view hop-by-hop information from the local provider router to the remote provider router, when TTL decrementing is disabled on the core network using the **no-propagate-ttl** configuration statement.



**NOTE:** Using **propagate-ttl** with **traceroute** on the CE router does not show hop-by-hop information.

---

**routing-instance routing-instance-name**—(Optional) Name of the routing instance for the traceroute attempt.

**source source-address**—(Optional) Source address of the outgoing traceroute packets.

**tos value**—(Optional) Value to include in the IP type-of-service (ToS) field. The range of values is 0 through 255.

**ttl value**—(Optional) Maximum time-to-live value to include in the traceroute request. The range of values is 0 through 128.

**wait seconds**—(Optional) Maximum time to wait for a response to the traceroute request.

**Required Privilege Level**    network

**Related Documentation**

- *traceroute monitor*

**List of Sample Output**

- [traceroute on page 4291](#)
- [traceroute as-number-lookup host on page 4291](#)
- [traceroute no-resolve on page 4291](#)
- [traceroute propagate-ttl on page 4292](#)
- [traceroute \(Between CE Routers, Layer 3 VPN\) on page 4292](#)
- [traceroute \(Through an MPLS LSP\) on page 4292](#)

**Output Fields** [Table 439 on page 4291](#) describes the output fields for the **traceroute** command. Output fields are listed in the approximate order in which they appear.

**Table 439: traceroute Output Fields**

| Field Name             | Field Description                                             |
|------------------------|---------------------------------------------------------------|
| <b>traceroute to</b>   | IP address of the receiver.                                   |
| <b>hops max</b>        | Maximum number of hops allowed.                               |
| <b>byte packets</b>    | Size of packets being sent.                                   |
| <b>number-of-hops</b>  | Number of hops from the source to the named router or switch. |
| <b>router-name</b>     | Name of the router or switch for this hop.                    |
| <b>address</b>         | Address of the router or switch for this hop.                 |
| <b>Round trip time</b> | Average round-trip time, in milliseconds (ms).                |

## Sample Output

### traceroute

```
user@host> traceroute santacruz
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1 blue23 (10.168.1.254)  2.370 ms  2.853 ms  0.367 ms
 2 red14 (10.168.255.250) 0.778 ms  2.937 ms  0.446 ms
 3 yellow (10.156.169.254) 7.737 ms  89.905 ms  0.834 ms
```

### traceroute as-number-lookup host

```
user@host> traceroute as-number-lookup 10.100.1.1
traceroute to 10.100.1.1 (10.100.1.1), 30 hops max, 40 byte packets
 1 10.39.1.1 (10.39.1.1) 0.779 ms  0.728 ms  0.562 ms
 2 10.39.1.6 (10.39.1.6) [AS 32] 0.657 ms  0.611 ms  0.617 ms
 3 10.100.1.1 (10.100.1.1) [AS 10, 40, 50] 0.880 ms  0.808 ms  0.774 ms
```

### traceroute no-resolve

```
user@host> traceroute santacruz no-resolve
```

```
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1  10.168.1.254  0.458 ms  0.370 ms  0.365 ms
 2  10.168.255.250  0.474 ms  0.450 ms  0.444 ms
 3  10.156.169.254  0.931 ms  0.876 ms  0.862 ms
```

### traceroute propagate-ttl

```
user@host> traceroute propagate-ttl 100.200.2.2 routing-instance VPN-A
traceroute to 100.200.2.2 (100.200.2.2) from 1.1.0.2, 30 hops max, 40 byte packets

 1  1.2.0.2 (1.2.0.2)  2.456 ms  1.753 ms  1.672 ms
    MPLS Label=299776 CoS=0 TTL=1 S=0
    MPLS Label=299792 CoS=0 TTL=1 S=1
 2  1.3.0.2 (1.3.0.2)  1.213 ms  1.225 ms  1.166 ms
    MPLS Label=299792 CoS=0 TTL=1 S=1
 3  100.200.2.2 (100.200.2.2)  1.422 ms  1.521 ms  1.443 ms
```

### traceroute (Between CE Routers, Layer 3 VPN)

```
user@host> traceroute vpn09
traceroute to vpn09.skybank.net (10.255.14.179), 30 hops max, 40
byte packets
 1  10.39.10.21 (10.39.10.21)  0.598 ms  0.500 ms  0.461 ms
 2  10.39.1.13 (10.39.1.13)  0.796 ms  0.775 ms  0.806 ms
    MPLS Label=100006 CoS=0 TTL=1 S=1
 3  vpn09.skybank.net (10.255.14.179)  0.783 ms  0.716 ms  0.686
```

### traceroute (Through an MPLS LSP)

```
user@host> traceroute mpls1
traceroute to 10.168.1.224 (10.168.1.224), 30 hops max, 40 byte packets
 1  mpls1-sr0.company.net (10.168.200.101)  0.555 ms  0.393 ms  0.367 ms
    MPLS Label=1024 CoS=0 TTL=1
 2  mpls5-lo0.company.net (10.168.1.224)  0.420 ms  0.394 ms  0.401 ms
```

---

## Operational Commands: RPM

- `show services rpm active-servers`
- `show services rpm history-results`
- `show services rpm probe-results`

## show services rpm active-servers

|                                 |                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show services rpm active-servers                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. |
| <b>Description</b>              | Display the protocols and corresponding ports for which a router or switch is configured as a real-time performance monitoring (RPM) server.                                                                  |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <a href="#">show services rpm active-servers on page 4293</a>                                                                                                                                                 |
| <b>Output Fields</b>            | <a href="#">Table 440 on page 4293</a> lists the output fields for the <b>show services rpm active-servers</b> command. Output fields are listed in the approximate order in which they appear.               |

**Table 440: show services rpm active-servers Output Fields**

| Field Name                        | Field Description                                                                                                                                   |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Protocol</b>                   | Protocol configured on the receiving probe server. The protocol can be the User Datagram Protocol (UDP) or the Transmission Control Protocol (TCP). |
| <b>Port</b>                       | Port configured on the receiving probe server.                                                                                                      |
| <b>Destination interface name</b> | Output interface name for the probes.                                                                                                               |

## Sample Output

### show services rpm active-servers

```
user@host> show services rpm active-servers
  Protocol: TCP, Port: 50000, Destination interface name: lt-0/0/0.0
  Protocol: UDP, Port: 50001, Destination interface name: lt-0/0/0.0
```

## show services rpm history-results

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show services rpm history-results &lt;brief   detail&gt; &lt;owner <i>owner</i>&gt; &lt;since <i>time</i>&gt; &lt;test <i>name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.</p>                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Display standard information about the results of the last 50 probes for each real-time performance monitoring (RPM) instance.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>none</b>—Display the results of the last 50 probes for all RPM instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>owner <i>owner</i></b>—(Optional) Display information for the specified probe owner.</p> <p><b>since <i>time</i></b>—(Optional) Display information from the specified time. Specify time as <i>yyyy-mm-dd.hh:mm:ss</i>.</p> <p><b>test <i>name</i></b>—(Optional) Display information for the specified test.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>    | <p><a href="#">show services rpm history-results on page 4295</a></p> <p><a href="#">show services rpm history-results detail on page 4295</a></p>                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | Table 441 on page 4294 lists the output fields for the <b>show services rpm history-results</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                            |

**Table 441: show services rpm history-results Output Fields**

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                          | Level of Output |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Owner</b>           | Probe owner.                                                                                                                                                                                                                                                                                                               | All levels      |
| <b>Test</b>            | Name of a test for a probe instance.                                                                                                                                                                                                                                                                                       | All levels      |
| <b>Probe received</b>  | Timestamp when the probe result was determined.                                                                                                                                                                                                                                                                            | All levels      |
| <b>Round trip time</b> | Average ping round-trip time (RTT), in microseconds.                                                                                                                                                                                                                                                                       | All levels      |
| <b>Probe results</b>   | <p>Result of a particular probe performed by a remote host. The following information is contained in the results:</p> <ul style="list-style-type: none"> <li><b>Response received</b>—Timestamp when the probe result was determined.</li> <li><b>Rtt</b>—Average ping round-trip time (RTT), in microseconds.</li> </ul> | <b>detail</b>   |

Table 441: show services rpm history-results Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Results over current test</b> | Displays the results for the current test by probe at the time each probe was completed, as well as the status of the current test at the time the probe was completed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail</b>   |
| <b>Probes sent</b>               | Number of probes sent with the current test.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>   |
| <b>Probes received</b>           | Number of probe responses received within the current test.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail</b>   |
| <b>Loss percentage</b>           | Percentage of lost probes for the current test.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |
| <b>Measurement</b>               | <p>Increment of measurement. Possible values are round-trip time delay and, for the probe type icmp-pin-timestamp, the egress and ingress delay:</p> <ul style="list-style-type: none"> <li>• <b>Minimum</b>—Minimum RTT, ingress delay, or egress delay measured over the course of the current test.</li> <li>• <b>Maximum</b>—Maximum RTT, ingress delay, or egress delay measured over the course of the current test.</li> <li>• <b>Average</b>—Average RTT, ingress delay, or egress delay measured over the course of the current test.</li> <li>• <b>Jitter</b>—Difference, in microseconds, between the maximum and minimum RTT measured over the course of the current test.</li> <li>• <b>Stddev</b>—Standard deviation of the round-trip time, in microseconds, measured over the course of the current test.</li> </ul> | <b>detail</b>   |

## Sample Output

### show services rpm history-results

```

user@host> show services rpm history-results
      Owner, Test                Probe received                Round trip time
p1, t1                Wed Aug 12 01:02:35 2009                315 usec
p1, t1                Wed Aug 12 01:02:36 2009                266 usec
p1, t1                Wed Aug 12 01:02:37 2009                314 usec
p1, t1                Wed Aug 12 01:02:38 2009                388 usec
p1, t1                Wed Aug 12 01:02:39 2009                316 usec
p1, t1                Wed Aug 12 01:02:40 2009                271 usec
p1, t1                Wed Aug 12 01:02:41 2009                314 usec
p1, t1                Wed Aug 12 01:02:42 2009                1180 usec

```

### show services rpm history-results detail

```

user@host> show services rpm history-results detail
Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
Probe results:
  Response received, Wed Aug 12 01:02:35 2009,
  Client and server hardware timestamps
  Rtt: 315 usec
Results over current test:
  Probes sent: 1, Probes received: 1, Loss percentage: 0
Measurement: Round trip time
  Samples: 1, Minimum: 315 usec, Maximum: 315 usec, Average: 315 usec,
  Peak to peak: 0 usec, Stddev: 0 usec, Sum: 315 usec

```

Owner: p1, Test: t1, Probe type: icmp-ping-timestamp  
Probe results:  
Response received, Wed Aug 12 01:02:36 2009,  
Client and server hardware timestamps  
Rtt: 266 usec, Round trip jitter: -50 usec,  
Round trip interarrival jitter: 3 usec  
Results over current test:  
Probes sent: 2, Probes received: 2, Loss percentage: 0  
Measurement: Round trip time  
Samples: 2, Minimum: 266 usec, Maximum: 315 usec, Average: 291 usec,  
Peak to peak: 49 usec, Stddev: 24 usec, Sum: 581 usec  
Measurement: Negative round trip jitter  
Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,  
Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec

Owner: p1, Test: t1, Probe type: icmp-ping-timestamp  
Probe results:  
Response received, Wed Aug 12 01:02:37 2009,  
Client and server hardware timestamps  
Rtt: 314 usec, Round trip jitter: 49 usec,  
Round trip interarrival jitter: 6 usec  
Results over current test:  
Probes sent: 3, Probes received: 3, Loss percentage: 0  
Measurement: Round trip time  
Samples: 3, Minimum: 266 usec, Maximum: 315 usec, Average: 298 usec,  
Peak to peak: 49 usec, Stddev: 23 usec, Sum: 895 usec  
Measurement: Positive round trip jitter  
Samples: 1, Minimum: 49 usec, Maximum: 49 usec, Average: 49 usec,  
Peak to peak: 0 usec, Stddev: 0 usec, Sum: 49 usec  
Measurement: Negative round trip jitter  
Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,  
Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec

Owner: p1, Test: t1, Probe type: icmp-ping-timestamp  
Probe results:  
Response received, Wed Aug 12 01:02:38 2009,  
Client and server hardware timestamps  
Rtt: 388 usec, Round trip jitter: 74 usec,  
Round trip interarrival jitter: 10 usec  
Results over current test:  
Probes sent: 4, Probes received: 4, Loss percentage: 0  
Measurement: Round trip time  
Samples: 4, Minimum: 266 usec, Maximum: 388 usec, Average: 321 usec,  
Peak to peak: 122 usec, Stddev: 44 usec, Sum: 1283 usec  
Measurement: Positive round trip jitter  
Samples: 2, Minimum: 49 usec, Maximum: 74 usec, Average: 62 usec,  
Peak to peak: 25 usec, Stddev: 12 usec, Sum: 123 usec  
Measurement: Negative round trip jitter  
Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,  
Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec



## show services rpm probe-results

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show services rpm probe-results<br><owner <i>owner</i> ><br><test <i>name</i> >                                                                                                                                                                        |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Series Routers.                                   |
| <b>Description</b>              | Display the results of the most recent real-time performance monitoring (RPM) probes.                                                                                                                                                                  |
| <b>Options</b>                  | <b>none</b> —Display all results of the most recent RPM probes.<br><br><b>owner <i>owner</i></b> —(Optional) Display information for the specified probe owner.<br><br><b>test <i>name</i></b> —(Optional) Display information for the specified test. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show services rpm probe-results on page 4300</a><br><a href="#">show services rpm probe-results (BGP Neighbor Discovery) on page 4302</a>                                                                                                  |
| <b>Output Fields</b>            | <a href="#">Table 442 on page 4297</a> lists the output fields for the <b>show services rpm probe-results</b> command. Output fields are listed in the approximate order in which they appear.                                                         |

**Table 442: show services rpm probe-results Output Fields**

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Owner</b>          | Owner name. When you configure the probe owner statement at the <b>[edit services rpm]</b> hierarchy level, this field displays the configured owner name. When you configure BGP neighbor discovery through RPM, the output for this field is <b>Rpm-Bgp-Owner</b> .                                                                                                     |
| <b>Test</b>           | Name of a test representing a collection of probes. When you configure the test test-name statement at the <b>[edit services rpm probe owner]</b> hierarchy level, the field displays the configured test name. When you configure BGP neighbor discovery through RPM, the output for this field is <b>Rpm-BGP-Test-<i>n</i></b> , where <i>n</i> is a cumulative number. |
| <b>Target address</b> | Destination address used for the probes.                                                                                                                                                                                                                                                                                                                                  |
| <b>Source address</b> | Source address used for the probes.                                                                                                                                                                                                                                                                                                                                       |
| <b>Probe type</b>     | Protocol configured on the receiving probe server: <b>http-get</b> , <b>http-metadata-get</b> , <b>icmp-ping</b> , <b>icmp-ping-timestamp</b> , <b>tcp-ping</b> , <b>udp-ping</b> , or <b>udp-ping-timestamp</b> .                                                                                                                                                        |
| <b>Test size</b>      | Number of probes within a test.                                                                                                                                                                                                                                                                                                                                           |

Table 442: show services rpm probe-results Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Routing Instance Name</b>     | <p>(BGP neighbor discovery) Name of the configured (if any) routing instance, logical system name, or both, in which the probe is configured:</p> <ul style="list-style-type: none"> <li>When a routing instance is defined within a logical system, the logical system name is followed by the routing instance name. A slash ( / ) is used to separate the two entities. For example, if the routing instance called <b>R1</b> is configured within the logical system called <b>LS</b>, the name in the output field is <b>LS/R1</b>.</li> <li>When a routing instance is configured but the default logical system is used, the name in the output field is the name of the routing instance.</li> <li>When a logical system is configured but the default routing instance is used, the name in the output field is the name of the logical system followed by <b>default</b>. A slash ( / ) is used to separate the two entities. For example, <b>LS/default</b>.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Probe results</b>             | <p>Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement:</p> <ul style="list-style-type: none"> <li><b>Response received</b>—Timestamp when the probe result was determined.</li> <li><b>Client and server hardware timestamps</b>—If timestamps are configured, an entry appears at this point.</li> <li><b>Rtt</b>—Average ping round-trip time (RTT), in microseconds.</li> <li><b>Egress jitter</b>—Egress jitter, in microseconds.</li> <li><b>Ingress jitter</b>—Ingress jitter, in microseconds.</li> <li><b>Round trip jitter</b>—Round-trip jitter, in microseconds.</li> <li><b>Egress interarrival jitter</b>—Egress interarrival jitter, in microseconds.</li> <li><b>Ingress interarrival jitter</b>—Ingress interarrival jitter, in microseconds.</li> <li><b>Round trip interarrival jitter</b>—Round-trip interarrival jitter, in microseconds.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Results over current test</b> | <p>Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes.</p> <ul style="list-style-type: none"> <li><b>Probes sent</b>—Number of probes sent within the current test.</li> <li><b>Probes received</b>—Number of probe responses received within the current test.</li> <li><b>Loss percentage</b>—Percentage of lost probes for the current test.</li> <li><b>Measurement</b>—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type <b>icmp-ping-timestamp</b>, the egress delay and ingress delay.</li> </ul> <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> <li><b>Samples</b>—Number of probes.</li> <li><b>Minimum</b>—Minimum RTT, ingress delay, or egress delay measured over the course of the current test.</li> <li><b>Maximum</b>—Maximum RTT, ingress delay, or egress delay measured over the course of the current test.</li> <li><b>Average</b>—Average RTT, ingress delay, or egress delay measured over the course of the current test.</li> <li><b>Peak to peak</b>—Peak-to-peak difference, in microseconds.</li> <li><b>Stddev</b>—Standard deviation, in microseconds.</li> <li><b>Sum</b>—Statistical sum.</li> </ul> |

Table 442: show services rpm probe-results Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Results over last test</b> | <p>Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed</p> <ul style="list-style-type: none"> <li>• <b>Probes sent</b>—Number of probes sent for the most recently completed test.</li> <li>• <b>Probes received</b>—Number of probe responses received for the most recently completed test.</li> <li>• <b>Loss percentage</b>—Percentage of lost probes for the most recently completed test.</li> <li>• <b>Test completed</b>—Time the most recent test was completed.</li> <li>• <b>Measurement</b>—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type <b>icmp-ping-timestamp</b>, the egress delay and ingress delay.</li> </ul> <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> <li>• <b>Samples</b>—Number of probes.</li> <li>• <b>Minimum</b>—Minimum RTT, ingress delay, or egress delay measured for the most recently completed test.</li> <li>• <b>Maximum</b>—Maximum RTT, ingress delay, or egress delay measured for the most recently completed test.</li> <li>• <b>Average</b>—Average RTT, ingress delay, or egress delay measured for the most recently completed test.</li> <li>• <b>Peak to peak</b>—Peak-to-peak difference, in microseconds.</li> <li>• <b>Stddev</b>—Standard deviation, in microseconds.</li> <li>• <b>Sum</b>—Statistical sum.</li> </ul> |
| <b>Results over all tests</b> | <p>Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test.</p> <ul style="list-style-type: none"> <li>• <b>Probes sent</b>—Number of probes sent in all tests.</li> <li>• <b>Probes received</b>—Number of probe responses received in all tests.</li> <li>• <b>Loss percentage</b>—Percentage of lost probes in all tests.</li> <li>• <b>Measurement</b>—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe types <b>icmp-ping-timestamp</b> and <b>udp-ping-timestamp</b>, the egress delay and ingress delay.</li> </ul> <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> <li>• <b>Samples</b>—Number of probes.</li> <li>• <b>Minimum</b>—Minimum RTT, ingress delay, or egress delay measured over the course of the current test.</li> <li>• <b>Maximum</b>—Maximum RTT, ingress delay, or egress delay measured over the course of the current test.</li> <li>• <b>Average</b>—Average RTT, ingress delay, or egress delay measured over the course of the current test.</li> <li>• <b>Peak to peak</b>—Peak-to-peak difference, in microseconds.</li> <li>• <b>Stddev</b>—Standard deviation, in microseconds.</li> <li>• <b>Sum</b>—Statistical sum.</li> </ul>                                                                                                                                  |

Table 442: show services rpm probe-results Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Error Stats</b> | <p>Displays error statistics for each probe.</p> <ul style="list-style-type: none"> <li>• <b>Invalid client rcv timestamp</b>—Number of client receive timestamp less than client send timestamp.</li> <li>• <b>Invalid server send timestamp</b>—Number of server send timestamp less than server receive timestamp.</li> <li>• <b>Invalid server processing time</b>—Number of server side spent time greater than RTT.</li> </ul> <p><b>NOTE:</b> Error Stats is displayed in the output only if non-zero statistics exists.</p> |

## Sample Output

### show services rpm probe-results

```

user@host> show services rpm probe-results
Owner: ADSN-J4300.ADSN-J2300.D2, Test: 75300002
Target address: 172.16.54.172, Source address: 10.206.0.1,
Probe type: udp-ping-timestamp, Test size: 10 probes
Probe results:
  Response received, Tue Feb  6 14:53:15 2007,
  Client and server hardware timestamps
  Rtt: 575 usec, Egress jitter: 5 usec, Ingress jitter: 8 usec,
  Round trip jitter: 12 usec, Egress interarrival jitter: 8 usec,
  Ingress interarrival jitter: 7 usec, Round trip interarrival jitter: 7 usec,

  Round trip interarrival jitter: 669 usec
Results over current test:
  Probes sent: 10, Probes received: 10, Loss percentage: 0
  Measurement: Round trip time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive round trip jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative round trip jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
  Measurement: Egress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive Egress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative Egress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
  Measurement: Ingress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive Ingress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative Ingress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Results over last test:
  Probes sent: 10, Probes received: 10, Loss percentage: 0

```

```
Test completed on Tue Feb 6 14:53:16 2007
Measurement: Round trip time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive round trip jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative round trip jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Egress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Ingress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Results over all tests:
Probes sent: 560, Probes received: 560, Loss percentage: 0
Measurement: Round trip time
  Samples: 560, Minimum: 805 usec, Maximum: 3114 usec, Average: 1756 usec,

  Peak to peak: 2309 usec, Stddev: 519 usec, Sum: xxxx usec
Measurement: Positive round trip jitter
  Samples: 257, Minimum: 0 usec, Maximum: 2054 usec, Average: 597 usec,
  Peak to peak: 2054 usec, Stddev: 427 usec, Sum: xxxx usec
Measurement: Negative round trip jitter
  Samples: 302, Minimum: 1 usec, Maximum: 1812 usec, Average: 511 usec,
  Peak to peak: 1811 usec, Stddev: 408 usec, Sum: xxxx usec
Measurement: Egress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Ingress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Error Stats:
  Invalid client rcv timestamp: 3, Invalid server send timestamp: 0
  Invalid server processing time: 0
```

### show services rpm probe-results (BGP Neighbor Discovery)

```
user@host> show services rpm probe-results
Owner: Rpm-Bgp-Owner, Test: Rpm-Bgp-Test-1
Target address: 10.209.152.37, Probe type: icmp-ping, Test size: 5 probes
Routing Instance Name: LS1/RI1
Probe results:
  Response received, Fri Oct 28 05:20:23 2005
  Rtt: 662 usec
Results over current test:
  Probes sent: 5, Probes received: 5, Loss percentage: 0
  Measurement: Round trip time
    Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
    Jitter: 133 usec, Stddev: 53 usec
Results over all tests:
  Probes sent: 5, Probes received: 5, Loss percentage: 0
  Measurement: Round trip time
    Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
    Jitter: 133 usec, Stddev: 53 usec
```

## Operational Commands: SNMP

---

- [clear snmp rmon history](#)
- [clear snmp statistics](#)
- [request snmp spoof-trap](#)
- [show snmp health-monitor](#)
- [show snmp inform-statistics](#)
- [show snmp mib](#)
- [show snmp rmon](#)
- [show snmp rmon history](#)
- [show snmp statistics](#)
- [show snmp v3](#)

---

## clear snmp rmon history

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>clear snmp rmon history</b> < <i>interface-name</i>   all>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Delete the samples of Ethernet statistics collected, but do not delete the RMON history configuration.</p> <p>The <b>clear snmp rmon history</b> command deletes all the samples collected for the interface configured for the history group, but not the configuration of that group. If you want to delete the RMON history group configuration, you must use the <b>delete snmp rmon history</b> configuration-mode command.</p> |
| <b>Options</b>                  | <p><b>interface-name</b>—Delete the samples of Ethernet statistics collected for this interface.</p> <p><b>all</b>—Delete the samples of Ethernet statistics collected for all interfaces that have been configured for RMON monitoring.</p>                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show snmp rmon history on page 4328</a></li></ul>                                                                                                                                                                                                                                                                                                                                   |

## clear snmp statistics

---

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear snmp statistics                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Clear Simple Network Management Protocol (SNMP) statistics.                                                                                                                              |
| <b>Options</b>                  | This command has no options.                                                                                                                                                             |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show snmp statistics on page 4332</a></li></ul>                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">clear snmp statistics on page 4304</a>                                                                                                                                       |
| <b>Output Fields</b>            | See <a href="#">show snmp statistics</a> for an explanation of output fields.                                                                                                            |

## Sample Output

### clear snmp statistics

In the following example, SNMP statistics are displayed before and after the **clear snmp statistics** command is issued:

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 8, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 8, Total set varbinds: 0,
    Get requests: 0, Get nexts: 8, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 2298, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 8, Traps: 2290
```

```
user@host> clear snmp statistics
```

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 0, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
```



```
Total request varbinds: 0, Total set varbinds: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0,  
Silent drops: 0, Proxy drops 0  
Output:  
Packets: 0, Too bigs: 0, No such names: 0,  
Bad values: 0, General errors: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0
```

## request snmp spoof-trap

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>request snmp spoof-trap</b><br><b>&lt;trap&gt; variable-bindings &lt;object&gt; &lt;instance&gt; &lt;value&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.2.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Spoof (mimic) the behavior of a Simple Network Management Protocol (SNMP) trap.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>&lt;trap&gt;</b> —Name of the trap to spoof.<br><br><b>variable-bindings &lt;object&gt; &lt;instance&gt; &lt;value&gt;</b> —(Optional) List of variables and values to include in the trap. Each variable binding is specified as an object name, the object instance, and the value (for example, <b>ifIndex[14] = 14</b> ). Enclose the list of variable bindings in quotation marks ( " ") and use a comma to separate each object name, instance, and value definition (for example, <b>variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"</b> ). Objects included in the trap definition that do not have instances and values specified as part of the command are included in the trap and spoofed with automatically generated instances and values.<br><br><b>&lt;dummy name&gt;</b> —A dummy trap name to display the list of available traps.<br><br><b>Question mark (?)</b> —Question mark? to display possible completions. |
| <b>Required Privilege Level</b> | request                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">request snmp spoof-trap (with Variable Bindings) on page 4306</a><br><a href="#">request snmp spoof-trap (Illegal Trap Name) on page 4306</a><br><a href="#">request snmp spoof-trap (Question Mark ?) on page 4310</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## Sample Output

### request snmp spoof-trap (with Variable Bindings)

```
user@host> request snmpspoof-trap linkUp variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"
Spoof trap request result: trap sent successfully
```

### request snmp spoof-trap (Illegal Trap Name)

```
user@host> request snmp spoof-trap xx
Spoof trap request result: trap not found
```

```
Allowed Traps:
ads1AtucInitFailureTrap
ads1AtucPerfESsThreshTrap
ads1AtucPerfLofsThreshTrap
ads1AtucPerfLolsThreshTrap
ads1AtucPerfLossThreshTrap
ads1AtucPerfLprsThreshTrap
ads1AtucRateChangeTrap
ads1AturPerfESsThreshTrap
```

ads1AturPerfLossThreshTrap  
ads1AturPerfLossThreshTrap  
ads1AturPerfLprsThreshTrap  
ads1AturRateChangeTrap  
apsEventChannelMismatch  
apsEventFEPLF  
apsEventModeMismatch  
apsEventPSBF  
apsEventSwitchover  
authenticationFailure  
bfdSessDown  
bfdSessUp  
bgpBackwardTransition  
bgpEstablished  
coldStart  
dlswTrapCircuitDown  
dlswTrapCircuitUp  
dlswTrapTConnDown  
dlswTrapTConnPartnerReject  
dlswTrapTConnProtViolation  
dlswTrapTConnUp  
dsx1LineStatusChange  
dsx3LineStatusChange  
entConfigChange  
fallingAlarm  
frDLCIStatusChange  
ggsnTrapChanged  
ggsnTrapCleared  
ggsnTrapNew  
gmplsTunnelDown  
ifMauJabberTrap  
ipv6IfStateChange  
isisAreaMismatch  
isisAttemptToExceedMaxSequence  
isisAuthenticationFailure  
isisAuthenticationTypeFailure  
isisCorruptedLSPDetected  
isisDatabaseOverload  
isisIDLenMismatch  
isisLSPTooLargeToPropagate  
isisManualAddressDrops  
isisMaxAreaAddressesMismatch  
isisOriginatingLSPBufferSizeMismatch  
isisOwnLSPPurge  
isisProtocolsSupportedMismatch  
isisRejectedAdjacency  
isisSequenceNumberSkip  
isisVersionSkew  
jnxAccessAuthServerDisabled  
jnxAccessAuthServerEnabled  
jnxAccessAuthServiceDown  
jnxAccessAuthServiceUp  
jnxBfdSessDetectionTimeHigh  
jnxBfdSessTxIntervalHigh  
jnxBgpM2BackwardTransition  
jnxBgpM2Established  
jnxCmCfgChange  
jnxCmRescueChange  
jnxCollFlowOverload  
jnxCollFlowOverloadCleared  
jnxCollFtpSwitchover

jnxCollMemoryAvailable  
jnxCollMemoryUnavailable  
jnxCollUnavailableDest  
jnxCollUnavailableDestCleared  
jnxCollUnsuccessfulTransfer  
jnxDfcHardMemThresholdExceeded  
jnxDfcHardMemUnderThreshold  
jnxDfcHardPpsThresholdExceeded  
jnxDfcHardPpsUnderThreshold  
jnxDfcSoftMemThresholdExceeded  
jnxDfcSoftMemUnderThreshold  
jnxDfcSoftPpsThresholdExceeded  
jnxDfcSoftPpsUnderThreshold  
jnxEventTrap  
jnxExampleStartup  
jnxFEBSwitchover  
jnxFanFailure  
jnxFanOK  
jnxFruCheck  
jnxFruFailed  
jnxFruInsertion  
jnxFruOK  
jnxFruOffline  
jnxFruOnline  
jnxFruPowerOff  
jnxFruPowerOn  
jnxFruRemoval  
jnxHardDiskFailed  
jnxHardDiskMissing  
jnxJsAvPatternUpdateTrap  
jnxJsChassisClusterSwitchover  
jnxJsFwAuthCapacityExceeded  
jnxJsFwAuthFailure  
jnxJsFwAuthServiceDown  
jnxJsFwAuthServiceUp  
jnxJsNatAddrPoolThresholdStatus  
jnxJsScreenAttack  
jnxJsScreenCfgChange  
jnxLdpLspDown  
jnxLdpLspUp  
jnxLdpSesDown  
jnxLdpSesUp  
jnxMIMstCistPortLoopProtectStateChangeTrap  
jnxMIMstCistPortRootProtectStateChangeTrap  
jnxMIMstErrTrap  
jnxMIMstGenTrap  
jnxMIMstInvalidBpduRxdTrap  
jnxMIMstMstiPortLoopProtectStateChangeTrap  
jnxMIMstMstiPortRootProtectStateChangeTrap  
jnxMIMstNewRootTrap  
jnxMIMstProtocolMigrationTrap  
jnxMIMstRegionConfigChangeTrap  
jnxMIMstTopologyChgTrap  
jnxMacChangedNotification  
jnxMplsLdpInitSesThresholdExceeded  
jnxMplsLdpPathVectorLimitMismatch  
jnxMplsLdpSessionDown  
jnxMplsLdpSessionUp  
jnxOspfV3IfConfigError  
jnxOspfV3IfRxBadPacket  
jnxOspfV3IfStateChange

jnxOspfV3LsdbApproachingOverflow  
jnxOspfV3LsdbOverflow  
jnxOspfV3NbrRestartHelperStatusChange  
jnxOspfV3NbrStateChange  
jnxOspfV3NssaTranslatorStatusChange  
jnxOspfV3RestartStatusChange  
jnxOspfV3VirtIfConfigError  
jnxOspfV3VirtIfRxBadPacket  
jnxOspfV3VirtIfStateChange  
jnxOspfV3VirtNbrRestartHelperStatusChange  
jnxOspfV3VirtNbrStateChange  
jnxOtnAlarmCleared  
jnxOtnAlarmSet  
jnxOverTemperature  
jnxPMonOverloadCleared  
jnxPMonOverloadSet  
jnxPingEgressJitterThresholdExceeded  
jnxPingEgressStdDevThresholdExceeded  
jnxPingEgressThresholdExceeded  
jnxPingIngressJitterThresholdExceeded  
jnxPingIngressStdDevThresholdExceeded  
jnxPingIngressThresholdExceeded  
jnxPingRttJitterThresholdExceeded  
jnxPingRttStdDevThresholdExceeded  
jnxPingRttThresholdExceeded  
jnxPortBpduErrorStatusChangeTrap  
jnxPortLoopProtectStateChangeTrap  
jnxPortRootProtectStateChangeTrap  
jnxPowerSupplyFailure  
jnxPowerSupplyOK  
jnxRedundancySwitchover  
jnxRmonAlarmGetFailure  
jnxRmonGetOk  
jnxSecAccessIfMacLimitExceeded  
jnxSecAccessSdsRateLimitCrossed  
jnxSonetAlarmCleared  
jnxSonetAlarmSet  
jnxSpSvcSetCpuExceeded  
jnxSpSvcSetCpuOk  
jnxSpSvcSetZoneEntered  
jnxSpSvcSetZoneExited  
jnxStormEventNotification  
jnxSyslogTrap  
jnxTemperatureOK  
jnxVccpPortDown  
jnxVccpPortUp  
jnxVpnIfDown  
jnxVpnIfUp  
jnxVpnPwDown  
jnxVpnPwUp  
jnxl2aldGlobalMacLimit  
jnxl2aldInterfaceMacLimit  
jnxl2aldRoutingInstMacLimit  
linkDown  
linkUp  
lldpRemTablesChange  
mfrMibTrapBundleLinkMismatch  
mplsLspChange  
mplsLspDown  
mplsLspInfoChange  
mplsLspInfoDown

mplsLspInfoPathDown  
mplsLspInfoPathUp  
mplsLspInfoUp  
mplsLspPathDown  
mplsLspPathUp  
mplsLspUp  
mplsNumVrfRouteMaxThreshExceeded  
mplsNumVrfRouteMidThreshExceeded  
mplsNumVrfSecIllglLb1ThrshExcd  
mplsTunnelDown  
mplsTunnelReoptimized  
mplsTunnelRerouted  
mplsTunnelUp  
mplsVrfIfDown  
mplsVrfIfUp  
mplsXCDown  
mplsXCUp  
msdpBackwardTransition  
msdpEstablished  
newRoot  
ospfIfAuthFailure  
ospfIfConfigError  
ospfIfRxBadPacket  
ospfIfStateChange  
ospfLsdbApproachingOverflow  
ospfLsdbOverflow  
ospfMaxAgeLsa  
ospfNbrStateChange  
ospfOriginateLsa  
ospfTxRetransmit  
ospfVirtIfAuthFailure  
ospfVirtIfConfigError  
ospfVirtIfRxBadPacket  
ospfVirtIfStateChange  
ospfVirtIfTxRetransmit  
ospfVirtNbrStateChange  
pethMainPowerUsageOffNotification  
pethMainPowerUsageOnNotification  
pethPsePortOnOffNotification  
pingProbeFailed  
pingTestCompleted  
pingTestFailed  
ptopoConfigChange  
risingAlarm  
rpMauJabberTrap  
sd1cLSStatusChange  
sd1cPortStatusChange  
topologyChange  
traceRoutePathChange  
traceRouteTestCompleted  
traceRouteTestFailed  
vrrpTrapAuthFailure  
vrrpTrapNewMaster  
warmStart

#### request snmp spoof-trap (Question Mark ?)

```
user@host> request snmp spoof-trap ?
Possible completions:
<trap>                The name of the trap to spoof
ads1AtucInitFailureTrap
```

```
ads1AtucPerfESsThreshTrap
ads1AtucPerfLofsThreshTrap
ads1AtucPerfLolsThreshTrap
ads1AtucPerfLossThreshTrap
ads1AtucPerfLprsThreshTrap
ads1AtucRateChangeTrap
ads1AturPerfESsThreshTrap
ads1AturPerfLofsThreshTrap
ads1AturPerfLossThreshTrap
ads1AturPerfLprsThreshTrap
ads1AturRateChangeTrap
apsEventChannelMismatch
apsEventFEPLF
apsEventModeMismatch
apsEventPSBF
apsEventSwitchover
authenticationFailure
bfdSessDown
bfdSessUp
bgpBackwardTransition
bgpEstablished
coldStart
dlsWTrapCircuitDown
dlsWTrapCircuitUp
---(more 10%)---
```

## show snmp health-monitor

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show snmp health-monitor</code><br><code>&lt;alarms &lt;detail&gt;&gt;   &lt;logs&gt;</code>                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.0.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                             |
| <b>Description</b>              | Display information about Simple Network Management Protocol (SNMP) health monitor alarms and logs.                                                                                                                                                                |
| <b>Options</b>                  | <b>none</b> —Display information about all health monitor alarms and logs.<br><br><b>alarms &lt;detail&gt;</b> —(Optional) Display detailed information about health monitor alarms.<br><br><b>logs</b> —(Optional) Display information about health monitor logs. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>    | <a href="#">show snmp health-monitor on page 4314</a><br><a href="#">show snmp health-monitor alarms detail on page 4316</a>                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 443 on page 4312</a> describes the output fields for the <b>show snmp health-monitor</b> command. Output fields are listed in the approximate order in which they appear.                                                                        |

**Table 443: show snmp health-monitor Output Fields**

| Field Name                  | Field Description                                                           | Level of Output |
|-----------------------------|-----------------------------------------------------------------------------|-----------------|
| <b>Alarm Index</b>          | Alarm identifier.                                                           | All levels      |
| <b>Variable description</b> | Description of the health monitor object instance being monitored.          | All levels      |
| <b>Variable name</b>        | Name of the health monitor object instance being monitored.                 | All levels      |
| <b>Value</b>                | Current value of the monitored variable in the most recent sample interval. | All levels      |



Table 443: show snmp health-monitor Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>State</b>            | <p>State of the alarm or event entry:</p> <ul style="list-style-type: none"> <li>Alarms: <ul style="list-style-type: none"> <li><b>active</b>—Entry is fully configured and activated.</li> <li><b>falling threshold crossed</b>—Value of the variable has crossed the lower threshold limit.</li> <li><b>rising threshold crossed</b>—Value of the variable has crossed the upper threshold limit.</li> <li><b>under creation</b>—Entry is being configured and is not yet activated.</li> <li><b>startup</b>—Alarm is waiting for the first sample of the monitored variable.</li> <li><b>object not available</b>—Monitored variable of that type is not available to the health monitor agent.</li> <li><b>instance not available</b>—Monitored variable's instance is not available to the health monitor agent.</li> <li><b>object type invalid</b>—Monitored variable is not a numeric value.</li> <li><b>object processing errored</b>—An error occurred when the monitored variable was processed.</li> <li><b>unknown</b>—State is not one of the above.</li> </ul> </li> </ul> | All levels      |
| <b>Variable OID</b>     | Object ID to which the variable name is resolved. The format is x.x.x.x.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | detail          |
| <b>Sample type</b>      | Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of <b>absolute value</b> or <b>delta value</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | detail          |
| <b>Startup alarm</b>    | <p>Alarm that might be sent when this entry is first activated, depending on the following criteria:</p> <ul style="list-style-type: none"> <li>Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is either <b>rising alarm</b> or <b>rising or falling alarm</b>.</li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is either <b>falling alarm</b> or <b>rising or falling alarm</b>.</li> </ul> </li> <li>Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is <b>falling alarm</b>.</li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is <b>rising alarm</b>.</li> <li>Value of the alarm is between the thresholds.</li> </ul> </li> </ul>                                                                         | detail          |
| <b>Owner</b>            | Name of the entry configured by the user. If the entry was created through the CLI, the owner has <b>monitor</b> prepended to it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail          |
| <b>Creator</b>          | Mechanism by which the entry was configured ( <b>Health Monitor</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | detail          |
| <b>Sample interval</b>  | Time period between samples (in seconds).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | detail          |
| <b>Rising threshold</b> | Upper limit threshold value as a percentage of the maximum possible value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | detail          |

Table 443: show snmp health-monitor Output Fields (*continued*)

| Field Name          | Field Description                                                          | Level of Output |
|---------------------|----------------------------------------------------------------------------|-----------------|
| Falling threshold   | Lower limit threshold value as a percentage of the maximum possible value. | detail          |
| Rising event index  | Event triggered when the rising threshold is crossed.                      | detail          |
| Falling event index | Event triggered when the falling threshold is crossed.                     | detail          |

## Sample Output

### show snmp health-monitor

```
user@host> show snmp health-monitor
```

```

Alarm
Index  Variable description                               Value State

32768 Health Monitor: root file system utilization
      jnxHrStoragePercentUsed.1                      58 active

32769 Health Monitor: /config file system utilization
      jnxHrStoragePercentUsed.2                      0 active

32770 Health Monitor: RE 0 CPU utilization
      jnxOperatingCPU.9.1.0.0                        0 active

32773 Health Monitor: RE 0 Memory utilization
      jnxOperatingBuffer.9.1.0.0                     35 active

32775 Health Monitor: jkernel daemon CPU utilization
      Init daemon                                     0 active
      Chassis daemon                                  50 active
      Firewall daemon                                 0 active
      Interface daemon                                5 active
      SNMP daemon                                     11 active
      MIB2 daemon                                     42 active
      Sonet APS daemon                                0 active
      VRRP daemon                                     0 active
      Alarm daemon                                    3 active
      PFE daemon                                      0 active
      CRAFT daemon                                    0 active
      Traffic sampling control daemon                 0 active
      Ilmi daemon                                     0 active
      Remote operations daemon                        0 active
      CoS daemon                                      0 active
      Pic Services Logging daemon                     0 active
      Internal Routing Service Daemon                 3 active
      Network Access Service daemon                  0 active
      Forwarding UDP daemon                           0 active
      Routing socket proxy daemon                     0 active
      Disk Monitoring daemon                          1 active
      Inet daemon                                     0 active
      Syslog daemon                                   0 active
      Adaptive Services PIC daemon                    0 active
      ECC parity errors logging Daemon                0 active
      Layer 2 Tunneling Protocol daemon               0 active
      PPPoE daemon                                    3 active

```

|       |                                                   |              |
|-------|---------------------------------------------------|--------------|
|       | Redundancy device daemon                          | 0 active     |
|       | PPP daemon                                        | 0 active     |
|       | Dynamic Flow Capture Daemon                       | 0 active     |
| 32776 | Health Monitor: jroute daemon CPU utilization     |              |
|       | Routing protocol daemon                           | 1 active     |
|       | Management daemon                                 | 0 active     |
|       | Management daemon                                 | 0 active     |
|       | Command line interface                            | 4 active     |
|       | Periodic Packet Management daemon                 | 0 active     |
|       | Link Management daemon                            | 0 active     |
|       | Pragmatic General Multicast daemon                | 0 active     |
|       | Bidirectional Forwarding Detection daemon         | 0 active     |
|       | SRC daemon                                        | 0 active     |
|       | audit daemon                                      | 0 active     |
|       | Event daemon                                      | 0 active     |
| 32777 | Health Monitor: jcrypto daemon CPU utilization    |              |
|       | IPSec Key Management daemon                       | 0 active     |
| 32779 | Health Monitor: jkernel daemon Memory utilization |              |
|       | Init daemon                                       | 47384 active |
|       | Chassis daemon                                    | 20204 active |
|       | Firewall daemon                                   | 1956 active  |
|       | Interface daemon                                  | 3340 active  |
|       | SNMP daemon                                       | 4540 active  |
|       | MIB2 daemon                                       | 3880 active  |
|       | Sonet APS daemon                                  | 2632 active  |
|       | VRRP daemon                                       | 2672 active  |
|       | Alarm daemon                                      | 1856 active  |
|       | PFE daemon                                        | 2600 active  |
|       | CRAFT daemon                                      | 2000 active  |
|       | Traffic sampling control daemon                   | 3164 active  |
|       | Ilmi daemon                                       | 2132 active  |
|       | Remote operations daemon                          | 2964 active  |
|       | CoS daemon                                        | 3044 active  |
|       | Pic Services Logging daemon                       | 1944 active  |
|       | Internal Routing Service Daemon                   | 1392 active  |
|       | Network Access Service daemon                     | 1992 active  |
|       | Forwarding UDP daemon                             | 1876 active  |
|       | Routing socket proxy daemon                       | 1296 active  |
|       | Disk Monitoring daemon                            | 1180 active  |
|       | Inet daemon                                       | 1296 active  |
|       | Syslog daemon                                     | 1180 active  |
|       | Adaptive Services PIC daemon                      | 3220 active  |
|       | ECC parity errors logging Daemon                  | 1100 active  |
|       | Layer 2 Tunneling Protocol daemon                 | 3372 active  |
|       | PPPoE daemon                                      | 1424 active  |
|       | Redundancy device daemon                          | 1820 active  |
|       | PPP daemon                                        | 2060 active  |
|       | Dynamic Flow Capture Daemon                       | 10740 active |
| 32780 | Health Monitor: jroute daemon Memory utilization  |              |
|       | Routing protocol daemon                           | 8104 active  |
|       | Management daemon                                 | 13360 active |
|       | Management daemon                                 | 19252 active |
|       | Command line interface                            | 9912 active  |
|       | Periodic Packet Management daemon                 | 1484 active  |
|       | Link Management daemon                            | 2016 active  |
|       | Pragmatic General Multicast daemon                | 1968 active  |
|       | Bidirectional Forwarding Detection daemon         | 1956 active  |
|       | SRC daemon                                        | 1772 active  |

|              |             |
|--------------|-------------|
| audit daemon | 1772 active |
| Event daemon | 1808 active |

|                                                         |             |
|---------------------------------------------------------|-------------|
| 32781 Health Monitor: jcrypto daemon Memory utilization |             |
| IPSec Key Management daemon                             | 5600 active |

## show snmp health-monitor alarms detail

```
user@host> show snmp health-monitor alarms detail
```

```
Alarm Index 32768:
  Variable name      jnxHrStoragePercentUsed.1
  Variable OID       1.3.6.1.4.1.2636.3.31.1.1.1.1.1
  Sample type        absolute value
  Startup alarm       rising alarm
  Owner               Health Monitor: root file system
                     utilization
  Creator             Health Monitor
  State               active
  Sample interval     300 seconds
  Rising threshold    80
  Falling threshold   70
  Rising event index  32768
  Falling event index 32768
  Instance Value: 58
  Instance State: active

Alarm Index 32769:
  Variable name      jnxHrStoragePercentUsed.2
  Variable OID       1.3.6.1.4.1.2636.3.31.1.1.1.1.2
  Sample type        absolute value
  Startup alarm       rising alarm
  Owner               Health Monitor: /config file system
                     utilization
  Creator             Health Monitor
  State               active
  Sample interval     300 seconds
  Rising threshold    80
  Falling threshold   70
  Rising event index  32768
  Falling event index 32768
  Instance Value: 0
  Instance State: active

Alarm Index 32770:
  Variable name      jnxOperatingCPU.9.1.0.0
  Variable OID       1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0
  Sample type        absolute value
  Startup alarm       rising alarm
  Owner               Health Monitor: RE 0 CPU utilization

  Creator             Health Monitor
  State               active
  Sample interval     300 seconds
  Rising threshold    80
  Falling threshold   70
  Rising event index  32768
  Falling event index 32768
  Instance Value: 0
  Instance State: active
```

## Alarm Index 32773:

Variable name jnxOperatingBuffer.9.1.0.0  
 Variable OID 1.3.6.1.4.1.2636.3.1.13.1.11.9.1.0.0  
 Sample type absolute value  
 Startup alarm rising alarm  
 Owner Health Monitor: RE 0 Memory utilization

Creator Health Monitor  
 State active  
 Sample interval 300 seconds  
 Rising threshold 80  
 Falling threshold 70  
 Rising event index 32768  
 Falling event index 32768  
 Instance Value: 35  
 Instance State: active

## Alarm Index 32775:

Variable name sysAppElmtRunCPU.3  
 Variable OID 1.3.6.1.2.1.54.1.2.3.1.9.3  
 Sample type delta value  
 Startup alarm rising alarm  
 Owner Health Monitor: jkernel daemon CPU utilization

Creator Health Monitor  
 State active  
 Sample interval 300 seconds  
 Rising threshold 24000  
 Falling threshold 21000  
 Rising event index 32768  
 Falling event index 32768  
 Instance Name: sysAppElmtRunCPU.3.1.1  
 Instance Description: Init daemon  
 Instance Value: 0  
 Instance State: active

Instance Name: sysAppElmtRunCPU.3.2.2786  
 Instance Description: Chassis daemon  
 Instance Value: 50  
 Instance State: active

Instance Name: sysAppElmtRunCPU.3.3.2938  
 Instance Description: Firewall daemon  
 Instance Value: 0  
 Instance State: active

Instance Name: sysAppElmtRunCPU.3.4.2942  
 Instance Description: Interface daemon  
 Instance Value: 5  
 Instance State: active

Instance Name: sysAppElmtRunCPU.3.7.7332  
 Instance Description: SNMP daemon  
 Instance Value: 11  
 Instance State: active

Instance Name: sysAppElmtRunCPU.3.9.2914  
 Instance Description: MIB2 daemon  
 Instance Value: 42

```
Instance State: active

Instance Name: sysAppElemRunCPU.3.12.2916
Instance Description: Sonet APS daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElemRunCPU.3.13.2917
Instance Description: VRRP daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElemRunCPU.3.14.2787
Instance Description: Alarm daemon
Instance Value: 3
Instance State: active

Instance Name: sysAppElemRunCPU.3.15.2940
Instance Description: PFE daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElemRunCPU.3.16.2788
Instance Description: CRAFT daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElemRunCPU.3.17.2918
Instance Description: Traffic sampling control daemon
---(more 23%)---
```

## show snmp inform-statistics

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show snmp inform-statistics                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.           |
| <b>Description</b>              | Display information about Simple Network Management Protocol (SNMP) inform requests.                                                                                                           |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                   |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">show snmp inform-statistics on page 4319</a>                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 444 on page 4319</a> describes the output fields for the <b>show snmp inform-statistics</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 444: show snmp inform-statistics Output Fields**

| Field Name            | Field Description                                                                                                                                  |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Target Name</b>    | Name of the device configured to receive and respond to SNMP informs.                                                                              |
| <b>Address</b>        | IP address of the target device.                                                                                                                   |
| <b>Sent</b>           | Number of informs sent to the target device and acknowledged by the target device.                                                                 |
| <b>Pending</b>        | Number of informs held in memory pending a response from the target device.                                                                        |
| <b>Discarded</b>      | Number of informs discarded after the specified number of retransmissions to the target device were attempted.                                     |
| <b>Timeouts</b>       | Number of informs that did not receive an acknowledgement from the target device within the timeout specified.                                     |
| <b>Probe Failures</b> | Connection failures that occurred (for example, when the target server returned invalid content or you incorrectly configured the target address). |

## Sample Output

### show snmp inform-statistics

```

user@host> show snmp inform-statistics
Inform Request Statistics:
  Target Name: TA1_v3_md5_none Address: 172.17.20.184
    Sent: 176, Pending: 0
    Discarded: 0, Timeouts: 0, Probe Failures: 0
  Target Name: TA2_v3_sha_none Address: 192.168.110.59

```

Sent: 0, Pending: 4  
Discarded: 84, Timeouts: 0, Probe Failures: 258  
Target Name: TA5\_v2\_none Address: 172.17.20.184  
Sent: 0, Pending: 0  
Discarded: 2, Timeouts: 10, Probe Failures: 0



## show snmp mib

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show snmp mib (get   get-next   walk) (ascii   decimal) <i>object-id</i></code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>ascii</b> and <b>decimal</b> options introduced in Junos OS Release 9.6.</p> <p><b>ascii</b> and <b>decimal</b> options introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display local Simple Network Management Protocol (SNMP) Management Information Base (MIB) object values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>get</b>—Retrieve and display one or more SNMP object values.</p> <p><b>get-next</b>—Retrieve and display the next SNMP object values.</p> <p><b>walk</b>—Retrieve and display the SNMP object values that are associated with the requested object identifier (OID). When you use this option, the Junos OS displays the objects below the subtree that you specify.</p> <p><b>ascii</b>—Display the SNMP object's string indices as an ASCII-key representation.</p> <p><b>decimal</b>—Display the SNMP object values in the decimal (default) format. The <b>decimal</b> option is the default option for this command. Therefore, issuing the <b>show snmp mib (get   get-next   walk) decimal object-id</b> and the <b>show snmp mib (get   get-next   walk) object-id</b> commands display the same output.</p> <p><b>object-id</b>—The object can be represented by a sequence of dotted integers (such as 1.3.6.1.2.1.2) or by its subtree name (such as <b>interfaces</b>). When entering multiple objects, enclose the objects in quotation marks.</p> |
| <b>Required Privilege Level</b> | snmp—To view this statement in the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <p><a href="#">show snmp mib get on page 4322</a></p> <p><a href="#">show snmp mib get (Multiple Objects) on page 4322</a></p> <p><a href="#">show snmp mib get (Layer 2 Policer) on page 4322</a></p> <p><a href="#">show snmp mib get-next on page 4322</a></p> <p><a href="#">show snmp mib get-next (Specify an OID) on page 4322</a></p> <p><a href="#">show snmp mib walk on page 4322</a></p> <p><a href="#">show snmp mib walk (QFX Series) on page 4322</a></p> <p><a href="#">show snmp mib walk decimal on page 4323</a></p> <p><a href="#">show snmp mib walk (ASCII) on page 4323</a></p> <p><a href="#">show snmp mib walk (Multiple Indices) on page 4323</a></p> <p><a href="#">show snmp mib walk decimal (Multiple Indices) on page 4323</a></p>                                                                                                                                                                                                                                                                                                    |
| <b>Output Fields</b>            | <p>Table 445 on page 4322 describes the output fields for the <b>show snmp mib</b> command.</p> <p>Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 445: show snmp mib Output Fields

| Field Name          | Field Description                                                               |
|---------------------|---------------------------------------------------------------------------------|
| <i>name</i>         | Object name and numeric instance value.                                         |
| <i>object value</i> | Object value. The Junos OS translates OIDs into the corresponding object names. |

## Sample Output

### show snmp mib get

```
user@host> show snmp mib get sysObjectID.0
sysObjectID.0 = jnxProductNameM20
```

### show snmp mib get (Multiple Objects)

```
user@host> show snmp mib get ?sysObjectID.0 sysUpTime.0?
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
```

### show snmp mib get (Layer 2 Policer)

```
user@host> show snmp mib get ifInOctets.25970
ifInOctets.25970 = 7545720
```

### show snmp mib get-next

```
user@host> show snmp mib get-next jnxMibs
jnxBoxClass.0 = jnxProductLineM20.0
```

### show snmp mib get-next (Specify an OID)

```
user@host> show snmp mib get-next 1.3.6.1
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release: 2004-1 Build date: build date UTC Copyright (c) 1996-2004 Juniper
Networks, Inc.
```

### show snmp mib walk

```
user@host> show snmp mib walk system
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release #0: 2004-1 Build date: build date UTC Copyright (c) 1996-2004
Juniper Networks, Inc.
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
sysContact.0 = Your contact
sysName.0 = my router
sysLocation.0 = building 1
sysServices.0 = 4
```

### show snmp mib walk (QFX Series)

```
user@switch> show snmp mib walk system
sysDescr.0 = Juniper Networks, Inc. qfx3500s internet router, kernel JUNOS
11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC Build date: 2010-09-26 06:00:10
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0 = 138980301
sysContact.0 = System Contact
```

```

sysName.0      = LabQFX3500
sysLocation.0 = Lab
sysServices.0 = 4

```

#### show snmp mib walk decimal

```

user@host show snmp mib walk decimal jnxUtilData
jnxUtilCounter32Value.102.114.101.100 = 100

```

#### show snmp mib walk (ASCII)

```

show snmp mib walk ascii jnxUtilData
jnxUtilCounter32Value."fred" = 100

```

#### show snmp mib walk (Multiple Indices)

```

show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....

```

#### show snmp mib walk decimal (Multiple Indices)

```

show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....

```

## show snmp rmon

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show snmp rmon<br><alarms <brief   detail>   events <brief   detail>   logs>                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display information about Simple Network Management Protocol (SNMP) Remote Monitoring (RMON) alarms and events.                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>none</b>—Display information about all RMON alarms and events.</p> <p><b>alarms</b>—(Optional) Display information about RMON alarms.</p> <p><b>brief   detail</b>—(Optional) Display brief or detailed information about RMON alarms or events.</p> <p><b>events</b>—(Optional) Display information about RMON events.</p> <p><b>logs</b>—(Optional) Display information about RMON monitoring logs.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show snmp rmon on page 4326</a><br><a href="#">show snmp rmon alarms detail on page 4326</a><br><a href="#">show snmp rmon events detail on page 4327</a>                                                                                                                                                                                                                                           |
| <b>Output Fields</b>            | Table 446 on page 4324 describes the output fields for the <b>show snmp rmon</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                               |

Table 446: show snmp rmon Output Fields

| Field Name  | Field Description | Level of Output |
|-------------|-------------------|-----------------|
| Alarm Index | Alarm identifier. | All levels      |

Table 446: show snmp rmon Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>State</b>         | <p>State of the alarm or event entry:</p> <p>Alarms:</p> <ul style="list-style-type: none"> <li>• <b>active</b>—Entry is fully configured and activated.</li> <li>• <b>falling threshold crossed</b>—Value of the variable has crossed the lower threshold limit.</li> <li>• <b>rising threshold crossed</b>—Value of the variable has crossed the upper threshold limit.</li> <li>• <b>under creation</b>—Entry is being configured and is not yet activated.</li> <li>• <b>startup</b>—Alarm is waiting for the first sample of the monitored variable.</li> <li>• <b>object not available</b>—Monitored variable of that type is not available to the SNMP agent.</li> <li>• <b>instance not available</b>—Monitored variable's instance is not available to the SNMP agent.</li> <li>• <b>object type invalid</b>—Monitored variable is not a numeric value.</li> <li>• <b>object processing errored</b>—An error occurred when the monitored variable was processed.</li> <li>• <b>unknown</b>—State is not one of the above.</li> </ul> <p>Events:</p> <ul style="list-style-type: none"> <li>• <b>active</b>—Entry has been fully configured and activated.</li> <li>• <b>under creation</b>—Entry is being configured and is not yet activated.</li> <li>• <b>unknown</b>—State is not one of the above.</li> </ul> | All levels      |
| <b>Variable name</b> | Name of the SNMP object instance being monitored.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels      |
| <b>Event Index</b>   | Event identifier.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | All levels      |
| <b>Type</b>          | <p>Type of notification made when an event is triggered. It can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>log</b>—A system log message is generated and an entry is made to the log table.</li> <li>• <b>snmptrap</b>—An SNMP trap is sent to the configured destination.</li> <li>• <b>log and trap</b>—A system log message is generated, an entry is made to the log table, and an SNMP trap is sent to the configured destination.</li> <li>• <b>none</b>—Neither log nor trap will be sent.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b>   |
| <b>Last Event</b>    | Date and time of the last event. It has the format <i>yyyy-mm-dd hh:mm:ss timezone</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>brief</b>    |
| <b>Community</b>     | Identifies the trap group used for sending the SNMP trap.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail</b>   |
| <b>Variable OID</b>  | Object ID to which the variable name is resolved. The format is x.x.x.x.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b>   |
| <b>Sample type</b>   | Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of <b>absolute value</b> or <b>delta value</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail</b>   |

Table 446: show snmp rmon Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Startup alarm</b>       | Alarm that might be sent when this entry is first activated, depending on the following criteria: <ul style="list-style-type: none"> <li>Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is either <b>rising alarm</b> or <b>rising or falling alarm</b>.</li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is either <b>falling alarm</b> or <b>rising or falling alarm</b>.</li> </ul> </li> <li>Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is <b>falling alarm</b>.</li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is <b>rising alarm</b>.</li> <li>Value of the alarm is between the thresholds.</li> </ul> </li> </ul> | <b>detail</b>   |
| <b>Owner</b>               | Name of the entry configured by the user. If the entry was created through the CLI, the owner has <b>monitor</b> prepended to it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail</b>   |
| <b>Creator</b>             | Mechanism by which the entry was configured ( <b>CLI</b> or <b>SNMP</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail</b>   |
| <b>Sample interval</b>     | Time period between samples (in seconds).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b>   |
| <b>Rising threshold</b>    | Upper limit threshold value configured by the user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b>   |
| <b>Falling threshold</b>   | Lower limit threshold value configured by the user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b>   |
| <b>Rising event index</b>  | Event triggered when the rising threshold is crossed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |
| <b>Falling event index</b> | Event triggered when the falling threshold is crossed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail</b>   |
| <b>Current value</b>       | Current value of the monitored variable in the most recent sample interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail</b>   |

## Sample Output

### show snmp rmon

```

user@host> show snmp rmon
Alarm
Index  State                               Variable name
  1    falling threshold crossed         ifInOctets.1

Event
Index  Type                                Last Event
  1    log and trap                     2002-01-30 01:13:01 PST

```

### show snmp rmon alarms detail

```

user@host> show snmp rmon alarms detail

```

```
Alarm Index 1:
Variable name      ifInOctets.1
Variable OID       1.3.6.1.2.1.2.2.1.10.1
Sample type        delta value
Startup alarm      rising or falling alarm
Owner              monitor
Creator            CLI
State              falling threshold crossed
Sample interval    60 seconds
Rising threshold   100000
Falling threshold  80000
Rising event index 1
Falling event index 1
Current value      0
```

#### show snmp rmon events detail

```
user@host> show snmp rmon events detail
Event Index 1:
Type          log and trap
Community     boy-elroy
Last event    2002-01-30 01:13:01 PST
Creator       CLI
State         active
```

## show snmp rmon history

|                                 |                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show snmp rmon history</b><br><i>&lt;history-index&gt;</i><br><i>&lt;sample-index&gt;</i>                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Display the contents of the RMON history group.                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>none</b>—Display all the entries in the RMON history group.</p> <p><b>history-index</b>—(Optional) Display the contents of the specified entry in the RMON history group.</p> <p><b>sample-index</b>—(Optional) Display the statistics collected for the specified sample within the specified entry in the RMON history group.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear snmp rmon history on page 4303</a></li> </ul>                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>    | <a href="#">show snmp rmon history 1 on page 4329</a><br><a href="#">show snmp rmon history 1 sample 15 on page 4330</a>                                                                                                                                                                                                                  |
| <b>Output Fields</b>            | Table 447 on page 4328 lists the output fields for the <b>show smp rmon history</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                      |

Table 447: show smp rmon history Output Fields

| Field Name               | Field Description                                                                            |
|--------------------------|----------------------------------------------------------------------------------------------|
| History Index            | Identifies this RMON history entry within the RMON history group.                            |
| Owner                    | The entity that configured this entry. Range is 0 to 32 alphanumeric characters.             |
| Status                   | The status of the RMON history entry.                                                        |
| Interface or Data Source | The ifindex object that identifies the interface that is being monitored.                    |
| Interval                 | The interval (in seconds) configured for this RMON history entry.                            |
| Buckets Requested        | The requested number of buckets ( <b>intervals</b> ) configured for this RMON history entry. |
| Buckets Granted          | The number of buckets granted for this RMON history entry.                                   |



Table 447: show snmp rmon history Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Sample Index</b> | <p>The sample statistics taken at the specified interval.</p> <ul style="list-style-type: none"> <li>• <b>Drop Events</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Octets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Packets</b>—Total number of packets.</li> <li>• <b>Broadcast Packets</b>—Number of broadcast packets.</li> <li>• <b>Multicast Packets</b>—Number of multicast packets.</li> <li>• <b>CRC errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS error) or a bad FCS with a nonintegral number of octets (alignment error).</li> <li>• <b>Undersize Pkts</b>—Number of packets received during this sampling interval that were less than 64 octets long (excluding framing bits but including FCS octets) and were otherwise well formed.</li> <li>• <b>Oversize Pkts</b>—Number of packets received during the sampling interval that were longer than 1518 octets (excluding framing bits, but including FCS octets) but were otherwise well formed.</li> <li>• <b>Fragments</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>Jabbers</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Utilization(%)</b>—The best estimate of the mean physical layer network utilization on this interface during this sampling interval, in hundredths of a percent.</li> </ul> |

## Sample Output

### show snmp rmon history 1

```

user@host> show snmp rmon history 1
History Index 1:
Interface                171
Requested Buckets        50
Interval                 10

Sample Index 1: Interval Start: Tue Feb 12 04:12:32 2008
Drop Events              0
Octets                   486
Packets                  2

```

|                   |   |
|-------------------|---|
| Broadcast Packet  | 0 |
| Multicast Packets | 2 |
| CRC errors        | 0 |
| Undersize Pkts    | 0 |
| Oversize Pkts     | 0 |
| Fragments         | 0 |
| Jabbers           | 0 |
| Collisions        | 0 |
| Utilization(%)    | 0 |

Sample Index 2: Interval Start: Tue Feb 12 04:12:42 2008

|                   |     |
|-------------------|-----|
| Drop Events       | 0   |
| Octets            | 486 |
| Packets           | 2   |
| Broadcast Packet  | 0   |
| Multicast Packets | 2   |
| CRC errors        | 0   |
| Undersize Pkts    | 0   |
| Oversize Pkts     | 0   |
| Fragments         | 0   |
| Jabbers           | 0   |
| Collisions        | 0   |
| Utilization(%)    | 0   |

Sample Index 3: Interval Start: Tue Feb 12 04:12:52 2008

|                   |     |
|-------------------|-----|
| Drop Events       | 0   |
| Octets            | 486 |
| Packets           | 2   |
| Broadcast Packet  | 0   |
| Multicast Packets | 2   |
| CRC errors        | 0   |
| Undersize Pkts    | 0   |
| Oversize Pkts     | 0   |
| Fragments         | 0   |
| Jabbers           | 0   |
| Collisions        | 0   |
| Utilization(%)    | 0   |

### show snmp rmon history 1 sample 15

```
user@host> show snmp rmon history 1 sample 15
```

```
Index 1
Owner    = monitor
Status   = valid
Data Source = ifIndex.17
Interval = 1800
Buckets Requested = 50
Buckets Granted = 50
```

Sample Index 44: Interval Start: Thu Jan 1 00:08:35 1970

|                |     |
|----------------|-----|
| Drop Events    | = 0 |
| Octets         | = 0 |
| Packets        | = 0 |
| Broadcast Pkts | = 0 |
| Multicast Pkts | = 0 |
| CRC Errors     | = 0 |
| Undersize Pkts | = 0 |
| Oversize Pkts  | = 0 |
| Fragments      | = 0 |
| Jabbers        | = 0 |

```
Collisions = 0  
Utilization (%) = 0
```

## show snmp statistics

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show snmp statistics                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Display statistics about Simple Network Management Protocol (SNMP) packets sent and received by the router or switch.                                                                    |
| <b>Options</b>                  | This command has no options.                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear snmp statistics on page 4304</a></li> </ul>                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">show snmp statistics on page 4335</a>                                                                                                                                        |
| <b>Output Fields</b>            | <a href="#">Table 448 on page 4332</a> describes the output fields for the <b>show snmp statistics</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 448: show snmp statistics Output Fields**

| Field Name   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input</b> | <p>Information about received packets:</p> <ul style="list-style-type: none"> <li><b>Packets(snmplnPkts)</b>—Total number of messages delivered to the SNMP entity from the transport service.</li> <li><b>Bad versions—(snmplnBadVersions)</b> Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version.</li> <li><b>Bad community names—(snmplnBadCommunityNames)</b> Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity.</li> <li><b>Bad community uses—(snmplnBadCommunityUses)</b> Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message.</li> <li><b>ASN parse errors—(snmplnASNParseErrs)</b> Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.</li> <li><b>Too big—(snmplnTooBig)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>tooBig</b>.</li> <li><b>No such names—(snmplnNoSuchNames)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>noSuchName</b>.</li> <li><b>Bad values—(snmplnBadValues)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>badValue</b>.</li> <li><b>Read only—(snmplnReadOnly)</b> Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of <b>readOnly</b>. Only incorrect implementations of SNMP generate this error.</li> </ul> |

Table 448: show snmp statistics Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input (continued) | <ul style="list-style-type: none"> <li>• <b>General errors—(snmpInGenErrs)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>genErr</b>.</li> <li>• <b>Total requests varbinds—(snmpInTotalReqVars)</b> Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP <b>GetRequest</b> and <b>GetNext</b> PDUs.</li> <li>• <b>Total set varbinds—(snmpInSetVars)</b> Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP <b>SetRequest</b> PDUs.</li> <li>• <b>Get requests—(snmpInGetRequests)</b> Total number of SNMP <b>GetRequest</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Get nexts—(snmpInGetNexts)</b> Total number of SNMP <b>GetNext</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Set requests—(snmpInSetRequests)</b> Total number of SNMP <b>SetRequest</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Get responses—(snmpInGetResponses)</b> Total number of SNMP <b>GetResponse</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Traps—(snmpInTraps)</b> Total number of SNMP traps generated by the SNMP entity.</li> <li>• <b>Silent drops—(snmpSilentDrops)</b> Total number of <b>GetRequest</b>, <b>GetNextRequest</b>, <b>GetBulkRequest</b>, <b>SetRequests</b>, and <b>InformRequest</b> PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests.</li> <li>• <b>Proxy drops—(snmpProxyDrops)</b> Total number of <b>GetRequest</b>, <b>GetNextRequest</b>, <b>GetBulkRequest</b>, <b>SetRequests</b>, and <b>InformRequest</b> PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned.</li> <li>• <b>Commit pending drops</b>—Number of SNMP packets for <b>Set</b> requests dropped because of a previous pending SNMP <b>Set</b> request on the committed configuration.</li> <li>• <b>Throttle drops</b>—Number of SNMP packets for any requests dropped reaching the throttle limit.</li> </ul> |

Table 448: show snmp statistics Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| V3 Input   | <p>Information about SNMP version 3 packets:</p> <ul style="list-style-type: none"> <li>• <b>Unknown security models—(snmpUnknownSecurityModels)</b> Total number of packets received by the SNMP engine that were dropped because they referenced a security model that was not known to or supported by the SNMP engine.</li> <li>• <b>Invalid messages—(snmpInvalidMsgs)</b> Number of packets received by the SNMP engine that were dropped because there were invalid or inconsistent components in the SNMP message.</li> <li>• <b>Unknown pdu handlers—(snmpUnknownPDUHandlers)</b> Number of packets received by the SNMP engine that were dropped because the PDU contained in the packet could not be passed to an application responsible for handling the PDU type.</li> <li>• <b>Unavailable contexts—(snmpUnavailableContexts)</b> Number of requests received for a context that is known to the SNMP engine, but is currently unavailable.</li> <li>• <b>Unknown contexts—(snmpUnknownContexts)</b> Total number of requests received for a context that is unknown to the SNMP engine.</li> <li>• <b>Unsupported security levels—(usmStatsUnsupportedSecLevels)</b> Total number of packets received by the SNMP engine that were dropped because they requested a security level unknown to the SNMP engine (or otherwise unavailable).</li> <li>• <b>Not in time windows—(usmStatsNotInTimeWindows)</b> Total number of packets received by the SNMP engine that were dropped because they appeared outside the authoritative SNMP engine's window.</li> <li>• <b>Unknown user names—(usmStatsUnknownUserNames)</b> Total number of packets received by the SNMP engine that were dropped because they referenced a user that was not known to the SNMP engine.</li> <li>• <b>Unknown engine ids—(usmStatsUnknownEngineIDs)</b> Total number of packets received by the SNMP engine that were dropped because they referenced an SNMP engine ID that was not known to the SNMP engine.</li> <li>• <b>Wrong digests—(usmStatsWrongDigests)</b> Total number of packets received by the SNMP engine that were dropped because they did not contain the expected digest value.</li> <li>• <b>Decryption errors—(usmStatsDecryptionErrors)</b> Total number of packets received by the SNMP engine that were dropped because they could not be decrypted.</li> </ul> |

Table 448: show snmp statistics Output Fields (*continued*)

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Output</b> | <p>Information about transmitted packets:</p> <ul style="list-style-type: none"> <li>• <b>Packets—(snmpOutPkts)</b> Total number of messages passed from the SNMP entity to the transport service.</li> <li>• <b>Too big—(snmpOutTooBig)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>tooBig</b>.</li> <li>• <b>No such names—(snmpOutNoSuchNames)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>noSuchName</b>.</li> <li>• <b>Bad values—(snmpOutBadValues)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>badValue</b>.</li> <li>• <b>General errors—(snmpOutGenErrs)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>genErr</b>.</li> <li>• <b>Get requests—(snmpOutGetRequests)</b> Total number of SNMP <b>GetRequest</b> PDUs generated by the SNMP entity.</li> <li>• <b>Get nexts—(snmpOutGetNexts)</b> Total number of SNMP <b>GetNext</b> PDUs generated by the SNMP entity.</li> <li>• <b>Set requests—(snmpOutSetRequests)</b> Total number of SNMP <b>SetRequest</b> PDUs generated by the SNMP entity.</li> <li>• <b>Get responses—(snmpOutGetResponses)</b> Total number of SNMP <b>GetResponse</b> PDUs generated by the SNMP entity.</li> <li>• <b>Traps—(snmpOutTraps)</b> Total number of SNMP traps generated by the SNMP entity.</li> </ul> |

## Sample Output

### show snmp statistics

```

user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 246213, Bad versions: 12, Bad community names: 12,
    Bad community uses: 0, ASN parse errors: 96,
    Too big: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 227084, Total set varbinds: 67,
    Get requests: 44942, Get nexts: 190371, Set requests: 10712,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0,
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 1
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
  Output:
    Packets: 246093, Too big: 0, No such names: 31561,
    Bad values: 0, General errors: 2,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 246025, Traps: 0

```

## show snmp v3

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show snmp v3</code><br><code>&lt;access &lt;brief   detail&gt;   community   general   groups   notify &lt;filter&gt;   target &lt;address   parameters&gt;   users&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display the Simple Network Management Protocol version 3 (SNMPv3) operating configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>none</b>—Display all of the SNMPv3 operating configuration.</p> <p><b>access</b>—(Optional) Display SNMPv3 access information.</p> <p><b>brief   detail</b>—(Optional) Display brief or detailed information about SNMPv3 access information.</p> <p><b>community</b>—(Optional) Display SNMPv3 community information.</p> <p><b>general</b>—(Optional) Display SNMPv3 general information.</p> <p><b>groups</b>—(Optional) Display SNMPv3 security-to-group information.</p> <p><b>notify &lt;filter&gt;</b>—(Optional) Display SNMPv3 notify and, optionally, notify filter information.</p> <p><b>target &lt;address   parameters&gt;</b>—(Optional) Display SNMPv3 target and, optionally, either target address or target parameter information.</p> <p><b>users</b>—(Optional) Display SNMPv3 user information.</p> |
| <b>Additional Information</b>   | To edit the default display of the <b>show snmp v3</b> command, specify options in the <b>show</b> statement at the <b>[edit snmp v3]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show snmp v3 on page 4338</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 449 on page 4337</a> describes the output fields for the <b>show snmp v3</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



Table 449: show snmp v3 Output Fields

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Access control</b> | <p>Information about access control:</p> <ul style="list-style-type: none"> <li>• <b>Group</b>—Group name for which the configured access privileges apply. The group, together with the context prefix and the security model and security level, forms the index for this table.</li> <li>• <b>Context prefix</b>—SNMPv3 context for which the configured access privileges apply.</li> <li>• <b>Security model/level</b>—Security model and security level for which the configuration access privileges apply.</li> <li>• <b>Read view</b>—Identifies the MIB view applied to SNMPv3 read operations.</li> <li>• <b>Write view</b>—Identifies the MIB view applied to SNMPv3 write operations.</li> <li>• <b>Notify view</b>—Identifies the MIB view applied to outbound SNMP notifications.</li> </ul>                                                                                                                                                                                                                                       |
| <b>Engine</b>         | <p>Information about local engine configuration:</p> <ul style="list-style-type: none"> <li>• <b>Local engine ID</b>—Identifier that uniquely and unambiguously identifies the local SNMPv3 engine.</li> <li>• <b>Engine boots</b>—Number of times the local SNMPv3 engine has rebooted or reinitialized since the engine ID was last changed.</li> <li>• <b>Engine time</b>—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized.</li> <li>• <b>Max msg size</b>—Maximum message size the sender can accommodate.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Engine ID</b>      | <p>Information about engine ID:</p> <ul style="list-style-type: none"> <li>• <b>Local engine ID</b>—Identifier that uniquely and unambiguously identifies the local SNMPv3 engine.</li> <li>• <b>Engine boots</b>—Number of times the local SNMPv3 engine has rebooted or reinitialized since the engine ID was last changed.</li> <li>• <b>Engine time</b>—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized.</li> <li>• <b>Max msg size</b>—Maximum message size the sender can accommodate.</li> <li>• <b>Engine ID</b>—SNMPv3 engine ID associated with each user.</li> <li>• <b>User</b>—SNMPv3 user.</li> <li>• <b>Auth/Priv</b>—Authentication and encryption algorithm available for use by each user.</li> <li>• <b>Storage</b>—Indicates whether a user is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status.</li> <li>• <b>Status</b>—Status of the conceptual row. Only rows with an active status are used by the SNMPv3 engine.</li> </ul> |
| <b>Group name</b>     | Name of the group to which this entry belongs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Security model</b> | Identifies the security model context for the security name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Security name</b>  | Used with the security model; identifies a specific security name instance. Each security model/security name combination can be assigned to a specific group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Storage type</b>   | Indicates whether a user is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Status</b>         | Status of the conceptual row. Only rows with active status are used by the SNMPv3 engine.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## Sample Output

### show snmp v3

```
user@host> show snmp v3
Local engine ID: 80 00 0a 4c e04 31 32 33 34
Engine boots:      38
Engine time:       64583 seconds
Max msg size:      2048 bytes

Engine ID: local
  User          Auth/Priv  Storage  Status
  user1         md5/des   nonvolatile active
  user2         sha/none  nonvolatile active
  user3         none/none nonvolatile active

Engine ID: 81 00 0a 4c 04 64 64 64 64
  User          Auth/Priv  Storage  Status
  UNEW         md5/none  nonvolatile active
Group name      Security model  Security name  Storage type  Status
g1             usm          user1          nonvolatile active
g2             usm          user2          nonvolatile active
g3             usm          user3          nonvolatile active

Access control:
Group           Context prefix  Security model/level  Read view  Write view  Notify view
g1              usm/privacy v1              v1
g2              usm/authent v1              v1
g3              usm/none   v1              v1
```

---

## Operational Commands: Analyzers and Port Mirroring

- [show forwarding-options port-mirroring](#)
- [show forwarding-options analyzer](#)

## show forwarding-options port-mirroring

|                                 |                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show forwarding-options port-mirroring</b><br><b>&lt;terse   detail&gt;</b><br><b>&lt;instance-name&gt;</b>                                                                                                               |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br>Hierarchy level <b>[edit forwarding-options]</b> introduced in Junos OS Release 13.2X50-D10 (ELS). |
| <b>Description</b>              | Display current state of port-mirroring instances.                                                                                                                                                                           |
| <b>Options</b>                  | <b>terse   detail</b> —(Optional) Display the specified level of output.<br><br><b>instance-name</b> —(Optional) Display a single port-mirroring instance.                                                                   |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">[edit forwarding-options port-mirroring] Configuration Statement Hierarchy on page 4157</a></li> </ul>                                                                    |
| <b>List of Sample Output</b>    | <a href="#">show forwarding-options port-mirroring terse on page 4340</a><br><a href="#">show forwarding-options port-mirroring detail on page 4340</a>                                                                      |
| <b>Output Fields</b>            | <a href="#">Table 450 on page 4339</a> lists the output fields for the <b>show forwarding-options port-mirroring</b> command. Output fields are listed in the approximate order in which they appear.                        |

**Table 450: show forwarding-options port-mirroring Output Fields**

| Field Name            | Field Description                                   | Level of Output |
|-----------------------|-----------------------------------------------------|-----------------|
| Instance Name         | Name of port-mirroring instance.                    | All levels      |
| Instance Id           | Instance identification number.                     | All levels      |
| State                 | Instance state, either <b>up</b> or <b>down</b> .   | All levels      |
| Input parameters      |                                                     |                 |
| Rate                  | Rate (ratio of packets sampled).                    | detail          |
| Run-length            | Run length (number of consecutive packets sampled). | detail          |
| Maximum-packet-length | Maximum packet length.                              | detail          |
| Output parameters     |                                                     |                 |
| Family                | Protocol family.                                    | detail          |
| State                 | Instance state, either <b>up</b> or <b>down</b> .   | detail          |

Table 450: show forwarding-options port-mirroring Output Fields (*continued*)

| Field Name  | Field Description                  | Level of Output |
|-------------|------------------------------------|-----------------|
| Destination | Destination (next-hop group name). | detail          |

## Sample Output

### show forwarding-options port-mirroring terse

```

user@switch> show forwarding-options port-mirroring terse
Instance Name      Instance Id  State
&global_instance   1           up
inst1              2           up

```

### show forwarding-options port-mirroring detail

```

user@switch> show forwarding-options port-mirroring detail
Instance Name: &global_instance
Instance Id: 1      State: up
  Input parameters:
    Rate:          1
    Run-length:    0
    Maximum-packet-length: 0
  Output parameters:
    Family: ethernet-switching  State: up      Destination: ge-0/0/10.0

Instance Name: inst1
Instance Id: 2      State: up
  Input parameters:
    Rate:          1
    Run-length:    0
    Maximum-packet-length: 0
  Output parameters:
    Family: ethernet-switching  State: down    Destination: ge-0/0/10.0

```

## show forwarding-options analyzer

|                                 |                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show forwarding-options analyzer <i>analyzer-name</i></b>                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2 for EX Series switches.<br>Command introduced in Junos OS Release 14.1 for MX Series routers.                                                                        |
| <b>Description</b>              | Display information about analyzers configured for mirroring.                                                                                                                                                    |
| <b>Options</b>                  | <b><i>analyzer-name</i></b> —(Optional) Displays the status of a specific analyzer on the switch.                                                                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring and Analyzers on EX4300 Switches on page 3964</a></li> <li>• <a href="#">Understanding Port Mirroring Analyzers</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show forwarding-options analyzer on page 4341</a>                                                                                                                                                    |
| <b>Output Fields</b>            | <a href="#">Table 451 on page 4341</a> lists the output fields for the <b>show forwarding-options analyzer</b> command. Output fields are listed in the approximate order in which they appear.                  |

**Table 451: show forwarding-options analyzer Output Fields**

| Field Name                   | Field Description                                                                                                            |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| Analyzer name                | Displays the name of the analyzer.                                                                                           |
| Mirror rate                  | Displays the ratio of packets to be mirrored.                                                                                |
| Maximum packet length        | Displays the maximum packet length used for port mirroring.                                                                  |
| State                        | Specifies the Instance state, either <b>up</b> or <b>down</b> .                                                              |
| Ingress monitored interfaces | Displays interfaces for which traffic entering the interfaces is mirrored.                                                   |
| Output VLAN                  | Specifies a VLAN to which mirrored packets are sent. An analyzer can have output to either an interface or a VLAN, not both. |

## Sample Output

### show forwarding-options analyzer

```

user@switch> show forwarding-options analyzer
Analyzer name           : employee-monitor
Mirror rate             : 1
Maximum packet length   : 0
State                   : up
Ingress monitored interfaces : ge-0/0/0.0
Ingress monitored interfaces : ge-0/0/1.0

```

Output VLAN : default-switch/remote-analyzer

---

## Commands for Network Analytics

- [monitor start \(Analytics\)](#)
- [show analytics collector](#)
- [show analytics configuration](#)
- [show analytics queue-statistics](#)
- [show analytics status](#)
- [show analytics traffic-statistics](#)

## monitor start (Analytics)

**Syntax** `monitor start filename`

**Release Information** Command introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Start the display of the queue statistics or traffic statistics file if you had enabled queue or traffic monitoring on your device. The output is displayed in the JavaScript Object Notation (JSON) format.



**NOTE:** This topic describes the local file output in Junos OS Release 13.2X50-D15 and 13.2X51-D10 only. For information about 13.2X51-D15 and later, see [“Understanding Enhanced Analytics Local File Output” on page 3991](#)

**Options** *filename*—Name of the queue statistics or traffic statistics file.

**Required Privilege Level** trace

**Related Documentation**

- [Network Analytics Overview on page 3978](#)
- [analytics on page 4167](#)

**List of Sample Output**

[monitor start Using the Queue Statistics File \(Junos OS Release 13.2X51-D10\) on page 4344](#)  
[monitor start Using the Queue Statistics File \(Junos OS Release 13.2X50-D15\) on page 4345](#)  
[monitor start Using the Traffic Statistics File \(Junos OS Release 13.2X51-D10\) on page 4345](#)  
[monitor start Using the Traffic Statistics File \(Junos OS Release 13.2X50-D15\) on page 4345](#)

**Output Fields** [Table 452 on page 4343](#) describes the output fields for the **monitor start** command. Output fields are listed in the approximate order in which they appear.

**Table 452: monitor start Command Output Fields**

| Field                                                | Description                                                        |
|------------------------------------------------------|--------------------------------------------------------------------|
| hostname (used in Junos OS Release 13.2X50-D15 only) | Name of the network analytics host device.                         |
| record type                                          | Type of statistics. May be queue statistics or traffic statistics. |
| time                                                 | Time at which the statistics were captured.                        |
| router-id                                            | ID of the network analytics host device.                           |

Table 452: monitor start Command Output Fields (*continued*)

| Field                                         | Description                                                                                        |
|-----------------------------------------------|----------------------------------------------------------------------------------------------------|
| latency                                       | For queue statistics only. Traffic queue latency in milliseconds.                                  |
| port                                          | Name of the physical port configured for network analytics.                                        |
| queue depth                                   | For queue statistics only. Depth of the traffic queue in bytes.                                    |
| rxpkt                                         | For traffic statistics monitoring only. Total packets received.                                    |
| rxpps                                         | For traffic statistics monitoring only. Total packets received per second.                         |
| rxbyte                                        | For traffic statistics monitoring only. Total bytes received.                                      |
| rxbps                                         | For traffic statistics monitoring only. Total bytes received per second.                           |
| rxdrop                                        | For traffic statistics monitoring only. Total incoming packets dropped.                            |
| rxerr                                         | For traffic statistics monitoring only. Total packets with errors.                                 |
| rxutil (in Junos OS Release 13.2X50-D15 only) | For traffic statistics monitoring only. Total percent of traffic utilization for incoming traffic. |
| txpkt                                         | For traffic statistics monitoring only. Total packets transmitted.                                 |
| txpps                                         | For traffic statistics monitoring only. Total packets transmitted per second.                      |
| txbyte                                        | For traffic statistics monitoring only. Total bytes transmitted.                                   |
| txbps                                         | For traffic statistics monitoring only. Total bytes transmitted per second.                        |
| txdrop                                        | For traffic statistics monitoring only. Total transmitted bytes dropped.                           |
| txerr                                         | For traffic statistics monitoring only. Total transmitted packets with errors (dropped).           |
| txutil (in Junos OS Release 13.2X50-D15 only) | For traffic statistics monitoring only. Total percent of traffic utilization for outgoing traffic. |

## Sample Output

### monitor start Using the Queue Statistics File (Junos OS Release 13.2X51-D10)

```

user@host> monitor start analytics.qs
{"record-type":"queue-stats","time":"2013 Nov 3 4:40:42.840",
"router-id":"qfx5100-switch","port":"xe-0/0/18","latency":0,"queue-depth":208}

{"record-type":"queue-stats","time":"2013 Nov 3 4:40:44.887",
"router-id":"qfx5100-switch","port":"xe-0/0/18","latency": 1110,"queue-depth":
1387568}

```



### monitor start Using the Queue Statistics File (Junos OS Release 13.2X50-D15)

```
user@host> monitor start analytics.qs
{"hostname":"sw-la-pb-03","latency":566,"port":"xe-0/0/9","queue depth":708656,
"record type":"queue-stats","time":"Apr 11 20:18:40.329"}
```

### Sample Output

#### monitor start Using the Traffic Statistics File (Junos OS Release 13.2X51-D10)

```
user@host> monitor start analytics.ts
{"record-type":"traffic-stats","time":"2013 Nov 3 4:39:53.910",
"router-id":"qfx5100-switch","port":"xe-0/0/18","rxpkt":23193749091,"rxpps":8299889,

"rxbyte":2968799876957,"rxbps":824002992,"rxdrop":0,"rxerr":0,"txpkt":1029323986,
"txpps":82671,"txbyte":131753470470,"txbps":85598256,"txdrop":0,"txerr":0}
```

#### monitor start Using the Traffic Statistics File (Junos OS Release 13.2X50-D15)

```
user@host> monitor start analytics.ts
{"hostname":"sw-la-pb-03","port":"xe-0/0/9","record type":"traffic-statistics",
"time":"Apr 11 20:13:48.545", "rxpkt":601024640, "rxpps": 840315,
"rxbyte":76931153920,
"rxbps":863997032, "rxdrop":0, "rxerr":0, "rxutil":8.32,"txpkt":336551380309,
"txpps":405395,"txbyte":23369872265951,"txbps":3240000976,"txdrop":1010566660824,
"txerr":69920099883860,"txutil":32.76}
```

## show analytics collector

**Syntax** show analytics collector

**Release Information** Command introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Show the list of network analytics remote collectors and related information. Remote collectors can be configured to receive streaming output for queue statistics and traffic statistics from the network analytics process (Analyticsd) running on the Routing Engine.



**NOTE:** The show analytics collector command is not available in Junos OS Releases prior to 13.2X51-D15.

**Required Privilege Level** interface-control

**Related Documentation**

- [Network Analytics Overview on page 3978](#)
- [analytics on page 4167](#)
- [address \(Analytics Collector\) on page 4166](#)

**List of Sample Output** [show analytics collector on page 4347](#)

**Output Fields** [Table 453 on page 4346](#) describes the output fields for the **show analytics collector** command.

**Table 453: show analytics collector Command Output Fields**

| Field         | Description                                                                                                                                                                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Address       | IP Address of the collector that is configured for receiving the streaming data.                                                                                                                                                                                                        |
| Port          | Port number of the collector receiving the streaming data.                                                                                                                                                                                                                              |
| Transport     | Transport protocol: <ul style="list-style-type: none"> <li>• tcp—Transmission Control Protocol</li> <li>• udp—User Datagram Protocol</li> </ul> <p><b>NOTE:</b> The connection state of a port configured with the <b>udp</b> transport protocol is always displayed as <b>n/a</b>.</p> |
| Stream format | Format of the data that is sent to the server: <ul style="list-style-type: none"> <li>• csv—Comma-separated values</li> <li>• gpb—Google Protocol Buffer</li> <li>• json—JavaScript Object Notation</li> <li>• tsv—Tab-separated values</li> </ul>                                      |

Table 453: show analytics collector Command Output Fields (*continued*)

| Field | Description                                   |
|-------|-----------------------------------------------|
| State | Connection state of the streaming server.     |
| Sent  | Number of bytes sent to the streaming server. |

## Sample Output

### show analytics collector

```

user@host> show analytics collector
Address      Port    Transport Stream format State      Sent
10.94.184.25 50013   udp       gpb        n/a       8710
10.94.184.25 50040   tcp       gpb        Not initialized 0
10.94.184.25 50050   tcp       gpb        Established 405
10.94.184.62 50010   tcp       csv        Established 18
10.94.184.62 50020   udp       json       n/a       17

```

## show analytics configuration

|                                 |                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show analytics configuration                                                                                                                                                                               |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.                                                            |
| <b>Description</b>              | Show the network analytics configuration details for the global and interface configurations.                                                                                                              |
| <b>Required Privilege Level</b> | interface-control                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Network Analytics Overview on page 3978</a></li> <li>• <a href="#">analytics on page 4167</a></li> </ul>                                              |
| <b>List of Sample Output</b>    | <a href="#">show analytics configuration (Junos OS Release 13.2X51-D15 and Later) on page 4351</a><br><a href="#">show analytics configuration (Junos OS Release 13.2X51-D10 and Earlier) on page 4351</a> |
| <b>Output Fields</b>            | describes the output fields for the <b>show analytics configuration</b> command in Junos OS Release 13.2X51-D15 and later.                                                                                 |

**Table 454: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D15 and Later)**

| Field                               | Descriptions                                                                                                                                                                                                                                        |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Global Configurations</b>        |                                                                                                                                                                                                                                                     |
| Traffic monitoring status           | Settings are enabled or disabled. If traffic statistics monitoring is not enabled, this field is not shown.                                                                                                                                         |
| Traffic monitoring polling interval | Interval for traffic statistics polling in seconds.<br><br><b>NOTE:</b> Due to limitations and variations in hardware capability in different devices, there might be a difference in value between the actual interval and configured interval.    |
| Queue monitoring status             | Settings are enabled or disabled. If queue statistics monitoring is not enabled, this field is not shown.                                                                                                                                           |
| Queue monitoring polling interval   | Interval for queue statistics polling in milliseconds.<br><br><b>NOTE:</b> Due to limitations and variations in hardware capability in different devices, there might be a difference in value between the actual interval and configured interval. |
| Queue depth high threshold          | Upper limit of the depth threshold configuration in number of bytes.<br><br>If the queue depth threshold is not configured, this field is not shown.                                                                                                |
| Queue depth low threshold           | Lower limit of the depth threshold configuration in number of bytes.<br><br>If the queue depth threshold is not configured, this field is not shown.                                                                                                |

**Table 454: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D15 and Later) (continued)**

| Field                           | Descriptions                                                                                                                                          |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queue latency high threshold    | Upper limit of the latency threshold configuration in nanoseconds.<br><br>If the queue latency threshold is not configured, this field is not shown.  |
| Queue latency low threshold     | Lower limit of the latency threshold configuration in microseconds.<br><br>If the queue latency threshold is not configured, this field is not shown. |
| <b>Interface Configurations</b> |                                                                                                                                                       |
| Interface                       | Name of interface that is configured for network analytics. The interface configuration overrides the global network analytics configuration.         |
| Traffic Statistics              | Settings are Enabled or Disabled for the interface.                                                                                                   |
| Queue Statistics                | Settings are Enabled or Disabled for the interface.                                                                                                   |
| Queue depth threshold High      | Upper limit of the depth threshold configuration in number of bytes.<br><br>If the queue depth threshold is not configured, <b>n/a</b> is displayed.  |
| Queue depth threshold Low       | Lower limit of the depth threshold configuration in number of bytes.<br><br>If the queue depth threshold is not configured, <b>n/a</b> is displayed.  |
| Latency threshold High          | Upper limit of the latency threshold configuration in nanoseconds.<br><br>If the latency threshold is not configured, <b>n/a</b> is displayed.        |
| Latency threshold Low           | Lower limit of the latency threshold configuration in nanoseconds.<br><br>If the latency threshold is not configured, <b>n/a</b> is displayed.        |

[Table 455 on page 4349](#) describes the output fields for the **show analytics configuration** command in Junos OS Release 13.2X51-D10 and 13.2X50-D15.

**Table 455: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D10 and earlier)**

| Field                        | Descriptions                                                                                                               |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Global Configurations</b> |                                                                                                                            |
| Traffic statistics           | Settings are Auto, Enabled, or Disabled.<br><br>If <b>Auto</b> is displayed, traffic statistics monitoring is not enabled. |

**Table 455: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D10 and earlier) (continued)**

| Field                              | Descriptions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Poll interval (traffic statistics) | <p>Interval for traffic statistics polling in seconds.</p> <p>If the output displays a setting of 0 seconds, the polling interval was not configured, and the default interval applies.</p> <p><b>NOTE:</b> The default interval is 1 second in Junos OS Release 13.2X50-D15 and later, except for EX4300 switches, on which the default interval is 5 seconds, and 2 seconds in Junos OS Release 13.2X51-D10.</p> <p><b>NOTE:</b> Due to limitations and variations in hardware capability in different devices, there might be a difference in value between the actual interval and configured interval.</p> |
| Queue statistics                   | <p>Settings are Auto, Enabled, or Disabled.</p> <p>If <b>Auto</b> is displayed, queue statistics monitoring is not enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Poll interval (queue statistics)   | <p>Interval for queue statistics polling in milliseconds.</p> <p><b>NOTE:</b> The default interval is 8 milliseconds in Junos OS Release 13.2X50-D15 and later, and 10 milliseconds in Junos OS Release 13.2X51-D10 or later.</p> <p><b>NOTE:</b> Due to limitations and variations in hardware capability in different devices, there might be a difference in value between the actual interval and configured interval.</p>                                                                                                                                                                                  |
| Depth threshold high               | <p>Upper limit of the depth threshold configuration in number of bytes.</p> <p>If <b>0</b> is displayed, depth threshold is not enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Depth threshold low                | <p>Lower limit of the depth threshold configuration in number of bytes.</p> <p>If <b>0</b> is displayed, depth threshold is not enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Latency threshold high             | <p>Upper limit of the latency threshold configuration in microseconds.</p> <p>If <b>0</b> is displayed, latency threshold is not enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Latency threshold low              | <p>Lower limit of the latency threshold configuration in microseconds.</p> <p>If <b>0</b> is displayed, latency threshold is not enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Interface Configurations</b>    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Interface                          | Name of interface that is configured for network analytics. The interface configuration overrides the global network analytics configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Traffic Statistics                 | Settings are Enabled or Disabled for the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Queue Statistics                   | Settings are Enabled or Disabled for the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Depth-threshold High               | <p>Upper limit of the depth threshold configuration in number of bytes.</p> <p>If <b>0</b> is displayed, depth threshold is not enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

**Table 455: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D10 and earlier) (continued)**

| Field                  | Descriptions                                                                                                                    |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| Depth-threshold Low    | Lower limit of the depth threshold configuration in number of bytes.<br><br>If 0 is displayed, depth threshold is not enabled.  |
| Latency-threshold High | Upper limit of the latency threshold configuration in microseconds.<br><br>If 0 is displayed, latency threshold is not enabled. |
| Latency-threshold Low  | Lower limit of the latency threshold configuration in microseconds.<br><br>If 0 is displayed, latency threshold is not enabled. |

## Sample Output

### show analytics configuration (Junos OS Release 13.2X51-D15 and Later)

```

user@host> show analytics configuration
Traffic monitoring status is enabled
Traffic monitoring polling interval : 5 seconds
Queue monitoring status is enabled
Queue monitoring polling interval : 1000 milliseconds
Queue depth high threshold : 99999 bytes
Queue depth low threshold : 99 bytes

```

| Interface | Traffic<br>Statistics | Queue<br>Statistics | Queue depth<br>threshold |     | Latency<br>threshold  |     |
|-----------|-----------------------|---------------------|--------------------------|-----|-----------------------|-----|
|           |                       |                     | High<br>(bytes)          | Low | High<br>(nanoseconds) | Low |
| xe-0/0/16 | enabled               | enabled             | n/a                      | n/a | 2300                  | 20  |
| xe-0/0/18 | enabled               | enabled             | n/a                      | n/a | 2300                  | 20  |
| xe-0/0/19 | enabled               | enabled             | n/a                      | n/a | 2300                  | 20  |

### show analytics configuration (Junos OS Release 13.2X51-D10 and Earlier)

```

user@host> show analytics configuration
Global configurations:
  Traffic statistics: Enabled, Poll interval: 2 seconds
  Queue statistics: Auto, Poll interval: 10 milliseconds
  Depth threshold high: 0 bytes, low: 0 bytes
  Latency threshold high: 0 microseconds, low: 0 microseconds

```

| Interface | Traffic<br>Statistics | Queue<br>Statistics | Depth-threshold |     | Latency-threshold      |     |
|-----------|-----------------------|---------------------|-----------------|-----|------------------------|-----|
|           |                       |                     | High<br>(bytes) | Low | High<br>(microseconds) | Low |
| xe-0/0/0  | Auto                  | Auto                | 204800          | 10  | 0                      | 0   |

## show analytics queue-statistics

|                                 |                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show analytics queue-statistics</code><br><code>&lt;interface <i>interface-name</i>&gt;</code>                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.                                                                                                                                                     |
| <b>Description</b>              | Show the queue statistics (queue length and latency) that are collected for all interfaces that are enabled for network analytics on a device. Optionally, if you wish to see the queue statistics for one interface only, you may specify the interface.                                           |
| <b>Options</b>                  | <code>interface <i>interface-name</i></code> —(Optional) Display the queue statistics for the specified interface only.                                                                                                                                                                             |
| <b>Required Privilege Level</b> | interface-control                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Network Analytics Overview on page 3978</a></li> <li><a href="#">analytics on page 4167</a></li> </ul>                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">show analytics queue-statistics (Junos OS Release 13.2X51-D15 and Later) on page 4352</a><br><a href="#">show analytics queue-statistics (Junos OS Release 13.2X51-D10) on page 4353</a><br><a href="#">show analytics queue-statistics (Junos OS Release 13.2X50-D15) on page 4353</a> |
| <b>Output Fields</b>            | <a href="#">Table 425 on page 3991</a> describes the output fields for the <code>show analytics queue-statistics</code> command.                                                                                                                                                                    |

Table 456: show analytics queue-statistics Command Output Fields

| Field                               | Description                                                                                                                     |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| Time                                | Date and time at which the queue statistics are collected.                                                                      |
| Interface                           | Name of the interface at which the queue statistics are collected.                                                              |
| Queue-length or queue-depth (bytes) | Queue depth (length) in number of bytes.                                                                                        |
| Latency                             | Queue depth in nanoseconds (Junos OS Release 13.2X51-D15 and later) or microseconds (Junos OS Release 13.2X51-D10 and earlier). |

## Sample Output

### show analytics queue-statistics (Junos OS Release 13.2X51-D15 and Later)

```

user@host> show analytics queue-statistics
CLI issued at 2014-01-07 17:20:29.978561
Time                Interface      Queue-depth      Latency
                    (bytes)          (nanoseconds)
00:00:00.870058 ago  xe-0/0/19      1369680          1095744

```



|                     |           |         |         |
|---------------------|-----------|---------|---------|
| 00:00:01.875049 ago | xe-0/0/19 | 1381952 | 1105561 |
| 00:00:02.875053 ago | xe-0/0/19 | 1387776 | 1110220 |
| 00:00:03.876047 ago | xe-0/0/19 | 1387568 | 1110054 |
| 00:00:04.873045 ago | xe-0/0/19 | 1388192 | 1110553 |
| 00:00:05.871044 ago | xe-0/0/19 | 1385904 | 1108723 |
| 00:00:06.873354 ago | xe-0/0/19 | 1371552 | 1097241 |

#### show analytics queue-statistics (Junos OS Release 13.2X51-D10)

```
user@host> show analytics queue-statistics
```

| Time                   | Interface | Queue-length (bytes) | Latency (us) |
|------------------------|-----------|----------------------|--------------|
| 2013 Nov 3 3:52:26.272 | xe-0/0/9  | 208                  | 0            |
| 2013 Nov 3 3:52:26.292 | xe-0/0/9  | 208                  | 0            |
| 2013 Nov 3 3:52:26.372 | xe-0/0/9  | 208                  | 0            |
| 2013 Nov 3 3:52:26.392 | xe-0/0/9  | 208                  | 0            |
| 2013 Nov 3 3:52:26.432 | xe-0/0/9  | 208                  | 0            |
| 2013 Nov 3 3:52:26.492 | xe-0/0/9  | 208                  | 0            |
| 2013 Nov 3 3:52:26.572 | xe-0/0/9  | 208                  | 0            |
| 2013 Nov 3 4:30:24.584 | xe-0/0/9  | 1387152              | 1109         |
| 2013 Nov 3 4:30:24.604 | xe-0/0/9  | 1372384              | 1097         |
| 2013 Nov 3 4:30:24.624 | xe-0/0/9  | 1384864              | 1107         |


## Sample Output

#### show analytics queue-statistics (Junos OS Release 13.2X50-D15)

```
user@host> show analytics queue-statistics
```

| Time              | Interface | Queue-length (bytes) | Latency (us) |
|-------------------|-----------|----------------------|--------------|
| Apr 6 0:17:18.224 | xe-0/0/9  | 1043952              | 835          |
| Apr 6 0:17:18.234 | xe-0/0/9  | 1053520              | 842          |
| Apr 6 0:17:18.244 | xe-0/0/9  | 1055184              | 844          |

## show analytics status

|                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                           | show analytics status<br><global>                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>                                                                                                                                                                              | Command introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                      | Show the status of the network analytics components that are configured on a device.                                                                                                                                                                                                                 |
| <b>Options</b>                                                                                                                                                                                          | <b>none</b> —Show the global and interface status for network analytics.<br><br><b>global</b> —Show the global status only for network analytics.                                                                                                                                                    |
| <div>  <p><b>NOTE:</b> The <b>global</b> option is not available in Junos OS Releases prior to 13.2X51-D15.</p> </div> |                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                                         | interface-control                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>                                                                                                                                                                            | <ul style="list-style-type: none"> <li>• <a href="#">Network Analytics Overview on page 3978</a></li> <li>• <a href="#">analytics on page 4167</a></li> </ul>                                                                                                                                        |
| <b>List of Sample Output</b>                                                                                                                                                                            | <a href="#">show analytics status (Junos OS Release 13.2X51-D15 or Later) on page 4356</a><br><a href="#">show analytics status global (Junos OS Release 13.2X51-D15 or Later) on page 4356</a><br><a href="#">show analytics status (Junos OS Release 13.2X50-D15 and 13.2X51-D10) on page 4356</a> |
| <b>Output Fields</b>                                                                                                                                                                                    | Table 457 on page 4354 describes the output fields for the <b>show analytics status</b> command.                                                                                                                                                                                                     |

**Table 457: show analytics status Command Output Fields**

| Field                                                | Descriptions                                                                                                                                                                                                                                               |
|------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Global Configurations</b>                         |                                                                                                                                                                                                                                                            |
| Traffic statistics or Traffic monitoring status      | <p>Settings are Auto, Enabled, or Disabled.</p> <p>If <b>Auto</b> is displayed, traffic statistics monitoring is not enabled.</p> <p><b>NOTE:</b> The Disabled setting always supersedes the Enabled setting.</p>                                          |
| Poll interval or Traffic monitoring polling interval | <p>Interval for traffic statistics polling in seconds.</p> <p><b>NOTE:</b> Due to limitations and variations in the hardware capability of different devices, you might see a difference in value between the actual interval and configured interval.</p> |

Table 457: show analytics status Command Output Fields (*continued*)

| Field                                              | Descriptions                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queue statistics or Queue monitoring status        | <p>Can be Auto, Enabled, or Disabled.</p> <p>If <b>Auto</b> is displayed, queue statistics monitoring is not enabled.</p> <p><b>NOTE:</b> The Disabled setting always supersedes the Enabled setting.</p>                                                                                                            |
| Poll interval or Queue monitoring polling interval | <p>Interval for queue statistics polling in milliseconds.</p> <p><b>NOTE:</b> Due to limitations and variations in the hardware capability of different devices, you might see a difference in value between the actual interval and configured interval.</p>                                                        |
| Depth threshold high or Queue depth high threshold | <p>Upper limit of the depth threshold configuration in number of bytes.</p> <p>In Junos OS Release 13.2X51-D15 or later, if this parameter is not configured, this field is not shown.</p> <p>In Junos OS Release 13.2X51-D10 or earlier, if this parameter is not configured, a value of <b>0</b> is displayed.</p> |
| Depth threshold low or Queue depth low threshold   | <p>Lower limit of the depth threshold configuration in number of bytes.</p> <p>In Junos OS Release 13.2X51-D15 or later, if this parameter is not configured, this field is not shown.</p> <p>In Junos OS Release 13.2X51-D10 or earlier, if this parameter is not configured, a value of <b>0</b> is displayed.</p> |
| Latency threshold high                             | <p>Upper limit of the latency threshold configuration in microseconds.</p> <p>In Junos OS Release 13.2X51-D15 or later, if this parameter is not configured, this field is not shown.</p> <p>In Junos OS Release 13.2X51-D10 or earlier, if this parameter is not configured, a value of <b>0</b> is displayed.</p>  |
| Latency threshold low                              | <p>Lower limit of the latency threshold configuration in microseconds.</p> <p>In Junos OS Release 13.2X51-D15 or later, if this parameter is not configured, this field is not shown.</p> <p>In Junos OS Release 13.2X51-D10 or earlier, if this parameter is not configured, a value of <b>0</b> is displayed.</p>  |
| <b>Interface Configurations</b>                    |                                                                                                                                                                                                                                                                                                                      |
| Interface                                          | Name of an interface that is configured for network analytics. The interface configuration overrides the global network analytics configuration.                                                                                                                                                                     |
| Traffic Statistics                                 | <p>Settings are Enabled or Disabled for the interface.</p> <p><b>NOTE:</b> The Disabled setting always supersedes the Enabled setting.</p>                                                                                                                                                                           |
| Queue Statistics                                   | <p>Settings are Enabled or Disabled for the interface.</p> <p><b>NOTE:</b> The Disabled setting always supersedes the Enabled setting.</p>                                                                                                                                                                           |

Table 457: show analytics status Command Output Fields (*continued*)

| Field                                              | Descriptions                                                                                                                                                                                                                       |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Depth-threshold High or Queue depth threshold high | Upper limit of the depth threshold configuration in number of bytes.<br><br>If this parameter is not configured, an output of <b>n/a</b> or <b>0</b> is displayed in this column, depending on the software release.               |
| Depth-threshold Low or Queue depth threshold low   | Lower limit of the depth threshold configuration in number of bytes.<br><br>If this parameter is not configured, an output of <b>n/a</b> or <b>0</b> is displayed in this column, depending on the software release.               |
| Latency-threshold High                             | Upper limit of the latency threshold configuration in nanoseconds or microseconds.<br><br>If this parameter is not configured, an output of <b>n/a</b> or <b>0</b> is displayed in this column, depending on the software release. |
| Latency-threshold Low                              | Lower limit of the latency threshold configuration in nanoseconds or microseconds.<br><br>If this parameter is not configured, an output of <b>n/a</b> or <b>0</b> is displayed in this column, depending on the software release. |

## Sample Output

### show analytics status (Junos OS Release 13.2X51-D15 or Later)

```

user@host> show analytics status
Traffic monitoring status is auto
Traffic monitoring polling interval : 5 seconds
Queue monitoring status is enabled
Queue monitoring status polling interval : 1000 milliseconds
Queue depth high threshold : 1000000000 bytes
Queue depth low threshold : 99 bytes

```

| Interface | Traffic<br>Statistics | Queue<br>Statistics | Queue depth<br>threshold |     | Latency<br>threshold |     |
|-----------|-----------------------|---------------------|--------------------------|-----|----------------------|-----|
|           |                       |                     | High                     | Low | High                 | Low |
| xe-0/0/16 | enabled               | enabled             | 1000000000               | 99  | n/a                  | n/a |
| xe-0/0/18 | disabled              | enabled             | 1000000000               | 99  | n/a                  | n/a |
| xe-0/0/19 | enabled               | enabled             | 1000000000               | 99  | n/a                  | n/a |

### show analytics status global (Junos OS Release 13.2X51-D15 or Later)

```

user@host> show analytics status global

Traffic monitoring status is auto
Traffic monitoring polling interval : 5 seconds
Queue monitoring status is enabled
Queue monitoring status polling interval : 1000 milliseconds
Queue depth high threshold : 1000000000 bytes
Queue depth low threshold : 99 bytes

```

### show analytics status (Junos OS Release 13.2X50-D15 and 13.2X51-D10)

```

user@host> show analytics status

```

## Global configurations:

Traffic statistics: Auto, Poll interval: 2 seconds

Queue statistics: Auto, Poll interval: 10 milliseconds

Depth threshold high: 0 bytes, low: 0 bytes

Latency threshold high: 1000 microseconds, low: 50 microseconds

| Interface | Traffic<br>Statistics | Queue<br>Statistics | Depth-threshold |     | Latency-threshold |     |
|-----------|-----------------------|---------------------|-----------------|-----|-------------------|-----|
|           |                       |                     | High            | Low | High              | Low |
|           |                       |                     | (bytes)         |     | (microseconds)    |     |
| xe-0/0/6  | Enabled               | Enabled             | 0               | 0   | 1000              | 50  |
| xe-0/0/7  | Enabled               | Enabled             | 204800          | 10  | 0                 | 0   |
| xe-0/0/8  | Enabled               | Enabled             | 0               | 0   | 1000              | 50  |

## show analytics traffic-statistics

|                                 |                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show analytics traffic-statistics</code><br><code>&lt;interface <i>interface-name</i>&gt;</code>                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.                                                                                                                                                          |
| <b>Description</b>              | Show the traffic statistics that are collected for all interfaces that are enabled for network analytics on a device. Optionally, if you wish to see the traffic statistics for one interface only, you may specify the interface.                                                                       |
| <b>Options</b>                  | <code>interface <i>interface-name</i></code> —(Optional) Display the traffic statistics for the specified interface only.                                                                                                                                                                                |
| <b>Required Privilege Level</b> | interface-control                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Network Analytics Overview on page 3978</a></li> <li><a href="#">analytics on page 4167</a></li> </ul>                                                                                                                                                |
| <b>List of Sample Output</b>    | <a href="#">show analytics traffic-statistics (Junos OS Release 13.2X51-D15 or Later) on page 4359</a><br><a href="#">show analytics traffic-statistics (Junos OS Release 13.2X51-D10) on page 4359</a><br><a href="#">show analytics traffic-statistics (Junos OS Release 13.2X50-D15) on page 4359</a> |
| <b>Output Fields</b>            | <a href="#">Table 458 on page 4358</a> describes the output fields for the <code>show analytics traffic-statistics</code> command.                                                                                                                                                                       |

**Table 458: show analytics traffic-statistics Command Output Fields**

| Field                              | Description                                                                                                                                                                                                                                                                           |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time                               | The date and time at which the traffic statistics are generated.                                                                                                                                                                                                                      |
| Physical interface                 | Name of the interface at which the traffic statistics are collected.                                                                                                                                                                                                                  |
| Total octets                       | Total number of octets that are received and transmitted.                                                                                                                                                                                                                             |
| Total packets                      | Total number of packets that are received and transmitted.                                                                                                                                                                                                                            |
| Octets per second                  | Number of octets received and transmitted per second.                                                                                                                                                                                                                                 |
| Packet per second                  | Number of packets received and transmitted per second.                                                                                                                                                                                                                                |
| CRC/Align errors or Octets dropped | Number of cyclic redundancy check (CRC) errors or octets dropped. <ul style="list-style-type: none"> <li>Junos OS Release 13.2X51-D15 or later—Number of cyclic redundancy check (CRC) errors.</li> <li>Junos OS Release 13.2X51-D10 and earlier—Number of octets dropped.</li> </ul> |
| Packets dropped                    | Number of packets dropped.                                                                                                                                                                                                                                                            |

## Sample Output

### show analytics traffic-statistics (Junos OS Release 13.2X51-D15 or Later)

```

user@host> show analytics traffic-statistics
CLI issued at 2014-01-07 17:22:28.952677
Time: 00:00:03.480244 ago, Physical interface: xe-0/0/19
Traffic Statistics:
Total octets:          3929946593792      393001011519232
Total packets:         30702707784       3070320402462
Unicast packet:        30702707784       3070320402462
Multicast packets:     0                 0
Broadcast packets:     0                 0
Octets per second:     86407016          59044064
Packets per second:    84787             8469688
CRC/Align errors:      0                 392986110751744
Packets dropped:       0                 3070203990248

```

### show analytics traffic-statistics (Junos OS Release 13.2X51-D10)

```

user@host> show analytics traffic-statistics
Time: 2013 Nov 3 4:36:55.542, Physical interface: xe-0/0/8
Traffic Statistics:
Total octets:          2777524779008      101855533467
Total packet:          21699412289       795746503
Octets per second:     904001272         0
Packet per second:     8399574          0
Octets dropped:         0                 0
Packet dropped:        0                 0
Time: 2013 Nov 3 4:36:57.559, Physical interface: xe-0/0/10
Traffic Statistics:
Total octets:          2777546444381      129840936198
Total packet:          21699581650       1014382311
Octets per second:     90400211         86403728
Packet per second:     8400382          84438
Octets dropped:         0                 0
Packet dropped:        0                 0

```

### show analytics traffic-statistics (Junos OS Release 13.2X50-D15)

```

user@host> show analytics traffic-statistics
Time: Apr 5 19:52:48.549, Physical interface: xe-0/0/8
Traffic Statistics:
Total octets:          4797548752936      408886273632
Total packet:          5658257464       3190613435
Octets per second:     0                 0
Packet per second:     0                 0
Octets dropped:         0                 252901000
Packet dropped:         0                 252901
Utilization:           0.0%              0.0%
Time: Apr 5 19:52:48.549, Physical interface: xe-0/0/10
Traffic Statistics:
Total octets:          4790866253100      477139024
Total packet:          5624473639       477944
Octets per second:     0                 0
Packet per second:     0                 0
Octets dropped:         0                 166582000
Packet dropped:         0                 166582
Utilization:           0.0%              0.0%

```

## Operational Commands: sFlow

---

- [show sflow](#)
- [show sflow interface](#)
- [show sflow collector](#)
- [clear sflow collectors statistics](#)



## show sflow

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show sflow<br><collector><br><interface>                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Display default sFlow technology configuration information.                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>none</b>—Display default sFlow technology configuration information.</p> <p><b>collector</b>—(Optional) Display standard status information about the specified sFlow collector.</p> <p><b>interface</b>—(Optional) Display standard status information about the specified sFlow interface.</p>                                                                                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show sflow interface on page 4363</a></li> <li>• <a href="#">show sflow collector on page 4365</a></li> <li>• <a href="#">Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995</a></li> <li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show sflow on page 4362</a>                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 459 on page 4361</a> lists the output fields for the <b>show sflow</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                     |

**Table 459: show sflow Output Fields**

| Field Name          | Field Description                                                                                                   | Level of Output |
|---------------------|---------------------------------------------------------------------------------------------------------------------|-----------------|
| sFlow               | Status of the feature: <b>enabled</b> or <b>disabled</b> .                                                          | All levels      |
| Sample rate egress  | Rate at which egress packets are sampled.                                                                           | All levels      |
| Sample rate ingress | Rate at which ingress packets are sampled.                                                                          | All levels      |
| Sample limit        | Number of packets sampled per second. The sampling limit cannot be configured and is set to 300 packets per second. | All levels      |
| Polling interval    | Interval at which the sFlow agent polls the interface.                                                              | All levels      |
| Agent ID            | The IP address assigned to the sFlow agent.                                                                         | All levels      |

Table 459: show sflow Output Fields (*continued*)

| Field Name        | Field Description                      | Level of Output |
|-------------------|----------------------------------------|-----------------|
| Source IP address | The IP address for the sFlow datagram. | All levels      |

## Sample Output

show sflow

```
sFlow          : Enabled
Sample rate egress : 1:1000
Sample rate ingress : 1: 2048: Disabled
Sample limit      : 300 packets/second
Polling interval : 20 seconds
Agent ID         : 10.93.54.7
Source IP address : 10.93.54.7
```

## show sflow interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show sflow interface                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display the interfaces on which sFlow technology is enabled and the sampling parameters.                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show sflow on page 4361</a></li> <li>• <a href="#">show sflow collector on page 4365</a></li> <li>• <a href="#">Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995</a></li> <li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show sflow interface on page 4363</a>                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | Table 460 on page 4363 lists the output fields for the <b>show sflow interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                 |

**Table 460: show sflow interface Output Fields**

| Field Name                         | Field Description                                          | Level of Output |
|------------------------------------|------------------------------------------------------------|-----------------|
| <b>Interface</b>                   | Interfaces on which sFlow technology is enabled.           | All levels      |
| <b>Status Egress</b>               | Indicates whether egress sampling rate is enabled.         | All levels      |
| <b>Status Ingress</b>              | Indicates whether ingress sampling rate is enabled.        | All levels      |
| <b>Sample rate Egress</b>          | Rate at which egress packets are sampled.                  | All levels      |
| <b>Sample rate Ingress</b>         | Rate at which ingress packets are sampled.                 | All levels      |
| <b>Adapted sample rate Egress</b>  | Adapted rate at which egress packets are sampled.          | All levels      |
| <b>Adapted sample rate Ingress</b> | Adapted rate at which ingress packets are sampled.         | All levels      |
| <b>Polling-interval</b>            | The interval at which the sFlow agent polls the interface. | All levels      |

## Sample Output

### show sflow interface

```
Interface          Status      Sample rate  Adapted sample rate  Polling-interval
```

|            |                   |                     |                |                 |                |                 |    |
|------------|-------------------|---------------------|----------------|-----------------|----------------|-----------------|----|
| ge-0/0/0.0 | Egress<br>Enabled | Ingress<br>Disabled | Egress<br>1000 | Ingress<br>2048 | Egress<br>1000 | Ingress<br>2048 | 20 |
|------------|-------------------|---------------------|----------------|-----------------|----------------|-----------------|----|

## show sflow collector

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show sflow collector                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display a list of configured sFlow collectors and their properties.                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show sflow on page 4361</a></li> <li>• <a href="#">show sflow interface on page 4363</a></li> <li>• <a href="#">Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 3995</a></li> <li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show sflow collector on page 4365</a>                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 461 on page 4365</a> lists the output fields for the <b>show sflow collector</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                 |

**Table 461: show sflow collector Output Fields**

| Field Name    | Field Description                    | Level of Output |
|---------------|--------------------------------------|-----------------|
| IP address    | IP address of the collector.         | All levels      |
| UDP port      | UDP port number.                     | All levels      |
| No of samples | Number of samples sent to collector. | All levels      |

## Sample Output

### show sflow collector

```

IP-address    UDP-Port    No of samples
10.204.32.46  5600        1000
100.204.32.76 3400        1000

```

## clear sflow collectors statistics

---

|                                 |                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear sflow collectors statistics                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced in JUNOS Release 9.5 for EX Series switches.                                                                                                                                                |
| <b>Description</b>              | Clear the sFlow collector's statistics.                                                                                                                                                                        |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show sflow collector on page 4365</a></li><li>• <a href="#">Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 4049</a></li></ul> |

## Sample Output

```
clear sflow collectors statistics
```

## Operational Commands: Ethernet OAM Connectivity Fault Management

---

- [clear oam ethernet connectivity-fault-management delay-statistics](#)
- [clear oam ethernet connectivity-fault-management sla-iterator-statistics](#)
- [clear oam ethernet connectivity-fault-management statistics](#)
- [monitor ethernet delay-measurement](#)
- [show oam ethernet connectivity-fault-management delay-statistics](#)
- [show oam ethernet connectivity-fault-management forwarding-state](#)
- [show oam ethernet connectivity-fault-management interfaces](#)
- [show oam ethernet connectivity-fault-management path-database](#)
- [show oam ethernet connectivity-fault-management mep-database](#)
- [show oam ethernet connectivity-fault-management mip](#)
- [show oam ethernet connectivity-fault-management sla-iterator-statistics](#)

## clear oam ethernet connectivity-fault-management delay-statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>clear oam ethernet connectivity-fault-management delay-statistics</b><br><b>maintenance-association</b> <i>maintenance-association-name</i><br><b>maintenance-domain</b> <i>maintenance-domain-name</i><br><b>&lt;logical-system</b> <i>logical-system-name</i> <b>&gt;</b><br><b>&lt;one-way&gt;</b><br><b>&lt;two-way&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.6.<br>Command introduced in Junos OS Release 11.4 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | On MX Series routers and EX Series switches, clear ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM) delay statistics and ETH-DM frame counts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b>maintenance-association</b> <i>maintenance-association-name</i> —Clear ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance association.<br><br><b>maintenance-domain</b> <i>maintenance-domain-name</i> —Clear ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance domain.<br><br><b>logical-system</b> <i>logical-system-name</i> —(MX Series routers only) (Optional) Clear ETH-DM delay statistics and ETH-DM frame counts for the specified logical system.<br><br><b>one-way</b> —(Optional) Clear one-way ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance association, maintenance domain, or (on the routers only) logical system.<br><br><b>two-way</b> —(Optional) Clear two-way ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance association, maintenance domain, or (on the routers only) logical system. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear oam ethernet connectivity-fault-management statistics</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management delay-statistics on page 4376</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management interfaces</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>    | <a href="#">clear oam ethernet connectivity-fault-management delay statistics on page 4367</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

### Sample Output

#### clear oam ethernet connectivity-fault-management delay statistics

```
user@switch> clear oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md1 maintenance-association mal
Delay statistics entries cleared
```





## clear oam ethernet connectivity-fault-management sla-iterator-statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear oam ethernet connectivity-fault-management sla-iterator-statistics<br>maintenance-association <i>maintenance-association-name</i><br>maintenance-domain <i>maintenance-domain-name</i><br><local-mep <i>local-mep-id</i> ><br><remote-mep <i>remote-mep-id</i> ><br>sla-iterator <i>sla-iterator</i>                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2 for MX Series routers.                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Clear Ethernet Operation, Administration, and Maintenance (OAM) service-level agreement (SLA) iterator statistics. For MX Series routers, clear the SLA iterator statistics and proactive Ethernet synthetic loss measurement (ETH-SLM) statistics.                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>maintenance-association</b> <i>maintenance-association-name</i>—Name of the maintenance association.</p> <p><b>maintenance-domain</b> <i>maintenance-domain-name</i>—Name of the maintenance domain.</p> <p><b>local-mep</b> <i>local-mep-id</i>—(Optional) Identifier of the local MEP.</p> <p><b>remote-mep</b> <i>remote-mep-id</i>—(Optional) Identifier of the remote MEP.</p> <p><b>sla-iterator</b> <i>sla-iterator</i>— Name of the SLA iterator profile.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring an Iterator Profile on a Switch (CLI Procedure) on page 4067</a></li> </ul>                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">clear oam ethernet connectivity-fault-management sla-iterator-statistics on page 4369</a>                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                       |

### Sample Output

clear oam ethernet connectivity-fault-  
-management sla-iterator- statistics

```
user@switch> clear oam ethernet connectivity-fault-management sla-iterator-statistics
maintenance-domain md1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator
i1
Iterator statistics entries cleared
```

## clear oam ethernet connectivity-fault-management statistics

---

|                                 |                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear oam ethernet connectivity-fault-management statistics</code><br><code>&lt;interface ethernet-interface-name&gt;</code><br><code>&lt;level md-level&gt;</code>                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Clear all statistics maintained by CFM.                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><code>interface ethernet-interface-name</code>—(Optional) Clear CFM statistics only for MEPs attached to the specified Ethernet physical interface.</p> <p><code>level level</code>—(Optional) Clear CFM statistics only for MEPs within CFM maintenance domains (MDs) of the specified level.</p>                                     |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show oam ethernet connectivity-fault-management interfaces on page 4384</a></li><li>• <a href="#">show oam ethernet connectivity-fault-management path-database on page 4390</a></li><li>• <a href="#">show oam ethernet connectivity-fault-management mip on page 4399</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">clear oam ethernet connectivity-fault-management statistics on page 4370</a>                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                     |

### Sample Output

clear oam ethernet connectivity-fault-management statistics

```
user@host> clear oam ethernet connectivity-fault-management statistics
Cleared statistics of all CFM sessions
```

## monitor ethernet delay-measurement

**Syntax** monitor ethernet delay-measurement maintenance-domain *md-name*  
 maintenance-association *ma-name* (one-way | two-way) (*remote-mac-address* | mep  
*remote-mep-id*) <count *count*> <no-session-id-tlv> <priority *802.1p value*> <size *size*>  
 <wait *time*>

**Release Information** Command introduced in Junos OS Release 11.4 for EX Series switches.

**Description** Start an ITU-T Y.1731 Ethernet frame delay measurement session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a summary of the frames exchanged in the measurement session. Frame delay measurement statistics are stored at one of the MEPs for later retrieval.



**NOTE:** If you attempt to monitor delays to a nonexistent MAC address, you must type Ctrl +C to explicitly quit the **monitor ethernet delay-measurement** command and return to the CLI command prompt.

To start an Ethernet frame delay measurement session, the switch initiates an exchange of frames carrying one-way or two-way frame delay measurement protocol data units (PDUs) between the local and remote MEPs. The frame counts—the types of and number of Ethernet frame delay measurement PDU frames exchanged to measure frame delay times—are displayed as the run-time output of the **monitor ethernet delay-measurement** command and are also stored at both the initiator and receiver MEPs for later retrieval. Ethernet frame delay measurement statistics, described below, are measured and stored at only one of the MEPs:

**Frame delay**—The difference, in microseconds, between the time a frame is sent and when it is received.

**Frame delay variation**—The difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.”

For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet frame delay measurement, only the initiator MEP (on the local system) collects statistics.

**Options** **count** *count*—(Optional) Number of frames to send to the specified peer MEP. The range of values is 1 through 65,535 frames. The default value is 10 frames.

**maintenance-association** *ma-name*—Name of an existing CFM maintenance association.

**maintenance-domain** *md-name*—Name of an existing CFM maintenance domain.

**mep** *remote-mep-id*—Numeric identifier of the peer MEP with which to perform Ethernet frame delay measurement. The discovered MAC address of the peer MEP is used. The range of values is 1 through 8191.

**no-session-id-tlv**—(Optional) Prevent insertion of the session ID TLV in the request frame.

**one-way**—Measurement type is one-way Ethernet frame delay measurement, which is based on the difference between the time at which the initiator MEP sends a one-way delay measurement request (IDM) frame and the time at which the receiver MEP receives the frame.

**priority 802.1p value**—(Optional) Priority of the delay measurement request frame supported by both one-way delay measurement and two-way delay measurement. The range of values is from 0 through 7. The default value is zero.

**remote-mac-address**—Unicast MAC address of the peer MEP with which to perform Ethernet frame delay measurement. Specify the MAC address as six hexadecimal bytes in *nn:nn:nn:nn:nn:nn* format. Multicast MAC addresses are not supported.

**size size** —(Optional) Size of the data TLV to be included in the request frame. The range of values is from 1 through 1400 bytes.

**two-way**—Measurement type is two-way Ethernet frame delay measurement, which is based on the difference between the time at which the initiator MEP sends a two-way delay measurement message (DMM) frame and the time at which the initiator MEP receives an associated two-way delay measurement reply (DMR) frame from the responder MEP, subtracting the time elapsed at the responder MEP.

**wait time**—(Optional) Number of seconds to wait between sending frames. The range of values is from 1 through 255 seconds. The default value is 1 second.

**Required Privilege Level** trace and maintenance

**Related Documentation**

- [Configuring an Iterator Profile on a Switch \(CLI Procedure\) on page 4067](#)
- [show oam ethernet connectivity-fault-management mep-database on page 4393](#)
- [show oam ethernet connectivity-fault-management mep-statistics](#)
- [show oam ethernet connectivity-fault-management delay-statistics on page 4376](#)
- [clear oam ethernet connectivity-fault-management statistics on page 4370](#)

**List of Sample Output**

- [monitor ethernet delay-measurement one-way on page 4374](#)
- [monitor ethernet delay-measurement two-way on page 4374](#)
- [monitor ethernet delay-measurement two-way \(Invalid DMR Frames Received\) on page 4374](#)

**Output Fields** The **monitor ethernet delay-measurement** command displays different output at the CLI, depending on whether you start a one-way or two-way frame delay measurement:

- [Table 462 on page 4373](#) lists the run-time output fields for the **monitor ethernet delay-measurement one-way** command.
- [Table 463 on page 4373](#) lists the run-time output fields for the **monitor ethernet delay-measurement two-way** command.

Output fields are listed in the approximate order in which they appear.

**Table 462: monitor ethernet delay-measurement one-way Output Fields**

| Output Field Name                | Output Field Description                                                                       |
|----------------------------------|------------------------------------------------------------------------------------------------|
| <b>One-way ETH-DM request to</b> | Unicast MAC address of the remote peer MEP.                                                    |
| <b>Interface</b>                 | Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached. |
| <b>IDM Frames sent</b>           | PDU frames sent to the remote MEP in this ETH-DM session.                                      |
| <b>Packets transmitted</b>       | Total number of IDM PDU frames sent to the remote MEP during this measurement session.         |
| <b>Average delay</b>             | Average two-way frame delay measured in this session.                                          |
| <b>Average delay variation</b>   | Average frame jitter measured in this session.                                                 |
| <b>Best case delay</b>           | Lowest two-way frame delay measured in this session.                                           |
| <b>Worst case delay</b>          | Highest two-way frame delay measured in this session.                                          |

**NOTE:** For one-way delay measurement, these CLI output fields display **NA** ("not applicable") at the initiator MEP because one-way frame delay measurements occur at the receiver MEP.

**Table 463: monitor ethernet delay-measurement two-way Output Fields**

| Output Field Name                                          | Output Field Description                                                                                          |
|------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>Two-way Ethernet frame delay measurement request to</b> | Unicast MAC address of the remote peer MEP.                                                                       |
| <b>Interface</b>                                           | Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.                    |
| <b>DMR received from</b>                                   | Unicast MAC address of the remote MEP that transmitted this DMR frame in response to a DMM frame.                 |
| <b>Delay</b>                                               | Two-way delay, in microseconds, for the initiator-transmitted DMM frame.                                          |
| <b>Delay variation</b>                                     | Difference, in microseconds, between the current and previous delay values. This is also known as <i>jitter</i> . |
| <b>Packets transmitted</b>                                 | Total number of DMM PDU frames sent to the remote MEP in this measurement session.                                |
| <b>Valid packets received</b>                              | Total number of DMR PDU frames received from the remote MEP in this measurement session.                          |
| <b>Average delay</b>                                       | Average two-way frame delay measured in this session.                                                             |
| <b>Average delay variation</b>                             | Average frame jitter measured in this session.                                                                    |

Table 463: monitor ethernet delay-measurement two-way Output Fields (*continued*)

| Output Field Name       | Output Field Description                              |
|-------------------------|-------------------------------------------------------|
| <b>Best case delay</b>  | Lowest two-way frame delay measured in this session.  |
| <b>Worst case delay</b> | Highest two-way frame delay measured in this session. |

## Sample Output

### monitor ethernet delay-measurement one-way

```

user@switch> monitor ethernet delay-measurement one-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
One-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA

```

### monitor ethernet delay-measurement two-way

```

user@switch> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a Delay: 100 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 111 usec Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a Delay: 110 usec Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a Delay: 119 usec Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a Delay: 122 usec Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 108 usec Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 10
Average delay: 103 usec, Average delay variation: 8 usec
Best case delay: 92 usec, Worst case delay: 122 usec

```

### monitor ethernet delay-measurement two-way (Invalid DMR Frames Received)

```

user@switch> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a Delay: 100 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 111 usec Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a Delay: 110 usec Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a Delay: 119 usec Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a Delay: 122 usec Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a with invalid timestamp(s).
DMR received from 00:05:85:73:39:4a Delay: 108 usec Delay variation: 16 usec

--- Delay measurement statistics ---

```

Packets transmitted: 10, Valid packets received: 9, Invalid packets received: 1  
Average delay: 105 usec, Average delay variation: 9 usec  
Best case delay: 92 usec, Worst case delay: 122 usec

## show oam ethernet connectivity-fault-management delay-statistics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show oam ethernet connectivity-fault-management delay-statistics</b><br><b>&lt;count <i>entry-count</i>&gt;</b><br><b>&lt;local-mep <i>local-mep-id</i>&gt;</b><br><b>maintenance-association <i>ma-name</i></b><br><b>maintenance-domain <i>md-name</i></b><br><b>&lt;remote-mep <i>remote-mep-id</i>&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.5.<br>Command introduced in Junos OS Release 11.4 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | On MX Series routers with Ethernet interfaces on Dense Port Concentrators (DPCs), display ETH-DM delay statistics.<br><br>On EX Series switches, display delay measurement results.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <b>count <i>entry-count</i></b> —(Optional) Number of entries to display from the statistics table. The range of values is 1 through 100. The default value is 100 entries.<br><br><b>local-mep <i>local-mep-id</i></b> —(Optional) Numeric identifier of the local MEP. On MX Series routers, the range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.<br><br><b>maintenance-association <i>ma-name</i></b> —Name of an existing CFM maintenance association.<br><br><b>maintenance-domain <i>md-name</i></b> —Name of an existing connectivity fault management (CFM) maintenance domain.<br><br><b>remote-mep <i>remote-mep-id</i></b> —(Optional) Numeric identifier of the remote MEP. On MX Series routers, the range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>clear oam ethernet connectivity-fault-management statistics</i></li><li>• <a href="#">clear oam ethernet connectivity-fault-management delay-statistics on page 4367</a></li><li>• <i>show oam ethernet connectivity-fault-management interfaces</i></li><li>• <i>show oam ethernet connectivity-fault-management mep-database</i></li><li>• <i>show oam ethernet connectivity-fault-management mep-statistics</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet connectivity-fault-management delay-statistics on page 4378</a><br><a href="#">show oam ethernet connectivity-fault-management delay-statistics remote-mep on page 4378</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 464 on page 4377</a> lists the output fields for the <b>show oam ethernet connectivity-fault-management delay-statistics</b> command and the <b>show oam ethernet</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |



**connectivity-fault-management mep-statistics** command. Output fields are listed in the approximate order in which they appear.

**Table 464: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields**

| Output Field Name                      | Field Description                                                                                                                                                                                                                                                                                                        |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MEP identifier</b>                  | Maintenance association end point (MEP) numeric identifier.                                                                                                                                                                                                                                                              |
| <b>MAC address</b>                     | Unicast MAC address configured for the MEP.                                                                                                                                                                                                                                                                              |
| <b>Remote MEP count</b>                | Number of remote MEPs (unless you specify the <b>remote-mep</b> option).                                                                                                                                                                                                                                                 |
| <b>Remote MEP identifier</b>           | Numeric identifier of the remote MEP.                                                                                                                                                                                                                                                                                    |
| <b>Remote MAC address</b>              | Unicast MAC address of the remote MEP.                                                                                                                                                                                                                                                                                   |
| <b>Index</b>                           | Index number that corresponds to the ETH-DM entry in the CFM database.                                                                                                                                                                                                                                                   |
| <b>One-way delay (usec)</b>            | <p>For a one-way ETH-DM session, the frame delay time, in microseconds, measured at the receiver MEP.</p> <p>For a detailed description of one-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p>  |
| <b>Two-way delay (usec)</b>            | <p>For a two-way ETH-DM session, the frame delay time, in microseconds, measured at the initiator MEP.</p> <p>For a detailed description of two-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p> |
| <b>Average one-way delay</b>           | Average one-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                |
| <b>Average one-way delay variation</b> | Average one-way “frame jitter” for the statistics displayed.                                                                                                                                                                                                                                                             |
| <b>Best-case one-way delay</b>         | Lowest one-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                 |
| <b>Worst-case one-way delay</b>        | Highest one-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                |
| <b>Average two-way delay</b>           | Average two-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                |
| <b>Average two-way delay variation</b> | Average two-way “frame jitter” for the statistics displayed.                                                                                                                                                                                                                                                             |
| <b>Best-case two-way delay</b>         | Lowest two-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                 |
| <b>Worst-case two-way delay</b>        | Highest two-way frame delay calculated in this session.                                                                                                                                                                                                                                                                  |

## Sample Output

show oam ethernet connectivity-fault-  
management  
delay-statistics

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6 maintenance-association ma6
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 2
Remote MEP identifier: 101
Remote MAC address: 00:05:85:73:39:4a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)      (usec)
    1      259        519
    2      273        550
    3      287        571
    4      299        610
    5      313        650
Average one-way delay      : 286 usec
Average one-way delay variation: 62 usec
Best case one-way delay    : 259 usec
Worst case one-way delay   : 313 usec
Average two-way delay      : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay    : 519 usec
Worst case two-way delay   : 650 usec

Remote MEP identifier: 102
Remote MAC address: 00:04:55:63:39:5a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)      (usec)
    1      29        58
    2      23        59
    3      27        56
    4      29        62
    5      33        68
Average one-way delay      : 28 usec
Average one-way delay variation: 3 usec
Best case one-way delay    : 23 usec
Worst case one-way delay   : 33 usec
Average two-way delay      : 60 usec
Average two-way delay variation: 3 usec
Best case two-way delay    : 56 usec
Worst case two-way delay   : 68 usec
```

show oam ethernet connectivity-fault-  
management delay-statistics remote-mep

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6 maintenance-association ma6 remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39

Remote MEP identifier: 101
Remote MAC address: 00:05:85:73:39:4a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)      (usec)
    1      259        519
```

|   |     |     |
|---|-----|-----|
| 2 | 273 | 550 |
| 3 | 287 | 571 |
| 4 | 299 | 610 |
| 5 | 313 | 650 |

Average one-way delay : 286 usec  
Average one-way delay variation: 62 usec  
Best case one-way delay : 259 usec  
Worst case one-way delay : 313 usec  
Average two-way delay : 580 usec  
Average two-way delay variation: 26 usec  
Best case two-way delay : 519 usec  
Worst case two-way delay : 650 usec

## show oam ethernet connectivity-fault-management forwarding-state

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show oam ethernet connectivity-fault-management forwarding-state</b><br><b>interface <i>interface-name</i></b><br><b>&lt;brief   detail   extensive&gt;</b>                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management forwarding state information for Ethernet interfaces.                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>interface <i>interface-name</i></b> —Display forwarding state information for the specified Ethernet interface only.<br><br><b>brief   detail   extensive</b> —(Optional) Display the specified level of output.                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear oam ethernet connectivity-fault-management statistics on page 4370</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management path-database on page 4390</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management mip on page 4399</a></li> </ul>                                                                                                       |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet connectivity-fault-management forwarding-state on page 4381</a><br><a href="#">show oam ethernet connectivity-fault-management forwarding-state interface on page 4381</a><br><a href="#">show oam ethernet connectivity-fault-management forwarding-state interface detail on page 4382</a><br><a href="#">show oam ethernet connectivity-fault-management forwarding-state interface interface-name on page 4383</a> |
| <b>Output Fields</b>            | <a href="#">Table 465 on page 4380</a> lists the output fields for the <b>show oam ethernet connectivity-fault-management forwarding-state</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                      |

**Table 465: show oam ethernet connectivity-fault-management forwarding-state Output Fields**

| Field Name            | Field Description                        | Level of Output |
|-----------------------|------------------------------------------|-----------------|
| <b>Interface name</b> | Interface identifier.                    | All levels      |
| <b>Filter action</b>  | Filter action for messages at the level. | All levels      |
| <b>Nexthop type</b>   | Next-hop type.                           | All levels      |
| <b>Nexthop index</b>  | Next-hop index number.                   | <b>brief</b>    |
| <b>Level</b>          | Maintenance domain (MD) level.           | <b>detail</b>   |

Table 465: show oam ethernet connectivity-fault-management forwarding-state Output Fields (*continued*)

| Field Name | Field Description                        | Level of Output |
|------------|------------------------------------------|-----------------|
| Direction  | MEP direction configured.                | none            |
| CEs        | Number of customer edge (CE) interfaces. | All levels      |

## Sample Output

show oam ethernet  
connectivity-fault-  
management forwarding-  
state

```
user@host> show oam ethernet connectivity-fault-management forwarding-state
CEs: 3
```

Maintenance domain forwarding state:

| Level | Direction | Filter action | Nexthop<br>type | Nexthop<br>index |
|-------|-----------|---------------|-----------------|------------------|
| 0     |           | Drop          | none            |                  |
| 1     |           | Drop          | none            |                  |
| 2     |           | Drop          | none            |                  |
| 3     |           | Drop          | none            |                  |
| 4     |           | Drop          | none            |                  |
| 5     |           | Drop          | none            |                  |
| 6     |           | Drop          | none            |                  |
| 7     |           | Drop          | none            |                  |

show oam ethernet  
connectivity-fault-  
management forwarding-  
state interface

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface
Interface name: ge-3/0/0.0
```

Maintenance domain forwarding state:

| Level | Direction | Filter action | Nexthop<br>type | Nexthop<br>index |
|-------|-----------|---------------|-----------------|------------------|
| 0     |           | Drop          | none            |                  |
| 1     |           | Drop          | none            |                  |
| 2     |           | Drop          | none            |                  |
| 3     |           | Drop          | none            |                  |
| 4     |           | Drop          | none            |                  |
| 5     |           | Drop          | none            |                  |
| 6     |           | Drop          | none            |                  |
| 7     | down      | Receive       | none            |                  |

```
Interface name: xe-0/0/0.0
```

```
Instance name: __+bd1__
```

Maintenance domain forwarding state:

| Level | Direction | Filter action | Nexthop<br>type | Nexthop<br>index |
|-------|-----------|---------------|-----------------|------------------|
| 0     |           | Drop          | none            |                  |

|   |      |         |      |
|---|------|---------|------|
| 1 |      | Drop    | none |
| 2 |      | Drop    | none |
| 3 |      | Drop    | none |
| 4 |      | Drop    | none |
| 5 |      | Drop    | none |
| 6 |      | Drop    | none |
| 7 | down | Receive | none |

show oam ethernet  
connectivity-fault-  
management forwarding-  
state interface detail

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface  
detail
```

```
Interface name: ge-3/0/0.0
```

```
Level: 0  
Filter action: Drop  
Nexthop type: none
```

```
Level: 1  
Filter action: Drop  
Nexthop type: none
```

```
Level: 2  
Filter action: Drop  
Nexthop type: none
```

```
Level: 3  
Filter action: Drop  
Nexthop type: none
```

```
Level: 4  
Filter action: Drop  
Nexthop type: none
```

```
Level: 5  
Filter action: Drop  
Nexthop type: none
```

```
Level: 6  
Filter action: Drop  
Nexthop type: none
```

```
Level: 7  
Direction: down  
Filter action: Receive  
Nexthop type: none
```

```
Interface name: xe-0/0/0.0
```

```
Level: 0  
Filter action: Drop  
Nexthop type: none
```

```
Level: 1  
Filter action: Drop  
Nexthop type: none
```

```
...
```

```
show oam ethernet  
connectivity-fault-  
management forwarding-  
state interface  
interface-name
```

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface  
interface-name ge-3/0/0.0  
Interface name: ge-3/0/0.0
```

Maintenance domain forwarding state:

| Level | Direction | Filter action | Nexthop<br>type | Nexthop<br>index |
|-------|-----------|---------------|-----------------|------------------|
| 0     |           | Drop          | none            |                  |
| 1     |           | Drop          | none            |                  |
| 2     |           | Drop          | none            |                  |
| 3     |           | Drop          | none            |                  |
| 4     |           | Drop          | none            |                  |
| 5     |           | Drop          | none            |                  |
| 6     |           | Drop          | none            |                  |
| 7     | down      | Receive       | none            |                  |

## show oam ethernet connectivity-fault-management interfaces

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show oam ethernet connectivity-fault-management interfaces</b><br><b>&lt;ethernet-interface-name&gt;</b><br><b>&lt;level md-level&gt;</b><br><b>&lt;brief   detail   extensive&gt;</b>                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for Ethernet interfaces.                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>ethernet-interface-name</b>—(Optional) Display CFM information only for CFM entities attached to the specified Ethernet interface.</p> <p><b>level md-level</b>—(Optional) Display CFM information for CFM identities enclosed within a maintenance domain of the specified level.</p>                                                                                                                       |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear oam ethernet connectivity-fault-management statistics on page 4370</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management path-database on page 4390</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management mip on page 4399</a></li> </ul>                                                                                                                                                                |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet connectivity-fault-management interfaces on page 4387</a><br><a href="#">show oam ethernet connectivity-fault-management interfaces detail on page 4387</a><br><a href="#">show oam ethernet connectivity-fault-management interfaces extensive on page 4388</a><br><a href="#">show oam ethernet connectivity-fault-management interfaces level on page 4389</a><br><a href="#">show oam ethernet connectivity-fault-management interfaces (Trunk Interfaces) on page 4389</a> |
| <b>Output Fields</b>            | <a href="#">Table 466 on page 4384</a> lists the output fields for the <b>show oam ethernet connectivity-fault-management interfaces</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                     |

**Table 466: show oam ethernet connectivity-fault-management interfaces Output Fields**

| Field Name                     | Field Description                         | Level of Output  |
|--------------------------------|-------------------------------------------|------------------|
| <b>Interface</b>               | Interface identifier.                     | All levels       |
| <b>Interface status</b>        | Local interface status.                   | All levels       |
| <b>Link status</b>             | Local link status. Up, down, or oam-down. | All levels       |
| <b>Maintenance domain name</b> | Maintenance domain name.                  | detail extensive |



**Table 466: show oam ethernet connectivity-fault-management interfaces Output Fields (*continued*)**

| Field Name                              | Field Description                                                                                                                                                                                              | Level of Output         |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Format (Maintenance domain)</b>      | Maintenance domain name format configured.                                                                                                                                                                     | <b>detail extensive</b> |
| <b>Level</b>                            | Maintenance domain level configured.                                                                                                                                                                           | All levels              |
| <b>Maintenance association name</b>     | Maintenance association name.                                                                                                                                                                                  | <b>detail extensive</b> |
| <b>Format (Maintenance association)</b> | Maintenance association name format configured.                                                                                                                                                                | <b>detail extensive</b> |
| <b>Continuity-check status</b>          | Continuity-check status.                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Interval</b>                         | Continuity-check message interval.                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Loss-threshold</b>                   | Lost continuity-check message threshold.                                                                                                                                                                       | <b>detail extensive</b> |
| <b>MEP identifier</b>                   | Maintenance association end point (MEP) identifier.                                                                                                                                                            | All levels              |
| <b>Neighbours</b>                       | Number of MEP neighbors.                                                                                                                                                                                       | All levels              |
| <b>Direction</b>                        | MEP direction configured.                                                                                                                                                                                      | <b>detail extensive</b> |
| <b>MAC address</b>                      | MAC address configured for the MEP.                                                                                                                                                                            | <b>detail extensive</b> |
| <b>MEP status</b>                       | Indicates the status of the Connectivity Fault Management (CFM) protocol running on the MEP: <b>Running</b> , <b>inactive</b> , <b>disabled</b> , or <b>unsupported</b> .                                      | <b>detail extensive</b> |
| <b>Remote MEP not receiving CCM</b>     | Whether the remote MEP is not receiving connectivity check messages (CCMs).                                                                                                                                    | <b>detail extensive</b> |
| <b>Erroneous CCM received</b>           | Whether erroneous CCMs have been received.                                                                                                                                                                     | <b>detail extensive</b> |
| <b>Cross-connect CCM received</b>       | Whether cross-connect CCMs have been received.                                                                                                                                                                 | <b>detail extensive</b> |
| <b>RDI sent by some MEP</b>             | Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs. | <b>detail extensive</b> |
| <b>CCMs sent</b>                        | Number of CCMs transmitted.                                                                                                                                                                                    | <b>detail extensive</b> |
| <b>CCMs received out of sequence</b>    | Number of CCMs received out of sequence.                                                                                                                                                                       | <b>detail extensive</b> |

**Table 466: show oam ethernet connectivity-fault-management interfaces Output Fields (continued)**

| Field Name                                 | Field Description                                                                                                                                                                                                               | Level of Output         |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>LBM sent</b>                            | Number of loopback request messages (LBMs) sent.                                                                                                                                                                                | <b>detail extensive</b> |
| <b>Valid in-order LBRs received</b>        | Number of loopback response messages (LBRs) received that were valid messages and in sequence.                                                                                                                                  | <b>detail extensive</b> |
| <b>Valid out-of-order LBRs received</b>    | Number of LBRs received that were valid messages and not in sequence.                                                                                                                                                           | <b>detail extensive</b> |
| <b>LBRs received with corrupted data</b>   | Number of LBRs received that were corrupted.                                                                                                                                                                                    | <b>detail extensive</b> |
| <b>LBRs sent</b>                           | Number of LBRs transmitted.                                                                                                                                                                                                     | <b>detail extensive</b> |
| <b>LTMs sent</b>                           | Linktrace messages (LTMs) transmitted.                                                                                                                                                                                          | <b>detail extensive</b> |
| <b>LTMs received</b>                       | Linktrace messages received.                                                                                                                                                                                                    | <b>detail extensive</b> |
| <b>LTRs sent</b>                           | Linktrace responses (LTRs) transmitted.                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>LTRs received</b>                       | Linktrace responses received.                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Sequence number of next LTM request</b> | Sequence number of next LTM request to be transmitted.                                                                                                                                                                          | <b>detail extensive</b> |
| <b>1DMs sent</b>                           | If the interface is attached to an initiator MEP for a one-way ETH-DM session:<br>Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session.<br><br>For all other cases, this field displays 0. | <b>detail extensive</b> |
| <b>Valid 1DMs received</b>                 | If the interface is attached to a receiver MEP for a one-way ETH-DM session:<br>Number of valid 1DM frames received.<br><br>For all other cases, this field displays 0.                                                         | <b>detail extensive</b> |
| <b>Invalid 1DMs received</b>               | If the interface is attached to a receiver MEP for a one-way ETH-DM session:<br>Number of invalid 1DM frames received.<br><br>For all other cases, this field displays 0.                                                       | <b>detail extensive</b> |
| <b>DMMs sent</b>                           | If the interface is attached to an initiator MEP for a two-way ETH-DM session:<br>Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session.<br><br>For all other cases, this field displays 0. | <b>detail extensive</b> |
| <b>DMRs sent</b>                           | If the interface is attached to a responder MEP for a two-way ETH-DM session:<br>Number of Delay Measurement Reply (DMR) frames sent.<br><br>For all other cases, this field displays 0.                                        | <b>detail extensive</b> |

**Table 466: show oam ethernet connectivity-fault-management interfaces Output Fields (continued)**

| Field Name                      | Field Description                                                                                                                                                     | Level of Output         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Valid DMRs received</b>      | If the interface is attached to an initiator MEP for a two-way ETH-DM session:<br>Number of valid DMRs received.<br><br>For all other cases, this field displays 0.   | <b>detail extensive</b> |
| <b>Invalid DMRs received</b>    | If the interface is attached to an initiator MEP for a two-way ETH-DM session:<br>Number of invalid DMRs received.<br><br>For all other cases, this field displays 0. | <b>detail extensive</b> |
| <b>Remote MEP count</b>         | Number of remote MEPs.                                                                                                                                                | <b>extensive</b>        |
| <b>Identifier (remote MEP)</b>  | MEP identifier of the remote MEP.                                                                                                                                     | <b>extensive</b>        |
| <b>MAC address (remote MEP)</b> | MAC address of the remote MEP.                                                                                                                                        | <b>extensive</b>        |
| <b>State (remote MEP)</b>       | State of the remote MEP.                                                                                                                                              | <b>extensive</b>        |
| <b>Interface (remote MEP)</b>   | Interface of the remote MEP.                                                                                                                                          | <b>extensive</b>        |

## Sample Output

### show oam ethernet connectivity-fault-management interfaces

```

user@host> show oam ethernet connectivity-fault-management interfaces
Interface      Link      Status      Level      MEP      Neighbours
               Identifier
ge-1/1/0.0     Up        Active      0          2        1
ge-1/1/0.1     Up        Active      0          2        1
ge-1/1/0.10    Up        Active      0          2        1
ge-1/1/0.100   Up        Active      0          2        1
ge-1/1/0.101   Up        Active      0          2        1
ge-1/1/0.102   Up        Active      0          2        1
ge-1/1/0.103   Up        Active      0          2        1
ge-1/1/0.104   Up        Active      0          2        1
ge-1/1/0.105   Up        Active      0          2        1
ge-1/1/0.106   Up        Active      0          2        1

```

...

### show oam ethernet connectivity-fault-management interfaces detail

```

user@host> show oam ethernet connectivity-fault-management interfaces detail
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Maintenance domain name: md0, Format: string, Level: 5
Maintenance association name: ma1, Format: string

```

```

Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
MEP status: running
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : yes
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : yes
Statistics:
  CCMs sent                                  : 76
  CCMs received out of sequence               : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received                : 0
  Valid out-of-order LBRs received            : 0
  LBRs received with corrupted data           : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                              : 0
  LTRs sent                                  : 0
  LTRs received                              : 0
  Sequence number of next LTM request         : 0
  1DMs sent                                  : 0
  Valid 1DMs received                        : 0
  Invalid 1DMs received                      : 0
  DMMs sent                                  : 0
  DMRs sent                                  : 0
  Valid DMRs received                       : 0
  Invalid DMRs received                     : 0
Remote MEP count: 2
  Identifier  MAC address  State  Interface
  2001       00:90:69:0b:7f:71  ok    ge-5/2/9.0
  4001       00:90:69:0b:09:c5  ok    ge-5/2/9.0

```

#### show oam ethernet connectivity-fault-management interfaces extensive

```

user@host> show oam ethernet connectivity-fault-management interfaces extensive
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Maintenance domain name: md0, Format: string, Level: 5
Maintenance association name: ma1, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
MEP status: running
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : yes
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : yes
Statistics:
  CCMs sent                                  : 76
  CCMs received out of sequence               : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received                : 0
  Valid out-of-order LBRs received            : 0
  LBRs received with corrupted data           : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                              : 0
  LTRs sent                                  : 0
  LTRs received                              : 0

```

```

Sequence number of next LTM request      : 0
1DMs sent                               : 0
Valid 1DMs received                      : 0
Invalid 1DMs received                   : 0
DMMs sent                               : 0
DMRs sent                               : 0
Valid DMRs received                    : 0
Invalid DMRs received                  : 0
Remote MEP count: 2
Identifier    MAC address      State   Interface
2001         00:90:69:0b:7f:71  ok     ge-5/2/9.0
4001         00:90:69:0b:09:c5  ok     ge-5/2/9.0

```

### show oam ethernet connectivity-fault-management interfaces level

```

user@host> show oam ethernet connectivity-fault-management interfaces level 7
Interface      Link      Status      Level      MEP      Neighbours
Identifier
ge-3/0/0.0     Up        Active      7          201      0
xe-0/0/0.0     Up        Active      7          203      1

```

### show oam ethernet connectivity-fault-management interfaces (Trunk Interfaces)

```

user@host> show oam ethernet connectivity-fault-management interfaces

Interface      Link      Status      Level      MEP      Neighbours
Identifier
ge-4/0/1.0, vln 100    Up        Active      5          100      0
ge-10/3/10.4091, vln 4091 Down      Inactive    4          400      0
ge-4/0/0.0         Up        Active      6          200      0

```

```

user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/0.0

```

```

Interface      Link      Status      Level      MEP      Neighbours
Identifier
ge-4/0/0.0     Up        Active      6          200      0

```

```

user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/1.0 vln 100

```

```

Interface      Link      Status      Level      MEP      Neighbours
Identifier
ge-4/0/1.0, vln 100    Up        Active      5          100      0

```

```

user@host> show oam ethernet connectivity-fault-management interfaces ge-10/3/10.4091
vln 4091

```

```

Interface      Link      Status      Level      MEP      Neighbours
Identifier
ge-10/3/10.4091, vln 4091 Down      Inactive    4          400      0

```

## show oam ethernet connectivity-fault-management path-database

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show oam ethernet connectivity-fault-management path-database</b> <i>host</i><br><b>maintenance-association</b> <i>ma-name</i> <b>maintenance-domain</b> <i>md-name</i> <i>mac-address</i>                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management maintenance linktrace database information.                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>mac-address</b>—Display connectivity fault management path database information for the specified MAC address of the remote host.</p> <p><b>maintenance-association</b> <i>ma-name</i>—Display connectivity fault management path database information for the specified maintenance association.</p> <p><b>maintenance-domain</b> <i>md-name</i>—Display connectivity fault management path database information for the specified maintenance domain.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear oam ethernet connectivity-fault-management statistics on page 4370</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management interfaces on page 4384</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management mip on page 4399</a></li> </ul>                                                                                                                       |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet connectivity-fault-management path-database on page 4391</a><br><a href="#">show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands) on page 4391</a>                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | Table 467 on page 4390 lists the output fields for the <b>show oam ethernet connectivity-fault-management path-database</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                      |

**Table 467: show oam ethernet connectivity-fault-management linktrace path-database Output Fields**

| Field Name                     | Field Description                                                           |
|--------------------------------|-----------------------------------------------------------------------------|
| <b>Linktrace to</b>            | MAC address of the 802.1ag node to which the linktrace message is targeted. |
| <b>Interface</b>               | Interface used by the local MEP to send the linktrace message (LTM).        |
| <b>Maintenance Domain</b>      | Maintenance domain identifier specified in the traceroute command.          |
| <b>Maintenance Association</b> | Maintenance association identifier specified in the traceroute command.     |
| <b>Level</b>                   | Maintenance domain level configured for the maintenance domain.             |

**Table 467: show oam ethernet connectivity-fault-management linktrace path-database Output Fields (continued)**

| Field Name                    | Field Description                                                                                                                                                                                                                                                             |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Local Mep</b>              | MEP identifier of the local MEP originating the linktrace.                                                                                                                                                                                                                    |
| <b>Hop</b>                    | Sequential hop count of the linktrace path.                                                                                                                                                                                                                                   |
| <b>TTL</b>                    | Number of hops remaining in the linktrace message (LTM). The time to live (TTL) is decremented at each hop.                                                                                                                                                                   |
| <b>Source MAC address</b>     | MAC address of the 802.1ag maintenance intermediate point (MIP) that is forwarding the LTM.                                                                                                                                                                                   |
| <b>Next hop MAC address</b>   | MAC address of the 802.1ag node that is the next hop in the LTM path.                                                                                                                                                                                                         |
| <b>Transaction Identifier</b> | 4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all maintenance domains. Use the transaction identifier to match an incoming linktrace responses (LTR), with a previously sent LTM. |

## Sample Output

### show oam ethernet connectivity-fault-management path-database

```

user@host> show oam ethernet connectivity-fault-management path-database
maintenance-domain MD1 maintenance-association MA1 00:01:02:03:04:05
Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
Maintenance Domain: MD1, Level: 7
Maintenance Association: MA1, Local Mep: 1

Hop      TTL      Source MAC address      Next hop MAC address
Transaction Identifier:100001
1         63      00:00:aa:aa:aa:aa      00:00:bb:bb:bb:bb
2         62      00:00:bb:bb:bb:bb      00:00:cc:cc:cc:cc
3         61      00:00:cc:cc:cc:cc      00:01:02:03:04:05
4         60      00:01:02:03:04:05      00:00:00:00:00:00

```

### show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands)

```

user@host> show oam ethernet connectivity-fault-management path-database
maintenance-domain MD2 maintenance-association MA2 00:06:07:08:09:0A
Linktrace to 00:06:07:08:09:0A, Interface : ge-5/0/1.0
Maintenance Domain: MD2, Level: 6
Maintenance Association: MA2, Local Mep: 10

Hop      TTL      Source MAC address      Next hop MAC address
Transaction Identifier:100002
1         63      00:00:aa:aa:aa:aa      00:00:bb:bb:bb:bb
2         62      00:00:bb:bb:bb:bb      00:00:cc:cc:cc:cc
3         61      00:00:cc:cc:cc:cc      00:06:07:08:09:0A
4         60      00:06:07:08:09:0A      00:00:00:00:00:00
Transaction Identifier:100003
1         63      00:00:aa:aa:aa:aa      00:00:bb:bb:bb:bb
2         62      00:00:bb:bb:bb:bb      00:00:cc:cc:cc:cc

```

|   |    |                   |                   |
|---|----|-------------------|-------------------|
| 3 | 61 | 00:00:cc:cc:cc:cc | 00:06:07:08:09:0A |
| 4 | 60 | 00:06:07:08:09:0A | 00:00:00:00:00:00 |



## show oam ethernet connectivity-fault-management mep-database

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show oam ethernet connectivity-fault-management mep-database<br>maintenance-domain <i>domain-name</i><br>maintenance-association <i>ma-name</i><br><local-mep <i>local-mep-id</i> ><br><remote-mep <i>remote-mep-id</i> >                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>maintenance-association <i>ma-name</i></b>—Display connectivity fault management information for the specified maintenance association.</p> <p><b>maintenance-domain <i>domain-name</i></b>—Display connectivity fault management information for the specified maintenance domain.</p> <p><b>local-mep <i>local-mep-id</i></b>—(Optional) Display connectivity fault management information for the specified local MEP only.</p> <p><b>remote-mep <i>remote-mep-id</i></b>—(Optional) Display connectivity fault management information for the specified remote MEP only.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear oam ethernet connectivity-fault-management statistics on page 4370</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management interfaces on page 4384</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management mip on page 4399</a></li> </ul>                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet connectivity-fault-management mep-database on page 4397</a><br><a href="#">show oam ethernet connectivity-fault-management mep-database local-mep remote-mep on page 4397</a><br><a href="#">show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event) on page 4397</a>                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 468 on page 4393</a> lists the output fields for the <b>show oam ethernet connectivity-fault-management mep-database</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                            |

Table 468: show oam ethernet connectivity-fault-management mep-database Output Fields

| Field Name              | Field Description        |
|-------------------------|--------------------------|
| Maintenance domain name | Maintenance domain name. |

Table 468: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                              |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Format (Maintenance domain)</b>      | Maintenance domain name format configured.                                                                                                                                                                     |
| <b>Level</b>                            | Maintenance domain level configured.                                                                                                                                                                           |
| <b>Maintenance association name</b>     | Maintenance association name.                                                                                                                                                                                  |
| <b>Format (Maintenance association)</b> | Maintenance association name format configured.                                                                                                                                                                |
| <b>Continuity-check status</b>          | Continuity-check status.                                                                                                                                                                                       |
| <b>Interval</b>                         | Continuity-check message interval.                                                                                                                                                                             |
| <b>MEP identifier</b>                   | Maintenance association end point (MEP) identifier.                                                                                                                                                            |
| <b>Direction</b>                        | MEP direction configured.                                                                                                                                                                                      |
| <b>MAC address</b>                      | MAC address configured for the MEP.                                                                                                                                                                            |
| <b>Auto-discovery</b>                   | Whether automatic discovery is enabled or disabled.                                                                                                                                                            |
| <b>Priority</b>                         | Priority used for CCMs and linktrace messages transmitted by the MEP.                                                                                                                                          |
| <b>Interface name</b>                   | Interface identifier.                                                                                                                                                                                          |
| <b>Interface status</b>                 | Local interface status.                                                                                                                                                                                        |
| <b>Link status</b>                      | Local link status.                                                                                                                                                                                             |
| <b>Remote MEP not receiving CCM</b>     | Whether the remote MEP is not receiving CCMs.                                                                                                                                                                  |
| <b>Erroneous CCM received</b>           | Whether erroneous CCMs have been received.                                                                                                                                                                     |
| <b>Cross-connect CCM received</b>       | Whether cross-connect CCMs have been received.                                                                                                                                                                 |
| <b>RDI sent by some MEP</b>             | Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs. |
| <b>CCMs sent</b>                        | Number of CCMs transmitted.                                                                                                                                                                                    |
| <b>CCMs received out of sequence</b>    | Number of CCMs received out of sequence.                                                                                                                                                                       |

**Table 468: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)**

| Field Name                                 | Field Description                                                                                                                                                                                             |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>LBMs sent</b>                           | Number of loopback messages (LBMs) sent.                                                                                                                                                                      |
| <b>Valid in-order LBRs received</b>        | Number of loopback response messages (LBRs) received that were valid messages and in sequence.                                                                                                                |
| <b>Valid out-of-order LBRs received</b>    | Number of LBRs received that were valid messages and not in sequence.                                                                                                                                         |
| <b>LBRs received with corrupted data</b>   | Number of LBRs received that were corrupted.                                                                                                                                                                  |
| <b>LBRs sent</b>                           | Number of LBRs transmitted.                                                                                                                                                                                   |
| <b>LTMs sent</b>                           | Linktrace messages (LTMs) transmitted.                                                                                                                                                                        |
| <b>LTMs received</b>                       | Linktrace messages received.                                                                                                                                                                                  |
| <b>LTRs sent</b>                           | Linktrace responses (LTRs) transmitted.                                                                                                                                                                       |
| <b>LTRs received</b>                       | Linktrace responses received.                                                                                                                                                                                 |
| <b>Sequence number of next LTM request</b> | Sequence number of the next linktrace message request to be transmitted.                                                                                                                                      |
| <b>1DMs sent</b>                           | <p>If the MEP is an initiator for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session.</p> <p>For all other cases, this field displays 0.</p> |
| <b>Valid 1DMs received</b>                 | <p>If the MEP is a receiver for a one-way ETH-DM session: Number of valid 1DM frames received.</p> <p>For all other cases, this field displays 0.</p>                                                         |
| <b>Invalid 1DMs received</b>               | <p>If the MEP is a receiver for a one-way ETH-DM session: Number of invalid 1DM frames received.</p> <p>For all other cases, this field displays 0.</p>                                                       |
| <b>DMMs sent</b>                           | <p>If the MEP is an initiator for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session.</p> <p>For all other cases, this field displays 0.</p> |
| <b>DMRs sent</b>                           | <p>If the MEP is a responder for a ETH-DM session: Number of Delay Measurement Reply (DMR) frames sent.</p> <p>For all other cases, this field displays 0.</p>                                                |

Table 468: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Valid DMRs received      | If the MEP is an initiator for a two-way ETH-DM session: Number of valid DMRs received.<br>For all other cases, this field displays 0.                                                                                                                                                                                                                                                                                                                                                                                                              |
| Invalid DMRs received    | If the MEP is an initiator for a two-way ETH-DM session: Number of invalid DMRs received.<br>For all other cases, this field displays 0.                                                                                                                                                                                                                                                                                                                                                                                                            |
| Remote MEP identifier    | MEP identifier of the remote MEP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| State (remote MEP)       | State of the remote MEP: <b>idle</b> , <b>start</b> , <b>ok</b> , or <b>failed</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| MAC address              | MAC address of the remote MEP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Type                     | Whether the remote MEP MAC address was learned using automatic discovery or configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Interface                | Interface of the remote MEP. A seven-digit number is appended if CFM is configured to run on a routing instance of type VPLS.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Last flapped             | Date, time, and how long ago the remote MEP interface went from down to up. The format is <b>Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                                                               |
| Remote defect indication | Whether the remote defect indication (RDI) bit is set in messages that have been received or transmitted.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Port status TLV          | <ul style="list-style-type: none"> <li>In the Maintenance domain section, displays the last transmitted port status TLV value.</li> <li>In the Remote MEP section, displays the last value of port status TLV received from the remote MEP.</li> </ul> <p>In the Action profile section, displays, the last occurred event <b>port-status-tlv blocked</b> event. This event occurred due to the reception of <b>blocked</b> value in the port status TLV from remote MEP.</p>                                                                       |
| Interface status TLV     | <ul style="list-style-type: none"> <li>In the Maintenance domain section, displays the last transmitted interface status TLV value.</li> <li>In the Remote MEP section, displays the last value of interface status TLV received from the remote MEP.</li> </ul> <p>In the Action profile section, if displays, the last occurred event interface-status-tlv event ( either <b>lower-layer-down</b> or <b>down</b>). This event occurred due to the reception of either lower or <b>down</b> value in the interface status TLV from remote MEP.</p> |
| Action profile           | Name of the action profile occurrence associated with a remote MEP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Last event               | When an action profile occurs, displays the last event that triggered it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Last event cleared       | When all the configured and occurred events (under action profile) are cleared, then the action taken gets reverted (such as down interface is made up) and the corresponding time is noted and displayed.                                                                                                                                                                                                                                                                                                                                          |
| Action                   | Action taken and the corresponding time of the action occurrence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## Sample Output

### show oam ethernet connectivity-fault-management mep-database

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
Maintenance association name: vpls-vlan200, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
Auto-discovery: enabled, Priority: 0
Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : no
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : no
Statistics:
  CCMs sent                                  : 1476
  CCMs received out of sequence              : 0
  LBMs sent                                 : 85
Remote MEP count: 1
Identifier   MAC address      State   Interface
  100       00:19:e2:b2:81:4b   ok     vt-0/1/10.1049088

```

### show oam ethernet connectivity-fault-management mep-database local-mep remote-mep

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200 local-mep 200
remote-mep 100
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
Maintenance association name: vpls-vlan200, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
Auto-discovery: enabled, Priority: 0
Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up

Remote MEP identifier: 100, State: ok
MAC address: 00:19:e2:b2:81:4b, Type: Learned
Interface: vt-0/1/10.1049088
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: none

```

### show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event)

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 remote-mep 200
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

```

Remote MEP identifier: 200, State: ok  
MAC address: 00:05:85:73:96:1f, Type: Configured  
Interface: ge-1/0/8.0  
Last flapped: Never  
Remote defect indication: false  
Port status TLV: none  
Interface status TLV: lower-layer-down  
Action profile: juniper  
  Last event: Interface-status-tlv lower-layer-down  
  Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)

## show oam ethernet connectivity-fault-management mip

|                                 |                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show oam ethernet connectivity-fault-management mip<br><qualifier>                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                               |
| <b>Description</b>              | Display all the maintenance association intermediate points (MIPs) created in the system. Specify qualifiers to display specific MIPs.                                                                                                            |
| <b>Options</b>                  | <i>qualifier</i> —(Optional) Display the specified MIP.                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show oam ethernet connectivity-fault-management interfaces on page 4384</a></li> <li>• <a href="#">show oam ethernet connectivity-fault-management path-database on page 4390</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet connectivity-fault-management mip on page 4399</a>                                                                                                                                                                  |
| <b>Output Fields</b>            | <a href="#">Table 469 on page 4399</a> lists the output fields for the <b>show oam ethernet connectivity-fault-management mip</b> command. Output fields are listed in the approximate order in which they appear.                                |

Table 469: show oam ethernet connectivity-fault-management mip Output Fields

| Field Name                   | Field Description                                    |
|------------------------------|------------------------------------------------------|
| MIP information for instance | Header for the MIP information showing the MIP name. |
| Interface                    | Interface type-dpc/pic/port.unit-number.             |
| Level                        | MIP level configured.                                |

### Sample Output

show oam ethernet  
connectivity-fault-  
management mip

```

user@host> show oam ethernet connectivity-fault-management mip
MIP information for  __mip_name__

MIP information for  default-switch bd1

    Interface      Level
    ge-3/0/0.0     7
    ge-3/0/1.0     6
    ge-3/0/3.0     6

```

## show oam ethernet connectivity-fault-management sla-iterator-statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show oam ethernet connectivity-fault-management sla-iterator-statistics</b><br><b>maintenance-domain</b> <i>md-name</i><br><b>maintenance-association</b> <i>ma-name</i><br><b>sla-iterator</b> <i>sla-iterator</i><br><b>&lt;local-mep</b> <i>local-mep-id</i> <b>&gt;</b><br><b>&lt;remote-mep</b> <i>remote-mep-id</i> <b>&gt;</b>                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 11.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series routers.</p> <p>Command introduced in Junos OS Release 13.2 for MX Series routers (not on MPC3E Hyperion cards).</p>                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Display the Ethernet Operation, Administration, and Maintenance (OAM) service-level agreement (SLA) iterator statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>maintenance-domain</b> <i>md-name</i>—Name of an existing connectivity fault management (CFM) maintenance domain.</p> <p><b>maintenance-association</b> <i>ma-name</i>—Name of an existing CFM maintenance association.</p> <p><b>sla-iterator</b> <i>sla-iterator</i>— Name of the iterator profile.</p> <p><b>local-mep</b> <i>local-mep-id</i>—(Optional) Numeric identifier of the local MEP. The range of values is 1 through 8191.</p> <p><b>remote-mep</b> <i>remote-mep-id</i>—(Optional) Numeric identifier of the remote MEP. The range of values is 1 through 8192.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an Iterator Profile on a Switch (CLI Procedure) on page 4067</a></li> <li>• <a href="#">clear oam ethernet connectivity-fault-management sla-iterator-statistics on page 4369</a></li> </ul>                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet connectivity-fault-management sla-iterator-statistics on page 4403</a><br><a href="#">show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) on page 4403</a>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | <p><a href="#">Table 470 on page 4400</a> lists the output fields for the <b>show oam ethernet connectivity-fault-management sla-iterator-statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                            |

**Table 470: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields**

| Output Field Name  | Output Field Description        |
|--------------------|---------------------------------|
| Maintenance domain | Name of the maintenance domain. |



Table 470: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields (*continued*)

| Output Field Name                     | Output Field Description                                                                                     |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------|
| Level                                 | Level of the maintenance domain level configured.                                                            |
| Maintenance association               | Name of the maintenance association.                                                                         |
| Local MEP id                          | Numeric identifier of the local MEP.                                                                         |
| Remote MEP id                         | Numeric identifier of the remote MEP.                                                                        |
| Remote MAC address                    | Unicast MAC address of the remote MEP.                                                                       |
| Iterator name                         | Name of iterator.                                                                                            |
| Iterator Id                           | Numeric identifier of the iterator.                                                                          |
| Iterator cycle time                   | Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection |
| Iteration period                      | Maximum number of cycles per iteration                                                                       |
| Iterator status                       | Current status of iterator whether running or stopped.                                                       |
| Infinite iterations                   | Status of iteration as infinite or finite.                                                                   |
| Counter reset time                    | Date and time when the counter was reset.                                                                    |
| Reset reason                          | Reason to reset counter.                                                                                     |
| Delay weight                          | Calculation weight of delay.                                                                                 |
| Delay variation weight                | Calculation weight of delay variation.                                                                       |
| DMM sent                              | Delay measurement message (DMM) PDU frames sent to the peer MEP in this session.                             |
| DMM skipped for threshold hit         | Number of DMM frames sent to the peer MEP in this session skipped during threshold hit.                      |
| DMM skipped for threshold hit window  | Number of DMM frames sent to the peer MEP in this session skipped during the last threshold hit window.      |
| DMR received                          | Number of delay measurement reply (DMR) frames received.                                                     |
| DMR out of sequence                   | Total number of DMR out of sequence packets received.                                                        |
| DMR received with invalid time stamps | Total number of DMR frames received with invalid timestamps.                                                 |

**Table 470: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields (*continued*)**

| Output Field Name                                 | Output Field Description                                                                                                                                                                                                                                                                                                          |
|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Average two-way delay                             | Average two-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                         |
| Average two-way delay variation                   | Average two-way “frame jitter” for the statistics displayed.                                                                                                                                                                                                                                                                      |
| Average one-way forward delay variation           | Average one-way forward delay variation for the statistics displayed in microseconds.                                                                                                                                                                                                                                             |
| Average one-way backward delay variation          | Average one-way backward delay variation for the statistics displayed in microseconds.                                                                                                                                                                                                                                            |
| Weighted average two-way delay                    | Weighted average two-way delay for the statistics displayed in microseconds.                                                                                                                                                                                                                                                      |
| Weighted average two-way delay variation          | Weighted average two-way delay variation for the statistics displayed in microseconds.                                                                                                                                                                                                                                            |
| Weighted average one-way backward delay variation | Weighted average one-way backward delay variation for the statistics displayed in microseconds.                                                                                                                                                                                                                                   |
| Weighted average one-way forward delay variation  | Weighted average one-way forward delay variation for the statistics displayed in microseconds.                                                                                                                                                                                                                                    |
| SLM packets sent                                  | Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.                                                                                                                                                                                                   |
| SLM packets received                              | Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.                                                                                                                                                                                             |
| SLR packets sent                                  | Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.                                                                                                                                                                                                 |
| SLR packets received                              | Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.                                                                                                                                                                                           |
| Local TXFC1 value                                 | Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.   |
| Local RXFC1 value                                 | Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet. |
| Last Received SLR frame TXFCf(tc)                 | Value of the local counter TxFC1 at the time of SLM frame transmission.                                                                                                                                                                                                                                                           |

Table 470: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields (*continued*)

| Output Field Name               | Output Field Description                                                |
|---------------------------------|-------------------------------------------------------------------------|
| Last Received SLR frame TXFCb(t | Value of the local counter RxFCI at the time of SLR frame transmission. |
| Frame loss (near-end)           | Count of frame loss associated with ingress data frames.                |
| Frame loss (far-end)            | Count of frame loss associated with egress data frames.                 |

## Sample Output

### show oam ethernet connectivity-fault-management sla-iterator-statistics

```

user@switch> show oam ethernet connectivity-fault-management sla-iterator-statistics
sla-iterator il maintenance-domain default-1 maintenance-association ma1 local-mep 1
remote-mep 2
Iterator statistics:
Maintenance domain: md6, Level: 6
Maintenance association: ma6, Local MEP id: 1000
Remote MEP id: 103, Remote MAC address: 00:90:69:0a:43:92
Iterator name: il, Iterator Id: 1
Iterator cycle time: 10ms, Iteration period: 1 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2010-03-19 20:42:39 PDT (2d 18:24 ago)
Reset reason: Adjacency flap

Iterator delay measurement statistics:
Delay weight: 1, Delay variation weight: 1
DMM sent : 23898520
DMM skipped for threshold hit : 11000
DMM skipped for threshold hit window : 0
DMR received : 23851165
DMR out of sequence : 1142
DMR received with invalid time stamps : 36540
Average two-way delay : 129 usec
Average two-way delay variation : 15 usec
Average one-way forward delay variation : 22 usec
Average one-way backward delay variation : 22 usec
Weighted average two-way delay : 134 usec
Weighted average two-way delay variation : 8 usec
Weighted average one-way forward delay variation : 6 usec
Weighted average one-way backward delay variation : 2 usec

```

## Sample Output

### show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers)

```

user@switch> show oam ethernet connectivity-fault-management sla-iterator-statistics
maintenance-domain md1 maintenance-association mau local-mep 4 remote-mep 3 sla-iterator
lm
Iterator statistics:
Maintenance domain: 2, Level: 2
Maintenance association: W-160432000-001, Local MEP id: 2
Remote MEP id: 1, Remote MAC address: 00:90:69:0a:43:39

```

```
Iterator name: iter1, Iterator Id: 1
Iterator cycle time: 100ms, Iteration period: 10 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2012-09-25 02:15:31 PDT (00:00:45 ago)
Reset reason: Adjacency flap
Iterator loss measurement statistics:
  LMM sent : 444
  LMM skipped for threshold hit : 0
  LMM skipped for threshold hit window: 0
  LMR received : 444
  LMR out of sequence : 0
  LMR forwarding-class mismatch : 0
Accumulated transmit statistics:
  Near-end (CIR) : 0
  Far-end (CIR) : 0
  Near-end (EIR) : 0
  Far-end (EIR) : 0
Accumulated receive statistics:
  Near-end (CIR) : 0
  Far-end (CIR) : 0
  Near-end (EIR) : 0
  Far-end (EIR) : 0
Accumulated loss statistics:
  Near-end loss (CIR) : 0
  Near-end loss-ratio (CIR) : 0 (0.00000%)
  Far-end loss (CIR) : 0
  Far-end loss-ratio (CIR) : 0 (0.00000%)
  Near-end loss (EIR) : 0
  Near-end loss-ratio (EIR) : 0 (0.00000%)
  Far-end loss (EIR) : 0
  Far-end loss-ratio (EIR) : 0 (0.00000%)
Last loss measurement statistics:
  Near-end (CIR) : 0
  Far-end (CIR) : 0
  Near-end (EIR) : 0
  Far-end (EIR) :
```

---

## Operational Commands: Ethernet OAM Link Fault Management

- `show oam ethernet link-fault-management`

## show oam ethernet link-fault-management

|                                 |                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show oam ethernet link-fault-management<br><brief   detail><br><interface-name>                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                    |
| <b>Description</b>              | Displays Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.                                                                                                                                            |
| <b>Options</b>                  | <b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>interface-name</b> —(Optional) Display link fault management information for the specified Ethernet interface only.                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 4027</a></li> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 4062</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet link-fault-management brief on page 4409</a><br><a href="#">show oam ethernet link-fault-management detail on page 4409</a>                                                                                                             |
| <b>Output Fields</b>            | Table 471 on page 4405 lists the output fields for the <b>show oam ethernet link-fault-management</b> command. Output fields are listed in the approximate order in which they appear.                                                                                |

**Table 471: show oam ethernet link-fault-management Output Fields**

| Field Name             | Field Description                                                                                                                                                                                               | Level of Output |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Status</b>          | Indicates the status of the established link. <ul style="list-style-type: none"> <li>• <b>Fail</b>—A link fault condition exists.</li> <li>• <b>Running</b>—A link fault condition does not exist.</li> </ul>   | All levels      |
| <b>Discovery state</b> | State of the discovery mechanism: <ul style="list-style-type: none"> <li>• <b>Passive Wait</b></li> <li>• <b>Send Any</b></li> <li>• <b>Send Local Remote</b></li> <li>• <b>Send Local Remote Ok</b></li> </ul> | All levels      |
| <b>Peer address</b>    | Address of the OAM peer.                                                                                                                                                                                        | All levels      |

Table 471: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Flags</b>                     | Information about the interface. <ul style="list-style-type: none"> <li><b>Remote-Stable</b>—Indicates remote OAM client acknowledgment of, and satisfaction with local OAM state information. <b>False</b> indicates that remote DTE has either not seen or is unsatisfied with local state information. <b>True</b> indicates that remote DTE has seen and is satisfied with local state information.</li> <li><b>Local-Stable</b>—Indicates local OAM client acknowledgment of, and satisfaction with remote OAM state information. <b>False</b> indicates that local DTE either has not seen or is unsatisfied with remote state information. <b>True</b> indicates that local DTE has seen and is satisfied with remote state information.</li> <li><b>Remote-State-Valid</b>—Indicates the OAM client has received remote state information found within Local Information TLVs of received Information OAM PDUs. <b>False</b> indicates that OAM client has not seen remote state information. <b>True</b> indicates that the OAM client has seen remote state information.</li> </ul>                                                                      | All levels      |
| <b>Remote loopback status</b>    | Indicates the remote loopback status. An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels      |
| <b>Remote entity information</b> | Remote entity information. <ul style="list-style-type: none"> <li><b>Remote MUX action</b>—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs.</li> <li><b>Remote parser action</b>—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs.</li> <li><b>Discovery mode</b>—Indicates whether discovery mode is active or inactive.</li> <li><b>Unidirectional mode</b>—Indicates the ability to operate a link in a unidirectional mode for diagnostic purposes.</li> <li><b>Remote loopback mode</b>—Indicates whether remote loopback is supported or not supported.</li> <li><b>Link events</b>—Indicates whether interpreting link events is supported or not supported on the remote peer.</li> <li><b>Variable requests</b>—Indicates whether variable requests are supported or not supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer.</li> </ul> | All levels      |
| <b>OAM Receive Statistics</b>    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                 |
| <b>Information</b>               | The number of information PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>   |
| <b>Event</b>                     | The number of loopback control PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |
| <b>Variable request</b>          | The number of variable request PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |
| <b>Variable response</b>         | The number of variable response PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail</b>   |
| <b>Loopback control</b>          | The number of loopback control PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |

Table 471: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                                         | Field Description                                                                                                                                                                  | Level of Output |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Organization specific</b>                       | The number of vendor organization specific PDUs received.                                                                                                                          | <b>detail</b>   |
| <b>OAM Transmit Statistics</b>                     |                                                                                                                                                                                    |                 |
| <b>Information</b>                                 | The number of information PDUs transmitted.                                                                                                                                        | <b>detail</b>   |
| <b>Event</b>                                       | The number of event notification PDUs transmitted.                                                                                                                                 | <b>detail</b>   |
| <b>Variable request</b>                            | The number of variable request PDUs transmitted.                                                                                                                                   | <b>detail</b>   |
| <b>Variable response</b>                           | The number of variable response PDUs transmitted.                                                                                                                                  | <b>detail</b>   |
| <b>Loopback control</b>                            | The number of loopback control PDUs transmitted.                                                                                                                                   | <b>detail</b>   |
| <b>Organization specific</b>                       | The number of vendor organization specific PDUs transmitted.                                                                                                                       | <b>detail</b>   |
| <b>OAM Received Symbol Error Event information</b> |                                                                                                                                                                                    |                 |
| <b>Events</b>                                      | The number of symbol error event TLVs that have been received after the OAM sublayer was reset.                                                                                    | <b>detail</b>   |
| <b>Window</b>                                      | The symbol error event window in the received PDU.<br><br>The protocol default value is the number of symbols that can be received in one second on the underlying physical layer. | <b>detail</b>   |
| <b>Threshold</b>                                   | The number of errored symbols in the period required for the event to be generated.                                                                                                | <b>detail</b>   |
| <b>Errors in period</b>                            | The number of symbol errors in the period reported in the received event PDU.                                                                                                      | <b>detail</b>   |
| <b>Total errors</b>                                | The number of errored symbols that have been reported in received event TLVs after the OAM sublayer was reset.<br><br>Symbol errors are coding symbol errors.                      | <b>detail</b>   |
| <b>OAM Received Frame Error Event Information</b>  |                                                                                                                                                                                    |                 |
| <b>Events</b>                                      | The number of errored frame event TLVs that have been received after the OAM sublayer was reset.                                                                                   | <b>detail</b>   |
| <b>Window</b>                                      | The duration of the window in terms of the number of 100 ms period intervals.                                                                                                      | <b>detail</b>   |
| <b>Threshold</b>                                   | The number of detected errored frames required for the event to be generated.                                                                                                      | <b>detail</b>   |
| <b>Errors in period</b>                            | The number of detected errored frames in the period.                                                                                                                               | <b>detail</b>   |

Table 471: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                                               | Field Description                                                                                                                                                                       | Level of Output |
|----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Total errors</b>                                      | The number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset.<br><br>A frame error is any frame error on the underlying physical layer. | <b>detail</b>   |
| <b>OAM Received Frame Period Error Event Information</b> |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset.                                                                                 | <b>detail</b>   |
| <b>Window</b>                                            | The duration of the frame seconds window.                                                                                                                                               | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of frame seconds errors in the period.                                                                                                                                       | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of frame seconds errors in the period.                                                                                                                                       | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset.                                                                     | <b>detail</b>   |
| <b>OAM Transmitted Symbol Error Event Information</b>    |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset.                                                                                      | <b>detail</b>   |
| <b>Window</b>                                            | The symbol error event window in the transmitted PDU.                                                                                                                                   | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of errored symbols in the period required for the event to be generated.                                                                                                     | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of symbol errors in the period reported in the transmitted event PDU.                                                                                                        | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset.                                                                       | <b>detail</b>   |
| <b>OAM Transmitted Frame Error Event Information</b>     |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset.                                                                                     | <b>detail</b>   |
| <b>Window</b>                                            | The duration of the window in terms of the number of 100 ms period intervals.                                                                                                           | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of detected errored frames required for the event to be generated.                                                                                                           | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of detected errored frames in the period.                                                                                                                                    | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of errored frames that have been detected after the OAM sublayer was reset.                                                                                                  | <b>detail</b>   |



## Sample Output

### show oam ethernet link-fault-management brief

```
user@host> show oam ethernet link-fault-management brief
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:72:2c:83
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote loopback status: Disabled on local port, Enabled on peer port
Remote entity information:
  Remote MUX action: discarding, Remote parser action: loopback
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported
```

### show oam ethernet link-fault-management detail

```
user@host> show oam ethernet link-fault-management detail
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:0a:07:14
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
OAM receive statistics:
  Information: 186365, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0
OAM transmit statistics:
  Information: 186347, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0
OAM received symbol error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame period error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM transmitted symbol error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
OAM transmitted frame error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
Remote entity information:
  Remote MUX action: forwarding, Remote parser action: forwarding
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported
```

## Operational Commands: Uplink Failure Detection

- `show uplink-failure-detection`

## show uplink-failure-detection

|                                 |                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show uplink-failure-detection</code><br><code>&lt;group group-name&gt;</code>                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for EX Series switches.                                                                                                                                |
| <b>Description</b>              | Display information about the uplink-failure-detection group, the member interfaces, and their status.                                                                                             |
| <b>Options</b>                  | <b>none</b> —Display information about all groups configured for uplink failure detection.<br><b>group group-name</b> —(Optional) Display information about the specified group only.              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 4064</a></li> </ul>                                                 |
| <b>List of Sample Output</b>    | <a href="#">show uplink-failure-detection on page 4410</a>                                                                                                                                         |
| <b>Output Fields</b>            | <a href="#">Table 472 on page 4410</a> lists the output fields for the <code>show uplink-failure-detection</code> command. Output fields are listed in the approximate order in which they appear. |

**Table 472: show uplink-failure-detection Output Fields**

| Field Name     | Field Description                                                                                                                                                                                                         |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Group          | Name of the group.                                                                                                                                                                                                        |
| Uplink         | The uplink interface or interfaces configured as link-to-monitor.<br><br><b>NOTE:</b> The asterisk (*) indicates that the link is up.                                                                                     |
| Downlink       | The downlink interface or interfaces configured as link-to-disable.<br><br><b>NOTE:</b> The asterisk (*) indicates that the link is up.                                                                                   |
| Failure Action | Status of uplink failure detection: <ul style="list-style-type: none"> <li>Active—The switch has detected an uplink failure and has brought the downlink down.</li> <li>Inactive—The uplink or uplinks are up.</li> </ul> |

## Sample Output

### show uplink-failure-detection

```

user@switch> show uplink-failure-detection
Group           : group1
Uplink          : ge-0/0/0*
Downlink        : ge-0/0/1*
Failure Action   : Inactive

```

Group : group2  
Uplink : ge-0/0/3.0  
Downlink : ge-0/0/4.0  
Failure Action : Active



## CHAPTER 69

# Troubleshooting

- [Troubleshooting Procedures on page 4413](#)
- [Operational Commands on page 4414](#)

## Troubleshooting Procedures

---

- [Troubleshooting Port Mirroring Configuration Error Messages on page 4413](#)

## Troubleshooting Port Mirroring Configuration Error Messages

Troubleshooting issues with port mirroring on EX Series switches:

1. [An Analyzer Configuration Returns a “Multiple interfaces cannot be configured as a member of Analyzer output VLAN” Error Message on page 4413](#)

### [An Analyzer Configuration Returns a “Multiple interfaces cannot be configured as a member of Analyzer output VLAN” Error Message](#)

---

**Problem** **Description:** In an analyzer configuration, if the VLAN to which mirrored traffic is sent contains more than one member interface, the following error message is displayed in the CLI when you commit the analyzer configuration and the commit fails:

```
Multiple interfaces cannot be configured as a member of Analyzer output VLAN <vlan
name>
```

**Solution** You must direct the mirrored traffic to a VLAN that has a single member interface. You can do this by completing either of these tasks:

- Reconfigure the existing VLAN to contain a single member interface. You can choose this method if you want to use the existing VLAN.
- Create a new VLAN with a single member interface and associate the VLAN with the analyzer.

To reconfigure the existing VLAN to contain only one member interface:

1. Remove member interfaces from the VLAN repeatedly by using either the **delete vlan** command or the **delete interface** command until the VLAN contains a single member interface:

- [edit]  
user@switch# **delete** vlan *vlan-id* interface *interface-name*
  - [edit]  
user@switch# **delete** interface *interface-name* unit 0 family *family-name* vlan member *vlan-id*
2. (Optional) Confirm that the VLAN contains only one interface:
- [edit]  
user@switch# **show** vlans *vlan-name*
- The output for this command must display only one interface.

To create a new VLAN with a single member interface:

1. Configure a VLAN to carry the mirrored traffic:
- [edit]  
user@switch# **set** vlans *vlan-name*
2. Associate an interface with the VLAN:
- [edit]  
user@switch# **set** interfaces *interface-name* unit *logical-unit-number* family *family-name* vlan members *vlan-name*
3. Associate the VLAN with the analyzer:
- [edit ethernet-switching-options]  
user@switch# **set** analyzer *analyzer-name* output vlan *vlan-name*

**Related  
Documentation**

- *Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches*
- *Configuring Port Mirroring to Analyze Traffic (CLI Procedure)*
- [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 4047](#)
- *Understanding Port Mirroring on EX Series Switches*

---

## Operational Commands

- [show pfe statistics bridge](#)

## show pfe statistics bridge

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show pfe statistics bridge</b><br><b>&lt;fpc slot&gt;</b>                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Display information about the number of packets discarded in the ingress pipeline of the Packet Forwarding Engine, packets discarded because of egress filtering or congestion filtering, number of control packets, and general counters for dropped packets. You can use this information to inform troubleshooting investigations.                                                                          |
| <b>Options</b>                  | <b>none</b> —Display bridge counter statistics for all Flexible PIC Concentrator (FPC) slots.<br><b>fpc slot</b> —(Optional) Display bridge counter statistics for a specific FPC slot.                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring System Process Information on page 850</a></li> <li>• <a href="#">Monitoring Switch Control Traffic on page 843</a></li> </ul>                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <a href="#">show pfe statistics bridge (EX3200 and EX4200 Switches) on page 4416</a><br><a href="#">show pfe statistics bridge (EX8200 Switches and EX8200 Virtual Chassis) on page 4417</a><br><a href="#">show pfe statistics bridge fpc (EX8200 Switches and EX8200 Virtual Chassis) on page 4418</a><br><a href="#">show pfe statistics bridge fpc (EX8200-40XS (40-port SFP+) Line Card) on page 4418</a> |
| <b>Output Fields</b>            | <a href="#">Table 438 on page 4284</a> lists the output fields for the <b>show pfe statistics bridge</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                      |

**Table 473: show pfe statistics bridge Output Fields**

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Ingress Counters</b> | <p>Information about ingress counters:</p> <ul style="list-style-type: none"> <li>• <b>Received</b>—Number of packets received by the bridge.</li> <li>• <b>VLAN Filtered</b>—Number of packets discarded because of VLAN filtering.</li> <li>• <b>Security Filtered</b>—Number of packets discarded because of security filtering.</li> <li>• <b>Other Discards</b>—Number of packets dropped by the bridge for reasons other than VLAN or security filtering.</li> </ul> |

Table 473: show pfe statistics bridge Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Egress Counters</b>  | <p>Information about egress counters:</p> <ul style="list-style-type: none"> <li>• <b>Unicast</b>—Number of unicast packets transmitted.</li> <li>• <b>Multicast</b>—Number of multicast packets transmitted.</li> <li>• <b>Broadcast</b>—Number of broadcast packets transmitted.</li> <li>• <b>Egress Filtered</b>—Number of egress-filtered packets (regardless of port, priority, or mode).</li> <li>• <b>TailDrop</b>—Number of packets filtered because of egress queue congestion.</li> <li>• <b>Forward Restrict</b>—Number of packets filtered because of egress forward restrictions.</li> <li>• <b>Congestion Filtered</b>—Number of packets filtered because of transmit queue (TxQ) congestion.</li> <li>• <b>Control Packets</b>—Number of control packets (sent to CPU, received from CPU, and sent to analyzer).</li> </ul> |
| <b>Drop Counters</b>    | <p>Information about drop counters:</p> <ul style="list-style-type: none"> <li>• <b>Drop Mode</b>—Count mode of the counter.</li> <li>• <b>Drop Counter</b>—Counter value.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>General Counters</b> | <p>Information about general counters:</p> <ul style="list-style-type: none"> <li>• <b>Drop Mode</b>—Count mode of the counter.</li> <li>• <b>Drop Counter</b>—Counter value.</li> <li>• <b>Source Not Learnt</b>—Number of source addresses that were not learnt because of internal congestion.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>MUX PFE</b>          | <p>Information about multiplexer PFE for oversubscribed cards:</p> <ul style="list-style-type: none"> <li>• <b>Drop Mode</b>—Count mode of the counter.</li> <li>• <b>Drop Count</b>—Counter value.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Sample Output

### show pfe statistics bridge (EX3200 and EX4200 Switches)

```

user@switch> show pfe statistics bridge
Slot 0

PFE:                0          1          2
-----
---- Ingress Counters ----
Received:            0          52          0
VLAN Filtered:       0          0          0
Security Filtered:   0          0          0
Other Discards:      0          0          0
---- Egress Counters ----
Unicast:             0         104         52
Multicast:            0          0          0
Broadcast:            0          0          0
Egress Filtered:     0          0          0
Congestion Filtered: 0          0          0
Control Packets:     5          0          0
---- General Counters ----
Drop Mode:           0          0          0

```



```
Drop Counter:          34217          36080          6367
Source Not Learnt:      0              0              0
```

### show pfe statistics bridge (EX8200 Switches and EX8200 Virtual Chassis)

```
user@switch> show pfe statistics bridge
```

```
Slot 0
```

```
PFE:                  0              1
```

```
----- Ingress Counters -----
```

```
Received:              946              48
VLAN Filtered:         0              0
Security Filtered:     0              0
Other Discards:        0              0
```

```
----- Egress Counters -----
```

```
Unicast:               0              0
Multicast:             0              0
Broadcast:             0              0
Egress Filtered:       0              0
TailDrop:              0              0
Forward Restrict:      0              0
Congestion Filtered:   0              0
Control Packets:       4103           896
```

```
----- Drop Counters -----
```

```
Drop Mode:             0              0
Drop Counter:          12528          2
```

```
Slot 1
```

```
PFE:                  0              1
```

```
----- Ingress Counters -----
```

```
Received:              0              0
VLAN Filtered:         0              0
Security Filtered:     0              0
Other Discards:        0              0
```

```
----- Egress Counters -----
```

```
Unicast:               0              0
Multicast:             0              0
Broadcast:             0              0
Egress Filtered:       0              0
TailDrop:              0              0
Forward Restrict:      0              0
Congestion Filtered:   0              0
Control Packets:       0              0
```

```
----- Drop Counters -----
```

```
Drop Mode:             0              0
Drop Counter:          0              0
```

```
Slot 2
```

```
PFE:                  0              1
```

```
----- Ingress Counters -----
```

```
Received:              0              0
VLAN Filtered:         0              0
Security Filtered:     0              0
Other Filtered:        0              0
```

```
----- Egress Counters -----
```

```
Unicast:               0              0
Multicast:             0              0
```

```

Broadcast:                0          0
Egress Filtered:          0          0
TailDrop:                 0          0
Forward Restrict:         0          0
Congestion Filtered:      0          0
Control Packets:          0          0
---- Drop Counters ----
Drop Mode:                0          0
Drop Counter:             0          0

```

#### show pfe statistics bridge fpc (EX8200 Switches and EX8200 Virtual Chassis)

```

user@switch> show pfe statistics bridge fpc 2
Slot 2

```

```

PFE:                0          1
-----
---- Ingress Counters ----
Received:           0          0
VLAN Filtered:      0          0
Security Filtered:  0          0
Other Discards:     0          0
---- Egress Counters ----
Unicast:            0          0
Multicast:          0          0
Broadcast:          0          0
Egress Filtered:    0          0
TailDrop:           0          0
Forward Restrict:   0          0
Congestion Filtered: 0          0
Control Packets:    0          0
---- Drop Counters ----
Drop Mode:          0          0
Drop Counter:       0          0

```

#### show pfe statistics bridge fpc (EX8200-40XS (40-port SFP+) Line Card)

```

user@switch> show pfe statistics bridge fpc 8
Slot 8

```

```

PFE:                0          1          2          3
-----
---- Ingress Counters ----
Received:           0          3          0          0
VLAN Filtered:      0          0          0          0
Security Filtered:  0          0          0          0
Other Discards:     0          1          0          0
---- Egress Counters ----
Unicast:            0          0          0          0
Multicast:          0          0          0          0
Broadcast:          0          0          0          0
Egress Filtered:    0          0          0          0
TailDrop:           0          0          0          0
Forward Restrict:   0          0          0          0
Congestion Filtered: 0          2          0          0
Control Packets:    4          0          0          0
---- Drop Counters ----
Drop Mode:          0          0          0          0
Drop Counter:       0          1          0          0

MUX PFE:            4          5

```

---

|             |   |   |
|-------------|---|---|
| Drop Mode:  | 0 | 0 |
| Drop Count: | 0 | 0 |



## PART 22

# PoE

- [Overview on page 4423](#)
- [Configuration on page 4433](#)
- [Administration on page 4467](#)
- [Troubleshooting Procedures on page 4493](#)



## CHAPTER 70

# Overview

- [Power over Ethernet Overview on page 4423](#)

## Power over Ethernet Overview

---

- [Understanding PoE on EX Series Switches on page 4423](#)

## Understanding PoE on EX Series Switches

Power over Ethernet (PoE) enables electric power, along with data, to be passed over a copper Ethernet LAN cable. Powered devices—such as VoIP telephones, wireless access points, video cameras, and point-of-sale devices—that support PoE can receive power safely from the same access ports that are used to connect personal computers to the network. This reduces the amount of wiring in a network, and also eliminates the need to position a powered device near an AC power outlet, making network design more flexible and efficient.



**NOTE:** We recommend that you do not connect an enabled PoE port on one switch to an enabled PoE port on a second switch. If there is a large voltage difference between the power supplies of the two switches, the resulting negative current will trigger a fail-safe mechanism on the second switch that prevents the power sourcing equipment (PSE) from delivering power to the other PoE ports on that switch.

This topic describes PoE on Juniper Networks EX Series Ethernet Switches.

- [PoE, PoE+, and Enhanced PoE on page 4423](#)
- [PoE Power Allocation on page 4425](#)

## PoE, PoE+, and Enhanced PoE

---

PoE was first defined in the IEEE 802.3af standard. In this standard, the amount of power that can be supplied to a powered device is limited to 15.4 W. A later standard, IEEE 802.3at, defined PoE+, which increases the amount of power to 30 W. The PoE+ standard provides support for legacy PoE devices—an IEEE 802.3af powered device can operate normally when connected to IEEE 802.3at (PoE+) power sourcing equipment.

Beginning in Juniper Networks Junos operating system (Junos OS) Release 11.1, Juniper Networks provides enhanced PoE on EX3200 and EX4200 switches. Enhanced PoE is a Juniper Networks extension to the IEEE 802.3af standard that provides power of up to 18.6 W per PoE port.

[Table 474 on page 4424](#) lists EX Series switches and line cards and the version of PoE they support.

**Table 474: PoE Version Support**

| Switch or Line Card                                                                                                                         | PoE Version                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EX2200 switch<br><br>(EX2200-C-12P-2G, EX2200-24P-4G, EX2200-48P-4G models)                                                                 | PoE+ (IEEE 802.3at)<br><br><b>NOTE:</b> Starting with Junos OS Release 12.2R1, PoE commands are enabled on all non-PoE-capable EX2200 switch models. The PoE commands do not provide any meaningful configuration on standalone non-PoE-capable switch models. However, in an EX2200 Virtual Chassis, you can execute PoE commands from a non-PoE-capable master switch to configure PoE on PoE-capable Virtual Chassis members. |
| EX3200 switch<br><br>(EX3300-24P, EX3300-24T, EX3300-48P, EX3300-48T models)                                                                | Enhanced PoE                                                                                                                                                                                                                                                                                                                                                                                                                     |
| EX3300 switch<br><br>(EX3300-24P, EX3300-48P models)                                                                                        | PoE+ (IEEE 802.3at)                                                                                                                                                                                                                                                                                                                                                                                                              |
| EX4200 switch—P and T models<br><br>(EX4200-24P, EX4200-24T, EX4200-48P, EX4200-48T)                                                        | Enhanced PoE                                                                                                                                                                                                                                                                                                                                                                                                                     |
| EX4200 switch—PX models<br><br>(EX4200-24PX and EX4200-48PX)                                                                                | PoE+ (IEEE 802.3at)                                                                                                                                                                                                                                                                                                                                                                                                              |
| EX4300 switch<br><br>(EX4300-24P and EX4300-48P)                                                                                            | PoE+ (IEEE 802.3at)                                                                                                                                                                                                                                                                                                                                                                                                              |
| EX4600 switch<br><br>(EX4600-40F-AFO and EX4600-40F-AFI)                                                                                    | PoE+ (IEEE 802.3at)<br><br><b>NOTE:</b> PoE is only supported on EX4600 switches operating in a mixed Virtual Chassis with EX4300 switches.                                                                                                                                                                                                                                                                                      |
| EX6200-48P (48-port PoE+) line card                                                                                                         | PoE+ (IEEE 802.3at)                                                                                                                                                                                                                                                                                                                                                                                                              |
| EX8200-2XS-40P (40-port PoE+ with 4-port SFP and 2-port SFP+) line card<br><br>EX8200-48PL (2-port SFP+ and 48-port PoE+ 20 Gbps) line card | PoE+ (IEEE 802.3at)—Ports 0 through 11, and PoE (IEEE 802.3af)—remaining PoE ports.                                                                                                                                                                                                                                                                                                                                              |





**NOTE:** This topic and its related topics use the term PoE as a generic term to refer to PoE, PoE+, and enhanced PoE.

## PoE Power Allocation

A switch or line card that supports PoE has a PoE controller that keeps track of the PoE power consumption on the switch or line card, and allocates power to the PoE ports. The following factors determine how the PoE controller allocates power to the PoE ports:

- [PoE Power Budget on page 4425](#)
- [PoE Interface Power Allocation on page 4429](#)
- [PoE Interface Power Priority on page 4431](#)

### **PoE Power Budget**

The PoE power budget is the total amount of power available for the PoE controller to allocate to the PoE ports. In allocating power, the PoE controller cannot exceed its PoE power budget and does not allocate power to a PoE port if the allocation exceeds the PoE power budget.

How the PoE power budget is determined depends on the switch model:

- [PoE Power Budget on EX2200, EX3200, EX3300, EX4200, EX4300, and EX4600 Switches on page 4425](#)
- [PoE Power Budget on EX6200 and EX8200 Switches on page 4428](#)

### **PoE Power Budget on EX2200, EX3200, EX3300, EX4200, EX4300, and EX4600 Switches**

The PoE power budget on EX2200, EX3200, EX3300, EX4200, and EX4300 switches depends on the switch model and the capacities of the power supplies installed. To find the PoE power budget for each switch model, see [Table 475 on page 4426](#) for EX2200 switch models, [Table 476 on page 4426](#) for EX3200 switch models, [Table 477 on page 4427](#) for EX3300 switch models, [Table 478 on page 4427](#) for EX4200 switch models, and [Table 479 on page 4428](#) for EX4300 switch models.



**NOTE:** EX4600 switches support PoE only in a mixed Virtual Chassis with EX4300 switches. The PoE controller on the EX4600 switch that is in the master role in the Virtual Chassis allocates power to the PoE-enabled ports on the EX4300 member switches by using the power budgets for those EX4300 switches.

Use the [show poe controller](#) command to display a switch's PoE power budget.

If your switch supports power supplies of different capacities, keep the following points in mind:

- If you change your existing power supply to a lower-capacity power supply, the PoE power budget might no longer be sufficient to power all the PoE ports on the switch.

- If your switch supports redundant power supplies and you have installed power supplies of different capacities, the PoE power budget is based on the wattage of the lowest-capacity power supply.
- You cannot increase the number of PoE-capable ports on a switch by installing a power supply that has a higher capacity.

[Table 475 on page 4426](#) lists the EX2200 switch models, number of PoE-enabled ports, power supply ratings, and PoE power budgets.

**Table 475: PoE Power Budget for EX2200 Switches**

| Switch Model Number | Number of PoE-Enabled Ports | Power Supply Rating | PoE Power Budget |
|---------------------|-----------------------------|---------------------|------------------|
| EX2200-C-12T        | –                           | 30 W                | –                |
| EX2200-C-12P        | 12                          | 180 W               | 100 W            |
| EX2200-24T          | –                           | 75 W                | –                |
| EX2200-24P          | 24                          | 550 W               | 405 W            |
| EX2200-24T-DC       | –                           | 100 W               | –                |
| EX2200-48T          | –                           | 75 W                | –                |
| EX2200-48P          | 48                          | 550 W               | 405 W            |

[Table 476 on page 4426](#) lists the EX3200 switch models, number of PoE-enabled ports, power supply ratings, and PoE power budgets.

**Table 476: PoE Power Budget for EX3200 Switch Models**

| Switch Model Number | Number of PoE-Enabled Ports | Power Supply Rating | PoE Power Budget |
|---------------------|-----------------------------|---------------------|------------------|
| EX3200-24T          | 8                           | 320 W               | 130 W            |
| EX3200-48T          | 8                           | 320 W               | 130 W            |
| EX3200-24P          | 24                          | 600 W               | 410 W            |
| EX3200-48P          | 48                          | 930 W               | 740 W            |
| EX3200-24T-DC       | –                           | 190 W               | –                |
| EX3200-48T-DC       | –                           | 190 W               | –                |

[Table 477 on page 4427](#) lists the EX3300 switch models, number of PoE-enabled ports, power supply ratings, and PoE power budgets.

**Table 477: PoE Power Budget EX3300 Switch Models**

| Switch Model Number | Number of PoE-Enabled Ports | Power Supply Rating | PoE Power Budget |
|---------------------|-----------------------------|---------------------|------------------|
| EX3300-24T          | –                           | 100 W               | –                |
| EX3300-24P          | 24                          | 550 W               | 405 W            |
| EX3300-24T-DC       | –                           | 100 W               | –                |
| EX3300-48T          | –                           | 100 W               | –                |
| EX3300-48T-BF       | –                           | 100 W               | –                |
| EX3300-48P          | 48                          | 900 W               | 740 W            |

[Table 478 on page 4427](#) lists the EX4200 switch models, number of PoE-enabled ports, power supply ratings, and PoE power budgets.

**Table 478: PoE Power Budget for EX4200 Switch Models**

| Switch Model Number | Number of PoE-Enabled Ports | Power Supply Rating | PoE Power Budget |
|---------------------|-----------------------------|---------------------|------------------|
| EX4200-24T          | 8                           | 320 W               | 130 W            |
| EX4200-48T          | 8                           | 320 W               | 130 W            |
| EX4200-24P          | 24                          | 600 W               | 410 W            |
| EX4200-48P          | 48                          | 930 W               | 740 W            |
| EX4200-24PX         | 24                          | 930 W               | 740 W            |
| EX4200-48PX         | 48                          | 930 W               | 740 W            |
| EX4200-24F          | –                           | 320 W               | –                |
| EX4200-24F-DC       | –                           | 190 W               | –                |
| EX4200-24T-DC       | –                           | 190 W               | –                |
| EX4200-48T-DC       | –                           | 190 W               | –                |

[Table 479 on page 4428](#) lists the EX4300 switch models, number of PoE-enabled ports, power supply ratings, and PoE power budgets.

Table 479: PoE Power Budget for EX4300 Switch Models

| Switch Model Number | Number of PoE-Enabled Ports | Power Supply Rating | PoE Power Budget |
|---------------------|-----------------------------|---------------------|------------------|
| EX4300-48P          | 48                          | 1100 W              | 900 W            |
| EX4300-48T          | 0                           | 350 W               | -                |
| EX4300-48T-AFI      | 0                           | 350 W               | -                |
| EX4300-24P          | 24                          | 715 W               | 565 W            |
| EX4300-24T          | 0                           | 350 W               | -                |
| EX4300-48T-DC       | 0                           | 550 W               | -                |
| EX4300-48T-DC-AFI   | 0                           | 550 W               | -                |

***PoE Power Budget on EX6200 and EX8200 Switches***

For EX6200 and EX8200 switches, each line card that supports PoE has its own PoE controller and PoE power budget. The PoE power budget is allocated to the line card by the switch's power management, while PoE power is allocated to the ports on the line card by the PoE controller. Because EX6200 and EX8200 switches can differ in the number and capacity of power supplies installed and in the number and types of line cards installed, the amount of power available for PoE power can vary for switches of the same model.

Power management allocates PoE power to line cards that support PoE only after it has allocated base power to and powered on all line cards. It then allocates the remaining power to the PoE power budgets of PoE line cards in order of line card power priority. (In a default configuration, power priority is determined by the line card slot number, with slot 0 having the highest priority.) If the remaining power is insufficient to provide PoE power to all PoE line cards, a low-priority line card might receive no PoE power or partial PoE power.

By default, power management allocates enough PoE power to a line card to power all PoE ports at their maximum supported power. If the powered devices connected to that line card require less power than that, you can configure a smaller PoE power budget for the line card. For example, power management normally allocates 915 W of PoE power to a 48-port PoE+ 20 Gbps (EX8200-48PL) line card. If the powered devices connected to that line card consume no more than a total of 250 W, you can set the PoE power budget for the line card to 250 W. Doing so frees 665 W, which then can be used to fulfill the PoE power needs of lower-priority line cards.

You can also configure the power priority of the PoE ports on a line card. If power management is unable to allocate enough power to a line card to meet its PoE power budget, the line card PoE controller turns off power to PoE ports in reverse priority order as required to meet the reduced power allocation.

Power management adjusts PoE power allocations as power availability and demand in a switch changes. As a general rule, power management allocates power to power on line cards before it allocates PoE power. For example, if you add a line card and there is insufficient power available to power it on, power management reduces the PoE power it provides to line cards, starting with the lowest priority line card, until it frees enough power to power on the new line card. When power management reduces the PoE power budget for a line card because of insufficient power, it logs a message in the system log.

Note that the actual power consumed by the powered devices does not affect power management's power allocation for a line card. If you have set the PoE budget for a line card to 500 W, power management allocates 500 W even if the powered devices are consuming less power than that. Similarly, the PoE power budget is not increased if you add additional powered devices: if the powered devices require more than the 500 W PoE budget that you have configured, lower-priority devices do not receive power.

You can display the switch's power budget maintained by power management, including its PoE power allocations, by using the [show chassis power-budget-statistics](#) command. You can also display the PoE power budget for each line card in a switch by using the [show poe controller](#) command.

For more information about how power management allocates power, including PoE power, see ["Understanding Power Management on EX Series Switches" on page 2495](#).

### **PoE Interface Power Allocation**

You can configure how the switch determines the maximum power for a PoE interface and how power is allocated to the PoE interfaces. If the power consumption of a powered device exceeds the maximum power allocated to the interface, the switch turns off power to the interface.

These PoE power allocation methods are available:

- [LLDP Power Negotiation on page 4429](#)
- [Class PoE Management Mode on page 4430](#)
- [Static PoE Management Mode on page 4430](#)

### **LLDP Power Negotiation**

Link Layer Discovery Protocol (LLDP) power negotiation enables the PoE controller to dynamically allocate power to LLDP-enabled powered devices based on their power needs. The PoE controller allocates to an interface only the power currently required by the connected powered device, and it can allocate the power in small increments.

When the PoE **management** is set to **class** and LLDP is enabled (both are default settings), LLDP power negotiation is enabled by default. If you disable LLDP power negotiation or the powered device does not support it, the switch uses the class of the powered device to determine the maximum power for interfaces.



**NOTE:** LLDP power negotiation is not supported on EX3200 and EX4200 (except EX4200 PX models) switches.

### Class PoE Management Mode

In the **class** PoE management mode, the maximum power for an interface is determined by the class of the connected powered device. The PoE standards IEEE 802.3af and IEEE 802.3at define classes of powered devices based on the levels of power that they require. [Table 480 on page 4430](#) lists the classes of powered devices and associated power levels.

**Table 480: Class of Powered Device and Power Levels**

| Standard                                   | Class | Maximum Power Delivered by PoE Port | Power Range of Powered Device |
|--------------------------------------------|-------|-------------------------------------|-------------------------------|
| IEEE 802.3af (PoE) and IEEE 802.3at (PoE+) | 0     | 15.4 W                              | 0.44 through 12.95 W          |
|                                            | 1     | 4.0 W                               | 0.44 through 3.84 W           |
|                                            | 2     | 7.0 W                               | 3.84 through 6.49 W           |
|                                            | 3     | 15.4 W                              | 6.49 through 12.95 W          |
| IEEE 802.3at (PoE+)                        | 4     | 30.0 W                              | 12.95 through 25.5 W          |

Because of line loss, the power range of the powered device is less than the maximum power delivered at the PoE port for each class. Line loss is influenced by cable length, quality, and other factors and is typically less than 16 percent of the maximum power.

The powered device communicates to the PoE controller which class it belongs to when it is connected. The PoE controller then allocates to the interface the maximum power required by the class (see [Table 480 on page 4430](#)). It does not allocate power to an interface until a powered device is connected. **Class 0** is the default class for powered devices that do not provide class information. Class 4 powered devices are supported only by PoE ports that support IEEE 802.3at (PoE+).

By default, when the **management** option is set to **class** and LLDP are enabled, LLDP power negotiation is also enabled on supported switches. See [“LLDP Power Negotiation” on page 4429](#) for more information.

### Static PoE Management Mode

In the **static** PoE management mode, you specify the maximum power for each PoE interface. The PoE controller then allocates this amount of power to the interface from its total budget. For example, if you specify a maximum value of 8.0 W for ge-0/0/3, the PoE controller allocates 8.0 W out of its total power budget for this interface. This amount is allocated to the interface irrespective of whether a powered device is connected to the interface or the connected powered device uses less power than 8.0 W.

Because of line loss, the power received by the powered device can be less than the power available at the PoE port. [Table 481 on page 4431](#) shows the maximum power available at a PoE port and the resulting power guaranteed to the powered device.

Table 481: Maximum Power per Port in Static Mode

| Switch or Line Card                                                                                                                   | Maximum Power Delivered by PoE Port                                                                                                                                    | Guaranteed Power to Powered Devices |
|---------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|
| EX2200 switches, EX3300 switches, EX4200 PX model switches, EX4300 switches, and EX4600 switches operating in a mixed Virtual Chassis | 30 W                                                                                                                                                                   | 25.5 W                              |
| EX3200 switches and EX4200 P and T model switches running Junos OS Release 10.4 or earlier                                            | 15.4 W                                                                                                                                                                 | 12.95 W                             |
| EX3200 switches and EX4200 P and T model switches running Junos OS Release 11.1 or later                                              | 18.6 W<br><i>NOTE:</i> Switches that are upgraded to Junos OS Release 11.1 from a previous release require an upgrade of the PoE controller software to obtain 18.6 W. | 15.64 W                             |
| EX6200-48P line cards                                                                                                                 | 30 W                                                                                                                                                                   | 25.5 W                              |
| EX8200-2XS-40P line cards and EX8200-48PL line cards                                                                                  | 30 W (ports 0 through 11)<br>15.4 W (remaining PoE ports)                                                                                                              | 25.5 W<br>12.95 W                   |

**PoE Interface Power Priority**

You can configure a PoE interface to have a power priority. The power priority determines which interfaces receive power if PoE power demands are greater than the PoE power budget. If the total power allocated for all interfaces exceeds the switch budget, PoE power to lower-priority interfaces is turned off and the power allocated to those interfaces drops to 0. Thus you must set interfaces that connect to critical powered devices, such as security cameras and emergency phones, to high priority.

Among PoE interfaces that have the same assigned priority, power priority is determined by the port number, with lower-numbered ports having higher priority.

For EX6200 and EX8200 switches, interface power priority determines the relative priority of the interfaces on a line card, not on the switch as a whole. The relative priority of interfaces residing on different line cards is determined by line card priority. For example, if line card 1 has a higher power priority than line card 2 and a power shortage occurs, power is removed from the PoE interfaces in this order:

- Low-priority interfaces on line card 2
- High-priority interfaces on line card 2
- Low-priority interfaces on line card 1
- High-priority interfaces on line card 1

You can manually configure PoE interface power priority, or you can enable LLDP power priority, which assigns each interface the power priority provided by the connected LLDP-enabled powered device. [Table 482 on page 4432](#) describes how the switch converts LLDP power priorities to switch power priorities.

Table 482: LLDP Power Priority Conversion

| LLDP Power Priority | Switch Power Priority |
|---------------------|-----------------------|
| Critical, High      | High                  |
| Low                 | Low                   |



NOTE: LLDP power priority requires LLDP power negotiation to be enabled, which is enabled by default when the PoE management option is set to class.



NOTE: LLDP power priority is not supported on EX3200 and EX4200 (except EX4200 PX model) switches.

Related Documentation

- [Example: Configuring PoE Interfaces on an EX Series Switch on page 4433](#)
- [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)
- [Example: Configuring PoE on an EX6200 or EX8200 Switch](#)
- [Upgrading the PoE Controller Software on page 4467](#)



## CHAPTER 71

# Configuration

- [Configuration Examples on page 4433](#)
- [Configuration Tasks on page 4440](#)
- [Configuration Statements on page 4450](#)

### Configuration Examples

---

- [Example: Configuring PoE Interfaces on an EX Series Switch on page 4433](#)
- [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)

#### Example: Configuring PoE Interfaces on an EX Series Switch

Power over Ethernet (PoE) ports supply electric power over the same ports that are used to connect network devices and enable you to plug in devices that require both network connectivity and electric power, such as VoIP phones, wireless access points, and some IP cameras. This reduces the amount of wiring in a network, and also eliminates the need to position a powered device near an AC power outlet, making network design more flexible and efficient.

You do not need to configure PoE unless you want to modify the default values or disable PoE on a specific interface.

This example describes a default configuration of PoE interfaces on an EX Series switch:

- [Requirements on page 4433](#)
- [Overview and Topology on page 4434](#)
- [Configuration on page 4434](#)
- [Verification on page 4435](#)

#### Requirements

---

This example uses the following hardware and software components:

- One EX Series switch that supports PoE



**NOTE:** EX4600 switches support PoE only when operating in a mixed Virtual Chassis with EX4300 switches.

- Avaya IP telephones
- Wireless access point
- Junos OS Release 9.0 or later for EX Series switches

Before you configure PoE, be sure you have:

- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* or *Connecting and Configuring an EX Series Switch (J-Web Procedure)* for details.

### Overview and Topology

The topology used in this example consists of a switch that has 24 ports. Eight of the ports support PoE (IEEE 802.3af), which means they provide both network connectivity and electric power for powered devices such as VoIP telephones, wireless access points, and IP security cameras that require 12.95 W or less. The remaining 16 ports provide only network connectivity. You use the standard ports to connect devices that have their own power sources, such as desktop and laptop computers, printers, and servers.

[Table 483 on page 4434](#) details the topology used in this configuration example.

**Table 483: Components of the PoE Configuration Topology**

| Property                                                                                                                           | Settings                                                                                                                                             |
|------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch hardware                                                                                                                    | EX Series switch with 24 Gigabit Ethernet ports: 8 PoE interfaces (ge-0/0/0 through ge-0/0/7) and 16 non-PoE interfaces (ge-0/0/8 through ge-0/0/23) |
| VLAN name                                                                                                                          | default                                                                                                                                              |
| Connection to a wireless access point (requires PoE)                                                                               | ge-0/0/0                                                                                                                                             |
| Connections to Avaya IP telephones with integrated hubs that allow phone and desktop PC to connect to a single port (requires PoE) | ge-0/0/1 through ge-0/0/7                                                                                                                            |
| Direct connections to desktop PCs, file servers, integrated printer/fax/copier machines (no PoE required)                          | ge-0/0/8 through ge-0/0/20                                                                                                                           |
| Unused ports (for future expansion)                                                                                                | ge-0/0/21 through ge-0/0/23                                                                                                                          |

### Configuration

To enable the default PoE configuration on the switch:

#### CLI Quick Configuration

To quickly enable the default configuration on the switch:

Simply connect the powered devices to the PoE ports.

#### Step-by-Step Procedure

To use the PoE interfaces with default values:

1. Make sure the switch is powered on.
2. Connect the wireless access point to interface ge-0/0/0.
3. Connect the Avaya phones to interfaces ge-0/0/1 through ge-0/0/7.

#### Verification

To verify that PoE interfaces have been created and are operational, perform this task:

- [Verifying That the PoE Interfaces Have Been Created on page 4435](#)

#### *Verifying That the PoE Interfaces Have Been Created*

**Purpose** Verify that the PoE interfaces have been created on the switch.

**Action** List all the PoE interfaces configured on the switch:

```
user@switch> show poe interface
```

| Interface | Admin status | Oper status | Max power | Priority | Power consumption | Class |
|-----------|--------------|-------------|-----------|----------|-------------------|-------|
| ge-0/0/0  | Enabled      | ON          | 15.4W     | Low      | 7.9W              | 0     |
| ge-0/0/1  | Enabled      | ON          | 15.4W     | Low      | 3.2W              | 2     |
| ge-0/0/2  | Enabled      | ON          | 15.4W     | Low      | 3.2W              | 2     |
| ge-0/0/3  | Enabled      | ON          | 15.4W     | Low      | 3.2W              | 2     |
| ge-0/0/4  | Enabled      | ON          | 15.4W     | Low      | 3.2W              | 2     |
| ge-0/0/5  | Enabled      | ON          | 15.4W     | Low      | 3.2W              | 2     |
| ge-0/0/6  | Enabled      | ON          | 15.4W     | Low      | 3.2W              | 2     |
| ge-0/0/7  | Enabled      | ON          | 15.4W     | Low      | 3.2W              | 2     |

**Meaning** The `show poe interface` command lists PoE interfaces configured on the switch, with their status, priority, power consumption, and class. This output shows that eight interfaces have been created with default values and are consuming power at the expected rates.

#### Related Documentation

- [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)
- [Configuring PoE on EX Series Switches \(CLI Procedure\) on page 4440](#)
- [Troubleshooting PoE Interfaces on page 4493](#)

### Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch

Power over Ethernet (PoE) ports supply electric power over the same ports that are used to connect network devices. These ports enable you to plug in devices that need both network connectivity and electric power, such as VoIP phones, wireless access points, and some IP cameras.

By default, PoE ports on EX Series switches are set to low power priority. You can configure a PoE port to have a high power priority setting. If a situation arises where there is not

sufficient power for all the PoE ports, the available power is directed to the higher priority ports, while power to the lower priority ports is shut down as needed. Thus you must set ports that connect to security cameras, emergency phones, and other high priority powered devices to high-priority.

This example describes how to configure a few high-priority PoE interfaces.

- [Requirements on page 4436](#)
- [Overview and Topology on page 4436](#)
- [Configuration on page 4437](#)
- [Verification on page 4439](#)

### Requirements

This example uses the following hardware and software components:

- One EX Series switch that supports PoE



**NOTE:** EX4600 switches support PoE only when operating in a mixed Virtual Chassis with EX4300 switches.

- Powered devices—wireless access point, VoIP telephones, and IP security cameras—that require PoE
- Junos OS Release 9.0 or later for EX Series switches

Before you configure PoE, be sure you have:

- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* or *Connecting and Configuring an EX Series Switch (J-Web Procedure)* for details.

### Overview and Topology

The topology used in this example consists of a switch that has 24 ports. Eight of the ports support PoE (IEEE 802.3af), which means they provide both network connectivity and electric power for powered devices such as VoIP telephones, wireless access points, and IP security cameras that require 12.95 W or less. The remaining 16 ports provide only network connectivity. You use the standard ports to connect devices that have their own power sources, such as desktop and laptop computers, printers, and servers.

[Table 484 on page 4436](#) details the topology used in this configuration example.

**Table 484: Components of the PoE Configuration Topology**

| Property        | Settings                                                                                                                                   |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| Switch hardware | Switch with 24 Gigabit Ethernet ports: 8 PoE interfaces (ge-0/0/0 through ge-0/0/7) and 16 non-PoE interfaces (ge-0/0/8 through ge-0/0/23) |
| VLAN name       | default                                                                                                                                    |

Table 484: Components of the PoE Configuration Topology (*continued*)

| Property                                                                                                  | Settings                    |
|-----------------------------------------------------------------------------------------------------------|-----------------------------|
| Connection to a wireless access point (requires PoE)                                                      | ge-0/0/0                    |
| Security IP Cameras (require PoE)                                                                         | ge-0/0/1 and ge-0/0/2 high  |
| Emergency VoIP phone (requires PoE)                                                                       | ge-0/0/3 high               |
| VoIP phone in Executive Office (requires PoE)                                                             | ge-0/0/4 high               |
| Other VoIP phones (require PoE)                                                                           | ge-0/0/5 through ge-0/0/7   |
| Direct connections to desktop PCs, file servers, integrated printer/fax/copier machines (no PoE required) | ge-0/0/8 through ge-0/0/20  |
| Unused ports (for future expansion)                                                                       | ge-0/0/21 through ge-0/0/23 |

### Configuration

To configure PoE interfaces:

#### CLI Quick Configuration

By default, PoE interfaces are created for all PoE ports and PoE is enabled. The default priority for PoE interfaces is **low**.

To quickly set some interfaces to high priority and to include descriptions of the interfaces, copy the following commands and paste them into the switch terminal window:

```
[edit]
set poe interface ge-0/0/1 priority high telemetries
set poe interface ge-0/0/2 priority high telemetries
set poe interface ge-0/0/3 priority high telemetries
set poe interface ge-0/0/4 priority high telemetries
set interfaces ge-0/0/0 description "wireless access point"
set interfaces ge-0/0/1 description "security camera front door"
set interfaces ge-0/0/2 description "security camera back door"
set interfaces ge-0/0/3 description "emergency phone"
set interfaces ge-0/0/4 description "Executive Office VoIP phone"
set interfaces ge-0/0/5 description "staff VoIP phone"
set interfaces ge-0/0/6 description "staff VoIP phone"
set interfaces ge-0/0/7 description "staff VoIP phone"
```

#### Step-by-Step Procedure

To configure PoE interfaces with different priorities:

1. Set the interfaces connected to high-priority powered devices to high priority. Include the **telemetries** statement for the high-priority interfaces, thus enabling the logging of power consumption on those interfaces:

```
[edit poe]
user@switch# set interface ge-0/0/1 priority high telemetries
user@switch# set interface ge-0/0/2 priority high telemetries
user@switch# set interface ge-0/0/3 priority high telemetries
user@switch# set interface ge-0/0/4 priority high telemetries
```

2. Provide descriptions for the PoE interfaces:

```
[edit interfaces]
```

```
user@switch# set ge-0/0/0 description "wireless access point"
user@switch# set ge-0/0/1 description "security camera front door"
user@switch# set ge-0/0/2 description "security camera back door"
user@switch# set ge-0/0/3 description "emergency phone"
user@switch# set ge-0/0/4 description "Executive Office VoIP phone"
user@switch# set ge-0/0/5 description "staff VoIP phone"
user@switch# set ge-0/0/6 description "staff VoIP phone"
user@switch# set ge-0/0/7 description "staff VoIP phone"
```

3. Connect the wireless access point to interface ge-0/0/0. This interface uses the default PoE settings.
4. Connect the two security cameras to interfaces ge-0/0/1 and ge-0/0/2. These interfaces are set to high priority with telemetries enabled.
5. Connect the emergency VoIP phone to interface ge-0/0/3. This interface is set to high priority with telemetries enabled.
6. Connect the Executive Office VoIP phone to interface ge-0/0/4. This interface is set to high priority with telemetries enabled.
7. Connect the staff VoIP phones to ge-0/0/5, ge-0/0/6, and ge-0/0/7. These interfaces use the default PoE settings.

### Results

Check the results of the configuration:

```
[edit]
user@switch# show
interfaces {
  ge-0/0/0 {
    description "wireless access point";
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/1 {
    description "security camera front door";
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/2 {
    description "security camera back door";
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/3 {
    description "emergency phone";
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/4 {
```

```

        description "Executive Office VoIP phone";
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/5 {
        description "staff VoIP phone";
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/6 {
        description "staff VoIP phone";
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/7 {
        description "staff VoIP phone";
        unit 0 {
            family ethernet-switching;
        }
    }
}
poe {
    interface all;
    interface ge-0/0/1 {
        priority high;
        telemetries;
    }
    interface ge-0/0/2 {
        priority high;
        telemetries;
    }
    interface ge-0/0/3 {
        priority high;
        telemetries;
    }
    interface ge-0/0/4 {
        priority high;
        telemetries;
    }
}

```

### Verification

To verify that PoE interfaces have been created and are operational, perform the following tasks:

- [Verifying That the PoE Interfaces Have Been Created with the Correct Priorities on page 4439](#)

#### *Verifying That the PoE Interfaces Have Been Created with the Correct Priorities*

**Purpose** Verify that the PoE interfaces on the switch are now set to the correct priority settings.

**Action** List all the PoE interfaces configured on the switch:

```
user@switch> show poe interface
```

| Interface | Admin<br>status | Oper<br>status | Max<br>power | Priority | Power<br>consumption | Class |
|-----------|-----------------|----------------|--------------|----------|----------------------|-------|
| ge-0/0/0  | Enabled         | ON             | 15.4W        | Low      | 7.9W                 | 0     |
| ge-0/0/1  | Enabled         | ON             | 15.4W        | High     | 4.8W                 | 0     |
| ge-0/0/2  | Enabled         | ON             | 15.4W        | High     | 4.8W                 | 0     |
| ge-0/0/3  | Enabled         | ON             | 15.4W        | High     | 3.3W                 | 2     |
| ge-0/0/4  | Enabled         | ON             | 15.4W        | High     | 4.7W                 | 2     |
| ge-0/0/5  | Enabled         | ON             | 15.4W        | Low      | 3.2W                 | 2     |
| ge-0/0/6  | Enabled         | ON             | 15.4W        | Low      | 3.3W                 | 2     |
| ge-0/0/7  | Enabled         | ON             | 15.4W        | Low      | 3.3W                 | 2     |

**Meaning** The **show poe interface** command lists PoE interfaces configured on the switch, with their status, priority, power consumption, and class. This output shows that eight PoE interfaces are enabled. Interfaces ge-0/0/1 through ge-0/0/4 are configured as priority **high**. The remaining PoE interfaces are configured with the default priority value of **low**.

- Related Documentation**
- [Example: Configuring PoE Interfaces on an EX Series Switch on page 4433](#)
  - [Configuring PoE on EX Series Switches \(CLI Procedure\) on page 4440](#)
  - [Troubleshooting PoE Interfaces on page 4493](#)

---

## Configuration Tasks

- [Configuring PoE on EX Series Switches \(CLI Procedure\) on page 4440](#)
- [Configuring PoE \(J-Web Procedure\) on page 4447](#)

### Configuring PoE on EX Series Switches (CLI Procedure)

Power over Ethernet (PoE) ports on EX Series switches supply electric power over the same ports that are used to connect network devices. These ports enable you to plug in devices that require both network connectivity and electric power, such as VoIP phones, wireless access points, and some IP cameras. This reduces the amount of wiring in a network, and also eliminates the need to position a powered device near an AC power outlet, making network design more flexible and efficient.

This topic describes:

- [PoE Configurable Options on page 4441](#)
- [Configuring the PoE Controller on EX2200, EX3200, EX3300, EX4200, and EX4300 Switches on page 4443](#)
- [Configuring the PoE Controllers on EX6200 and EX8200 Switches on page 4445](#)
- [Configuring PoE Interfaces on page 4446](#)



## PoE Configurable Options

For EX Series switches that support PoE ports, the factory default configuration enables PoE on the PoE-capable ports, with default settings in effect. You might not have to do any additional configuration if the default settings work for you. [Table 485 on page 4441](#) shows the configurable PoE options and their default settings for the PoE controller and for the PoE interfaces.



**NOTE:** On EX8200 switches, the factory default configuration enables PoE on all interfaces starting at Junos OS Release 11.2. Switches that have been upgraded to Release 11.2 from an earlier release might not have PoE enabled by default. To enable PoE on all PoE-capable ports on a switch, use the `set poe interface all` configuration command.



**NOTE:** EX4600 switches support PoE only in a mixed Virtual Chassis with EX4300 switches. EX4600 switches do not have PoE ports; therefore, the factory default configuration does not enable PoE.

**Table 485: Configurable PoE Options and Default Settings**

| Option                     | Default                               | Description                                                                                                                                                                                                                                                                                                 |
|----------------------------|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PoE Controller Options     |                                       |                                                                                                                                                                                                                                                                                                             |
| <code>guard-band</code>    | 0 W                                   | Reserves a specified amount of power from the PoE power budget to be used in the case of a spike in PoE power consumption: <ul style="list-style-type: none"> <li>Up to 15 W on EX6200 and EX8200 switches</li> <li>Up to 19 W on all other switches</li> </ul>                                             |
| <code>lldp-priority</code> | Not included in default configuration | When included in the configuration, assigns interfaces the power priority provided by the connected powered device by using Link Layer Discovery Protocol (LLDP) power negotiation rather than the power priority configured on the switch interface. <p>Requires LLDP power negotiation to be enabled.</p> |

Table 485: Configurable PoE Options and Default Settings (*continued*)

| Option                               | Default                                                                                                                                                                                                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>management</b>                    | <b>class</b>                                                                                                                                                                                                             | <p>Sets the PoE power management mode for the switch or line card. The power management mode determines how power to a PoE interface is allocated:</p> <ul style="list-style-type: none"> <li>• <b>class</b>—In this mode, the power allocated to a PoE interface is determined in one of two ways: <ul style="list-style-type: none"> <li>• If LLDP power negotiation is enabled, the PoE controller allocates PoE power by using LLDP power negotiation, which enables the PoE controller to dynamically allocate power to LLDP-enabled devices based on their power needs. LLDP power negotiation is enabled by default on supported switches when the <b>management</b> option is set to <b>class</b>. For information about configuring LLDP power negotiation, see <a href="#">“Configuring LLDP (CLI Procedure)” on page 1913</a>.</li> <li>• If LLDP power negotiation is disabled or not supported on the powered device or the switch, the maximum power delivered by an interface is determined by the class of the connected powered device. If there is no powered device connected, standard 15.4W power is allocated on the interface.</li> </ul> </li> <li>• <b>static</b>—The maximum power delivered by an interface is statically configured and is independent of the class of the connected powered device. The maximum power is allocated to the interface even if a powered device is not connected.</li> </ul> |
| <b>maximum-power</b>                 | <p>792 W for the EX8200-2XS-40P (40-port PoE+ with 4-port SFP and 2-port SFP+) line card</p> <p>915 W for the EX8200-48PL (48-port PoE+ 20 Gbps) line card</p> <p>1440 W for the EX6200-48P (48-port PoE+) line card</p> | <p>(EX6200 and EX8200 switches only) Sets the PoE power budget for the line card:</p> <ul style="list-style-type: none"> <li>• 37 W through 792 W for the EX8200-2XS-40P line card</li> <li>• 37 W through 915 W for the EX8200-48PL line card</li> <li>• 37 W through 1440 W for the EX6200-48P line card</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>notification-control</b>          | Not included in default configuration                                                                                                                                                                                    | When included in the configuration, enables the PoE controller to send PoE SNMP traps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Interface Options                    |                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>af-mode</b>                       | Not included in default configuration                                                                                                                                                                                    | (EX6200 switches only) When included in the configuration, restricts a PoE interface to supporting IEEE 802.3af only. The maximum power that can be delivered by the PoE interface is 15.4 W.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>disable (Power over Ethernet)</b> | Not included in default configuration                                                                                                                                                                                    | When included in the configuration, disables PoE on the interface. The interface maintains network connectivity but no longer supplies power to a connected powered device. Power is not allocated to the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 485: Configurable PoE Options and Default Settings (*continued*)

| Option                                | Default                                                                                                                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>maximum-power (Interface)</b>      | <p><b>30.0 W</b> for interfaces that support PoE+ (IEEE 802.3at)</p> <p><b>15.4 W</b> for interfaces that support PoE (IEEE 802.3af)</p> | <p>Sets the maximum power that can be delivered by a PoE interface when the power management mode is <b>static</b>:</p> <ul style="list-style-type: none"> <li>Up to 30 W for EX2200, EX3300, EX4200, EX4300, EX6200, and EX8200 switches</li> <li>Up to 18.6 W for EX3200 switches</li> </ul> <p>This setting is ignored if the power management mode is <b>class</b>.</p> <p><b>NOTE:</b> The <b>maximum-power</b> setting permitted by the CLI might be greater than the maximum power a given PoE port can deliver. For example, the CLI permits you to set any port on an EX8200 line card to 30 W; however, only ports 0 through 11 support 30 W. Similarly, the CLI permits you to set any port on an EX4200 switch to 30 W, but some EX4200 models support only 18.6 W per port. If you configure a <b>maximum-power</b> value that is greater than the maximum power supported by a port, the power allocated to the port will be the maximum supported.</p> |
| <b>priority (Power over Ethernet)</b> | <b>low</b>                                                                                                                               | <p>Sets an interface's power priority to either <b>low</b> or <b>high</b>. If power is insufficient for all PoE interfaces, the PoE power to low-priority interfaces is shut down before power to high-priority interfaces is shut down. Among interfaces that have the same assigned priority, the power priority is determined by port number, with lower-numbered ports having higher priority.</p> <p>If LLDP power priority is enabled, the switch assigns each interface the power priority provided by the connected LLDP-enabled powered device rather than the interface's configured priority.</p> <p>On EX6200 and EX8200 switches, <b>priority</b> determines the interface's power priority relative to the other interfaces on the line card, not the interfaces on the switch as a whole. If power management cannot provide the line card with its full PoE power budget, PoE power to interfaces with low priority is shut down first.</p>           |
| <b>telemetries</b>                    | Not included in default configuration                                                                                                    | <p>When included in the configuration, enables the logging of power consumption records on an interface. Logging occurs every 5 minutes for 1 hour unless you specify a different value for <b>interval (Power over Ethernet)</b> or <b>duration</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

### Configuring the PoE Controller on EX2200, EX3200, EX3300, EX4200, and EX4300 Switches

To configure the PoE controller on EX2200, EX3200, EX3300, EX4200, EX4300, and EX4600 switches:

- To change the management mode or to configure a guard band setting for a standalone switch or for all members of an EX3300 Virtual Chassis, an EX4200 Virtual Chassis, an EX4300 Virtual Chassis, a mixed EX4200 and EX4500 Virtual Chassis that supports PoE, or a mixed EX4300 and EX4600 Virtual Chassis that supports PoE:

```
[edit]
user@switch# set poe management mode guard-band watts
```

For example, to set the management mode to static and to configure a guard band of 15 W:

```
[edit]
user@switch# set poe management static guard-band 15
```



**NOTE:** If the PoE power budget for the switch is insufficient to provide maximum power to all the PoE ports, we recommend that you do not change the management mode from class to static. If you change the power management mode to static and do not change the other default settings, the PoE controller allocates maximum power to the PoE ports in the order of port number, which means PoE will be disabled on higher-numbered ports when the PoE power budget runs out.

In class mode, on the other hand, the PoE controller does not allocate power to a port until a powered device is connected. The class of the connected device determines the amount of power allocated. Thus in class mode, any PoE port can be used to power a device and all the PoE ports on the switch can be used as long as the combined power demand does not exceed the PoE power budget.



**NOTE:** On EX3200 and EX4200 switches that support enhanced PoE, you must change the management mode from class to static to take advantage of the higher per-port power limits of enhanced PoE.

- To enable PoE SNMP traps on a standalone switch or on a specific member of a Virtual Chassis:

```
[edit]
user@switch# set poe notification-control fpc number
```

For example, to enable PoE SNMP traps on a standalone switch or on member 0 of a Virtual Chassis:

```
[edit]
user@switch# set poe notification-control fpc 0
```

### Configuring the PoE Controllers on EX6200 and EX8200 Switches

On EX6200 and EX8200 switches, each line card that supports PoE has its own PoE controller. This means that the PoE controller options are configured separately for each line card.

In addition, each line card has its own separate, configurable PoE power budget. The default power budget for a line card is the amount of power needed to supply all PoE ports on the line card with their maximum supported power. Because there might not be enough power available in a switch to supply each PoE line card with the default PoE power budget, you can configure smaller power budgets for one or more line cards, freeing power for other line cards.

To configure the line card PoE controllers in an EX6200 or EX8200 switch:

- To configure a guard band setting, to change the management mode, or to configure the PoE power budget for a specific line card:

```
[edit]
user@switch# set poe fpc number guard-band watts management mode
maximum-power watts
```

For example, to configure a PoE budget of 350 W and a guard band of 15 W on line card 1:

```
[edit]
user@switch# set poe fpc 1 guard-band 15 maximum-power 350
```



**NOTE:** If you configure a PoE power budget for a line card that is smaller than the default power budget, we recommend that you do not change the management mode from class to static. If you change the power management mode to static and do not change the interface default settings, the PoE controller allocates maximum power to the PoE ports in the order of port number. As a result, PoE will be disabled on higher-numbered ports when the PoE power budget runs out.

In class mode, on the other hand, the PoE controller does not allocate power to a port until a powered device is connected. The class of the connected device determines the amount of power allocated. Thus in class mode, any PoE port can be used to power a device and all the PoE ports on the switch can be used as long as the combined power demand does not exceed the PoE power budget.

- To configure the same guard band value, management mode, or PoE power budget for all line cards in a switch:

```
[edit]
user@switch# set poe fpc all guard-band watts management mode maximum-power
watts
```

For example, to configure a PoE budget of 1000 W and static management mode for all line cards in a switch:

```
[edit]
user@switch# set poe fpc all management static maximum-power 1000
```

If you configure different settings for a specific line card, those settings override the settings configured with the **fpc all** statement.

- To enable PoE SNMP traps on a line card:

```
[edit]
user@switch# set poe notification-control fpc number
```

For example, to enable PoE SNMP traps on line card 7:

```
[edit]
user@switch# set poe notification-control fpc 7
```

## Configuring PoE Interfaces

---

To configure the PoE interfaces on a switch that supports PoE:

- To configure all PoE interfaces with the same setting or settings:

```
[edit]
user@switch# set poe interface all options
```

For example, to enable telemetry collection on all interfaces, using the default collection duration and interval:

```
[edit]
user@switch# set poe interface all telemetries
```



**NOTE:** For PoE to be enabled on all PoE-capable interfaces, the configuration must include the **interface all** statement in the **[edit poe]** hierarchy level. With the exception of EX8200 switches that were shipped from the factory with a Junos OS release earlier than Release 11.2, the factory default configurations of switches that support PoE include this statement.

- To configure individual PoE interfaces with different settings:

```
[edit]
user@switch# set poe interface interface-name options
```

For example:

```
[edit]
user@switch# set poe interface ge-0/0/0 priority high telemetries duration 24
```

```
[edit]
user@switch# set poe interface ge-0/0/1
```

```
[edit]
user@switch# set poe interface ge-0/0/5 maximum-power 18.6
```

```
[edit]
user@switch# set poe interface ge-5/0/7 disable
```

When you configure an individual interface, its configuration overrides any settings you configure with the **set poe interface all** command. For example, ge-0/0/1 in the preceding example retains the default settings, regardless of any settings configured with the **set poe interface all** command.

#### Related Documentation

- [Example: Configuring PoE Interfaces on an EX Series Switch on page 4433](#)
- [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)
- [Example: Configuring PoE on an EX6200 or EX8200 Switch](#)
- [Verifying PoE Configuration and Status \(CLI Procedure\) on page 4474](#)
- [Understanding PoE on EX Series Switches on page 4423](#)

## Configuring PoE (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

Power over Ethernet (PoE) ports supply electric power over the same ports that are used to connect network devices to EX Series switches. These ports allow you to plug in devices that require both network connectivity and electric power, such as VoIP phones, wireless access points, and some IP cameras. Using the Power over Ethernet (PoE) Configuration page in the J-Web interface, you can modify the settings of all interfaces that are PoE-enabled.

This topic includes:

- [Configuring PoE on EX2200, EX2200-C, EX3200, EX3300, EX4200, and EX4300 Switches on page 4447](#)
- [Configuring PoE on EX6200 Switches on page 4448](#)

### Configuring PoE on EX2200, EX2200-C, EX3200, EX3300, EX4200, and EX4300 Switches

To configure PoE:

1. Select **Configure > Power over Ethernet**.

The page displays a list of all PoE-capable interfaces except uplink ports. Specific operational details about an interface are displayed in the Details section of the page. The details include the PoE operational status and port class.



**NOTE:** If you are configuring a Virtual Chassis, the PoE configuring option is displayed if any member of the Virtual Chassis supports PoE, even if the Virtual Chassis master does not support PoE.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Edit**—Changes PoE settings for the selected port as described in [Table 486 on page 4448](#).
- **System Settings**—Modifies general PoE settings as described in [Table 487 on page 4448](#).

**Table 486: PoE Edit Settings**

| Field         | Description                                                                                      | Your Action                                                                                                             |
|---------------|--------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| Enable PoE    | Specifies that PoE is enabled on the interface.                                                  | Select this option to enable PoE or PoE+ on the interface.                                                              |
| Priority      | Lists the power priority (low or high) configured on the interface enabled for PoE.              | Set the priority as <b>High</b> or <b>Low</b> .                                                                         |
| Maximum Power | Specifies the maximum PoE wattage available to provision the active PoE interface on the switch. | Select a value in watts. If no value is specified, the default is 15.4 for PoE interfaces and 30.0 for PoE+ interfaces. |

**Table 487: System Settings**

| Field              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                | Your Action                                                                                                       |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| PoE Management     | Specifies the power management mode. The options are: <b>static</b> and <b>class</b> .<br><br><b>NOTE:</b> When the power management mode is set to <b>class</b> , the maximum power value is overridden by the maximum power value of the class of the powered device that is connected to the switch on the PoE port. When the power management mode is set to <b>static</b> , you can specify the maximum power for each PoE interface. | By default, the power management mode is <b>class</b> . Select <b>static</b> to change the power management mode. |
| Guard Band (watts) | Specifies the amount of power reserved for power spikes from the PoE power budget of the switch.                                                                                                                                                                                                                                                                                                                                           | Enter a value to set the guard band value in watts. The default value is 0.                                       |

### Configuring PoE on EX6200 Switches

To configure PoE:

1. Select **Configure > Power over Ethernet**.

The page displays a list of all PoE-capable interfaces for each FPC. Specific operational details about an interface are displayed in the Details section of the page. The details include the PoE operational status and port class.





**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Edit**—Changes PoE settings for the selected port as described in [Table 488 on page 4449](#).

- **FPC Settings**—Changes PoE settings of PoE-capable FPCs.

To configure FPC settings, click one of the following options:

- **Add**—Adds a PoE setting for an FPC as described in [Table 489 on page 4449](#).
- **Edit**—Modifies a PoE setting for an FPC as described in [Table 489 on page 4449](#).
- **Delete**—Deletes an existing PoE settings for an FPC.

**Table 488: Edit PoE Settings**

| Field         | Description                                                                              | Your Action                                                                                                             |
|---------------|------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| Enable PoE    | Specifies that PoE is enabled on the interface.                                          | Select this option to enable PoE or PoE+ on the interface.                                                              |
| Type          | Specifies whether the interface is PoE or PoE+.                                          | Select an option from the list.                                                                                         |
| Priority      | Lists the power priority (low or high) configured on the interface enabled for PoE.      | Set the priority as <b>High</b> or <b>Low</b> .                                                                         |
| Maximum Power | Specifies the maximum PoE wattage available to provision active PoE ports on the switch. | Select a value in watts. If no value is specified, the default is 15.4 for PoE interfaces and 30.0 for PoE+ interfaces. |

**Table 489: FPC PoE Settings**

| Field          | Description                                                                                                                                                                                                                                                                                                                                                                                                                  | Your Action                                                                                                       |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| FPC            | Specifies the FPC number.                                                                                                                                                                                                                                                                                                                                                                                                    | Select a value from the list.                                                                                     |
| PoE Management | <p>Specifies the power management mode. The options are <b>static</b> and <b>class</b>.</p> <p><b>NOTE:</b> When the power management mode is set to <b>class</b>, the maximum power value is overridden by the maximum power value for the interface that is connected to the switch on the PoE port. When the power management mode is set to <b>static</b>, you can specify the maximum power for each PoE interface.</p> | By default, the power management mode is <b>class</b> . Select <b>static</b> to change the power management mode. |

Table 489: FPC PoE Settings (*continued*)

| Field              | Description                                                                                                                                                                                                         | Your Action                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------|
| Guard Band (watts) | Specifies the amount of power reserved for power spikes from the PoE power budget of the switch.                                                                                                                    | Enter a value to set the guard band value in watts. The default value is 0. |
| Maximum Power      | Specifies the maximum PoE wattage available to provision active PoE ports on the FPC. For example, if you specify 1000 W, the PoE controller is limited to a power budget of 1000 W to distribute to the PoE ports. | Select a value in watts.                                                    |

**Related Documentation**

- [Configuring PoE on EX Series Switches \(CLI Procedure\) on page 4440](#)
- [Example: Configuring PoE Interfaces on an EX Series Switch on page 4433](#)
- [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)
- [Example: Configuring PoE on an EX6200 or EX8200 Switch](#)
- [Monitoring PoE on page 4470](#)
- [Understanding PoE on EX Series Switches on page 4423](#)

## Configuration Statements

- [\[edit poe\] Configuration Statement Hierarchy on EX Series Switches on page 4450](#)
- [af-mode on page 4452](#)
- [disable \(Power over Ethernet\) on page 4453](#)
- [duration on page 4454](#)
- [fpc \(Notification Control\) on page 4455](#)
- [guard-band on page 4456](#)
- [interface \(Power over Ethernet\) on page 4457](#)
- [interval \(Power over Ethernet\) on page 4458](#)
- [management on page 4459](#)
- [maximum-power \(Interface\) on page 4460](#)
- [notification-control on page 4462](#)
- [poe on page 4463](#)
- [priority \(Power over Ethernet\) on page 4464](#)
- [telemetries on page 4465](#)

### [edit poe] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported configuration statements in the **[edit poe]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit poe\] Hierarchy Level on page 4451](#)

### Supported Statements in the [edit poe] Hierarchy Level

The following hierarchy shows the **[edit poe]** configuration statements supported on EX Series switches except for EX6200 and EX8200 switches:

```
poe {
  guard-band watts;
  interface (all | interface-name) {
    disable;
    maximum-power watts;
    priority (high | low);
    telemetries {
      disable;
      duration hours;
      interval minutes;
    }
  }
  lldp-priority;
  management (class | static);
  notification-control {
    fpc slot-number {
      disable;
    }
  }
}
```

The following hierarchy shows the **[edit poe]** configuration statements supported on EX Series switches for EX6200 and EX8200 switches:

```
poe {
  fpc (all | slot-number) {
    guard-band watts;
    lldp-priority;
    management (class | static);
    maximum-power watts;
  }
  interface (all | interface-name) {
    af-mode;
    disable;
    maximum-power watts;
    priority (high | low);
    telemetries {
      disable;
      duration hours;
      interval minutes;
    }
  }
}
```

```
    }  
    notification-control {  
      fpc slot-number {  
        disable;  
      }  
    }  
  }  
}
```

**Related Documentation**

- [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)
- [Example: Configuring PoE on an EX6200 or EX8200 Switch](#)
- [Configuring PoE on EX Series Switches \(CLI Procedure\) on page 4440](#)
- [Understanding PoE on EX Series Switches on page 4423](#)

## **af-mode**

---

|                                 |                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | af-mode;                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <b>poe interface (Power over Ethernet)</b> (all   <i>interface-name</i> )]                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for EX Series switches.                                                                                   |
| <b>Description</b>              | Configure a PoE port on an EX6200 switch to support IEEE 802.3af only. The maximum power the port can deliver in either class or static mode is 15.4 W. |
| <b>Default</b>                  | PoE ports on an EX6200 switch support IEEE 802.3at (PoE+) by default.                                                                                   |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li></ul>                    |

## disable (Power over Ethernet)


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">poe interface</a> (all   <i>interface-name</i> )],<br>[edit <a href="#">poe interface</a> (all   <i>interface-name</i> ) <a href="#">telemetries</a> ],<br>[edit <a href="#">poe notification-control fpc slot-number</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Disable a PoE interface, disable the collection of power consumption data for a PoE interface, or disable the generation of the PoE SNMP traps. The action of the <b>disable</b> statement depends on which statement it is used with:</p> <ul style="list-style-type: none"> <li>When used with <b>interface</b>—Disable the PoE capability of this interface. The interface operates as a standard network access interface, and power is no longer allocated to it from the PoE power budget. Although the PoE capability is disabled, the PoE configuration for the interface is retained. To re-enable the PoE capability of this interface, delete the <b>disable</b> statement from the <b>interface</b> entry in the configuration.</li> <li>When used with <b>telemetries</b>—Disable the collection of PoE power consumption records for this interface. Any previously collected records are deleted. However, the <b>telemetries</b> configuration is retained, including the values for <a href="#">interval</a> and <a href="#">duration</a>. To re-enable record collection, delete the <b>disable</b> statement from the <b>telemetries</b> entry in the configuration.</li> <li>When used with <b>notification-control</b>—Disable the generation of PoE SNMP traps. To re-enable PoE traps, delete the <b>disable</b> statement from the <b>notification-control</b> entry in the configuration.</li> </ul> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435</a></li> <li><a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## duration

---

|                                 |                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>duration <i>hours</i>;</code>                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <a href="#">poe interface</a> (all   <i>interface-name</i> ) <a href="#">telemetries</a> ]                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                             |
| <b>Description</b>              | Modify the duration over which data is collected when you are monitoring the power consumption of a PoE interface.                                                                                                                                               |
| <b>Options</b>                  | <b><i>hours</i></b> —Number of hours over which the data is to be collected.<br><b>Range:</b> 1 through 24<br><b>Default:</b> 1                                                                                                                                  |
| <b>Required Privilege Level</b> | <b>system</b> —To view this statement in the configuration.<br><b>system-control</b> —To add this statement to the configuration.                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435</a></li><li>• <a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li></ul> |

## fpc (Notification Control)

|                                 |                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>fpc slot-number {<br/>    disable;<br/>}</code>                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">po e notification-control</a> ]                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Enable the specified PoE controller to generate PoE traps.                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | PoE traps are disabled by default.                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>slot-number</b>—Indicates the PoE controller by FPC slot number, where <b>slot-number</b> is:</p> <ul style="list-style-type: none"> <li>• 0—On an EX2200, EX3200, standalone EX3300, standalone EX4200, or standalone EX4300 switch</li> <li>• Member ID—On an EX3300, EX4200, EX4300, or EX4600 switch in a Virtual Chassis</li> </ul> |
|                                 | <p> <b>NOTE:</b> EX4600 switches support PoE only when operating in a mixed Virtual Chassis with EX4300 switches.</p>                                                                                                                                        |
|                                 | <ul style="list-style-type: none"> <li>• Line card slot number—On an EX6200 or EX8200 switch</li> </ul> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li> </ul>                                                                                                                                                                                                         |

## guard-band

---

|                                 |                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>guard-band <i>watts</i>;</code>                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <a href="#">poe</a> ],<br>[edit <a href="#">poe</a> (all   fpc <i>slot-number</i> )]                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.                                                                        |
| <b>Description</b>              | Reserve a specified amount of power from the PoE power budget for the switch or the line card in case of a spike in PoE consumption.                                                                                               |
| <b>Options</b>                  | <b>watts</b> —Amount of power to be reserved in case of a spike in PoE consumption.<br><b>Range:</b> 0 through 15 for EX6200 and EX8200 switches and 0 through 19 for ACX2000 routers and all other switches.<br><b>Default:</b> 0 |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li></ul>                                                                                               |



## interface (Power over Ethernet)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interface (all   <i>interface-name</i>) {   af-mode;   disable;   maximum-power <i>watts</i>;   priority (high   low);   telemetries {     disable;     duration <i>hours</i>;     interval <i>minutes</i>;   } } </pre>                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">poe</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Specify a PoE interface to be configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>all</b>—All PoE interfaces on the switch that have not been individually configured for PoE. If a PoE interface has been individually configured, that configuration overrides any settings specified with <b>all</b>.</p> <p><b><i>interface-name</i></b>—Name of the specific interface being configured.</p> <p>If you use the <b>interface</b> statement without any substatements, default values are used for the remaining statements.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435</a></li> <li>• <a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li> </ul>                                                                                                                                                                                                                                               |

## interval (Power over Ethernet)

---

|                                 |                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | interval <i>minutes</i> ;                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <a href="#">poe interface</a> (all   <i>interface-name</i> ) <a href="#">telemetries</a> ]                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Modify the interval at which data is collected when you are monitoring the power consumption of a PoE interface.                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>minutes</i> —Frequency of data collection.<br><b>Range:</b> 1 through 30<br><b>Default:</b> 5                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435</a></li><li>• <a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li><li>• <a href="#">Configuring PoE (J-Web Procedure) on page 4447</a></li></ul> |

## management

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | management (class   static   high-power);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <a href="#">poe</a> ],<br>[edit <a href="#">poe</a> (all   fpc slot-number)]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Designate how the PoE controller allocates power to the PoE interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>                  | class                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <ul style="list-style-type: none"> <li>• <b>class</b>—The amount of power allocated to the interface is determined by the class of the connected powered device. If no powered device is connected, standard 15.4 W power is allocated to the interface. See <a href="#">“Understanding PoE on EX Series Switches” on page 4423</a> for more information about classes of powered devices.</li> <li>• <b>static</b>—The amount of power allocated to the interface is determined by the value of the <a href="#">maximum-power</a> statement, not the class of the connected powered device. This amount is allocated even when a powered device is not connected to the interface, ensuring that power is available when needed.</li> <li>• <b>high-power</b>—(ACX2000 routers only) ACX2000 PoE interfaces support power delivery of up to 65 W per port using all four pairs of Ethernet RJ45 cables. Traditional PoE ports use only two pairs of Ethernet cable for power delivery. According to the IEEE 802.3af standard, each port can deliver a maximum power of up to 32 W. With <b>high-power</b> mode of power delivery over all four pairs, the power sourcing equipment (PSE) has an option to deliver up to 65 W per port, provided the powered devices request this high power over all four pairs of the Ethernet cable. By default, <b>high-power</b> mode is not enabled and has to be explicitly enabled. When the PoE controller is configured for <b>high-power</b> mode, the PoE controller does not deliver power to normal powered devices that request power over two pairs.</li> </ul> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li> <li>• <a href="#">Understanding PoE on EX Series Switches on page 4423</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## maximum-power (Interface)

---

|                            |                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>maximum-power watts;</code>                                                                                                          |
| <b>Hierarchy Level</b>     | [edit <code>poe interface</code> (all   <i>interface-name</i> )]                                                                           |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for ACX2000 routers. |
| <b>Description</b>         | Set the maximum amount of power that the switch can supply to the PoE port.                                                                |



**NOTE:** Although you can set this value when PoE power management is in class mode, it does not establish the maximum power for the port. Instead, the IEEE 802.3af (PoE) or IEEE 802.3at (PoE+) class of the connected device determines the maximum power for the port.

**Options** `watts`—The maximum power in watts that can be supplied to the ports..

For EX2200, EX3300, EX4200, EX4300, EX4600, EX6200, and EX8200 switches:

**Range:** 0.0 through 30.0

**Default:** 15.4 W for ports that support IEEE 802.3af and 30 W for ports that support IEEE 802.3at

For EX3200 switches:

**Range:** 0.0 through 18.6

**Default:** 15.4 W



**NOTE:** EX4600 switches support PoE only when operating in a mixed Virtual Chassis with EX4300 switches.

For ACX2000 routers:

**Range:** 1 through 65 W

**Default:** 32 W



**NOTE:** The `maximum-power` setting permitted by the CLI might be greater than the maximum power a given PoE port can deliver. For example, the CLI permits you to set any PoE port on an EX8200 line card to 30 W; however, only ports 0 through 11 support 30 W. Similarly, the CLI permits you to set any PoE port on an EX4200 switch to 30 W, but some models of EX4200 switch support only 18.6 W per port. If you configure a `maximum-power` value that is greater than the maximum power supported by a port, the power allocated to the port will be the maximum supported.

If you use the all option to set maximum-power to a value greater than 15.4 W on all interfaces on an EX8200 line card, the maximum power allocated to all ports is 15.4 W.



**NOTE:** Support for a maximum of 18.6 W per port instead of 15.4 W per port on EX3200 switches and P and T models of EX4200 switch requires Junos OS Release 11.1 or later. In addition to requiring an upgrade of Junos OS to Release 11.1 or later, switches that are running an earlier release of Junos OS release require the PoE controller software be upgraded as described in [“Upgrading the PoE Controller Software” on page 4467](#). If the controller software is not upgraded and you set maximum-power to a value greater than 15.4 W, the configuration is accepted when you commit it, but the actual power allocated to the port will be 15.4 W.



**NOTE:** On ACX2000 routers, the power sourcing equipment (PSE) delivers up to 65 W per port, provided the management mode is set to high-power mode, by using the high-power option at the [edit poe management] hierarchy level. By default, the management mode is set to static. In the static mode, the PSE can deliver power up to 32 W.

|                                 |                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration. |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                                                                                                           |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li> <li>• <a href="#">management on page 4459</a></li> </ul> |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## notification-control

---

**Syntax**    notification-control {  
              **fpc** *slot-number* {  
                  disable;  
              }  
          }

**Hierarchy Level**    [edit **poe**]

**Release Information**    Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description**    Enable or disable the generation of PoE SNMP traps. If PoE SNMP traps are enabled, an SNMP trap is sent whenever a PoE interface is enabled or disabled.

The remaining statements are explained separately.

**Required Privilege Level**    system—To view this statement in the configuration.  
                                  system-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)
- [Configuring PoE on EX Series Switches \(CLI Procedure\) on page 4440](#)

## poe

**Syntax** For switches other than EX6200 and EX8200 switches:

```
poe {
  guard-band watts;
  interface (all | interface-name) {
    disable;
    maximum-power watts;
    priority (high | low);
    telemetries {
      disable;
      duration hours;
      interval minutes;
    }
  }
  lldp-priority;
  management (class | static);
  notification-control {
    fpc slot-number {
      disable;
    }
  }
}
```

For EX6200 and EX8200 switches:

```
poe {
  fpc (all | slot-number) {
    guard-band watts;
    lldp-priority;
    management (class | static);
    maximum-power watts;
  }
  interface (all | interface-name) {
    af-mode;
    disable;
    maximum-power watts;
    priority (high | low);
    telemetries {
      disable;
      duration hours;
      interval minutes;
    }
  }
  notification-control {
    fpc slot-number {
      disable;
    }
  }
}
```

Hierarchy Level [\[edit\]](#)

|                                 |                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure PoE options.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435</a></li><li>• <a href="#">Example: Configuring PoE on an EX6200 or EX8200 Switch</a></li><li>• <a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li></ul> |

---

## priority (Power over Ethernet)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | priority (low   high);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">poe interface</a> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Set the power priority for individual interfaces when there is insufficient power for all PoE interfaces. If the switch needs to shut down powered devices because PoE demand exceeds the PoE budget, low-priority devices are shut down before high-priority devices. Among interfaces that have the same assigned priority, priority is determined by port number, with lower-numbered ports having higher priority.                                                                                                                                                                                                    |
| <b>Default</b>                  | low                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>high</b>—Specifies that this interface is to be treated as high-priority in terms of power allocation. If the switch needs to shut down powered devices because PoE demand exceeds the PoE budget, power is not shut down on this interface until it has been shut down on all the low-priority interfaces.</p> <p><b>low</b>—Specifies that this interface is to be treated as low-priority in terms of power allocation. If the switch needs to shut down powered devices because PoE demand exceeds the PoE budget, power is shut down on this interface before it is shut down on high-priority interfaces.</p> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435</a></li><li>• <a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li></ul>                                                                                                                                                                                                                                                                                                                                                          |



## telemetries

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>telemetries {   disable;   duration <i>hours</i>;   interval <i>minutes</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit <b>poe interface</b> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Enable the logging of power consumption of a PoE interface over time.</p> <p>If you want to log the power consumption of a PoE interface, you must explicitly specify the <b>telemetries</b> statement. When you commit the configuration, logging begins, with data being collected at the specified intervals. Logging stops at the end of the specified duration. If you do not specify the <b>duration</b> and <b>interval</b> statements, data is collected at the default interval of five minute intervals every hour.</p> <p>The remaining statements are explained separately.</p> |
| <b>Default</b>                  | Logging of power consumption is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435</a></li> <li>• <a href="#">Configuring PoE on EX Series Switches (CLI Procedure) on page 4440</a></li> <li>• <a href="#">show poe telemetries on page 4490</a></li> </ul>                                                                                                                                                                                                                                                               |



# Administration

- [Upgrading on page 4467](#)
- [Routine Monitoring on page 4470](#)
- [Operational Commands on page 4478](#)

## Upgrading

---

- [Upgrading the PoE Controller Software on page 4467](#)

### Upgrading the PoE Controller Software

Each Junos OS image for an EX Series switch that supports PoE contains the most recent version of the PoE controller software at the time the Junos OS image was built. When you upgrade Junos OS on your switch, the new image might contain a more recent version of the PoE controller software than is currently running on the PoE controller. You can upgrade your PoE controller software by requesting that the more recent version of the software contained in the Junos OS image be downloaded to the controller.



**NOTE:** Powered devices are not guaranteed to receive power while the new software is being downloaded to the PoE controller, a process that can take up to 10 minutes. In addition, during the software download, some PoE operational commands, such as [show poe interface](#), might not show correct output. We recommend that you upgrade your PoE controller software during a regularly scheduled maintenance window.



**NOTE:** On an EX8200 Virtual Chassis, you cannot execute PoE commands on the XRE200 External Routing Engine. You can execute PoE commands only on the member EX8200 switches. Use the [request session member member-id](#) command to open a CLI session on a member switch.

This topic covers:

- [Determining Whether the PoE Controller Software Needs Upgrading on page 4468](#)
- [Upgrading the PoE Controller Software on page 4468](#)
- [Monitoring the Upgrade Progress on page 4469](#)

---

### Determining Whether the PoE Controller Software Needs Upgrading

To determine whether the version of the PoE controller software supplied with Junos OS is more recent than the version of the software currently running on the PoE controller, enter the following command:

```
user@switch> show poe controller
Controller Maximum Power Guard Management Status Lldp
index power consumption band Class AF_MODE Priority
0** 130.00W 0.00W 0W Class AF_MODE Disabled
**New PoE software upgrade available.
Use 'request system firmware upgrade poe fpc-slot <slot>'
This procedure will take around 10 minutes (recommended to be performed during
maintenance)
```

The **New PoE software upgrade available** text in the output indicates that the PoE controller software is out-of-date and needs to be upgraded.

For Virtual Chassis or switches with PoE line cards, the output of the **show poe controller** command indicates which members of a Virtual Chassis or which PoE line cards have out-of-date PoE controller software:

```
user@switch> show poe controller
Controller Maximum Power Guard Management Status Lldp
index power consumption band Class AF_MODE Priority
2 130.00W 120.34W 0W Class AF_ENHANCE Disabled
4** 410.00W 182.80W 0W Class AF_MODE Disabled
**New PoE software upgrade available.
Use 'request system firmware upgrade poe fpc-slot slot'
This procedure will take around 10 minutes (recommended to be performed during
maintenance)
```

In the preceding example, member 4 of the Virtual Chassis has an out-of-date PoE controller software.



**NOTE:** We recommend that all member switches of a Virtual Chassis or all line cards in a switch run the same version of the PoE controller software.

---

---

### Upgrading the PoE Controller Software

To upgrade the PoE controller software for a standalone switch with built-in PoE interfaces, enter:

```
user@switch> request system firmware upgrade poe fpc-slot 0
Firmware upgrade initiated. Poe Upgrade takes about 10 minutes
Use 'show poe controller' to get the download status
```

To upgrade the PoE controller software on a specific Virtual Chassis member or line card on a switch, enter:

```
user@switch> request system firmware upgrade poe fpc-slot 8
Firmware upgrade initiated. Poe Upgrade takes about 10 minutes
Use 'show poe controller' to get the download status
```

To upgrade the PoE controller software on all members of a Virtual Chassis or all line cards on a switch, enter:

```
user@switch> request system firmware upgrade poe fpc-slot all-members
Firmware upgrade initiated. Poe Upgrade takes about 10 minutes
Use 'show poe controller' to get the download status
```

### Monitoring the Upgrade Progress

Use the **show poe controller** command to monitor the progress of the controller software upgrade:

```
user@switch> show poe controller
```

| Controller index | Maximum power | Power consumption | Guard band | Management | Status           | Lldp Priority |
|------------------|---------------|-------------------|------------|------------|------------------|---------------|
| 0**              | 130.00W       | 0.00W             | 0W         |            | SW_DOWNLOAD(14%) | Disabled      |

\*\*New PoE software upgrade available.

Use 'request system firmware upgrade poe fpc-slot <slot>'

This procedure will take around 10 minutes (recommended to be performed during maintenance)

The **Status** field is updated during the download process to show the following stages of the download:

- DOWNLOAD\_INIT
- SW\_DOWNLOAD (n%)

When the software upgrade is complete, the **New PoE software upgrade available** text is no longer displayed for the particular FPC.



**NOTE:** If you are upgrading the PoE controller software to enable enhanced PoE, the Status field for the controller shows AF\_ENHANCE after the upgrade completes, indicating that the controller now supports enhanced PoE. The default maximum power per port is not automatically increased as a result of the upgrade—it is still 15.4 W per port. You must explicitly set the maximum power for a port to 18.6 W. Enhanced PoE is supported in Junos OS Release 11.1 or later on EX3200 switches and on EX4200-P or EX4200-T model switches.

#### Related Documentation

- [Verifying PoE Configuration and Status \(CLI Procedure\) on page 4474](#)
- [Understanding PoE on EX Series Switches on page 4423](#)

## Routine Monitoring

---

- [Monitoring PoE on page 4470](#)
- [Monitoring PoE Power Consumption \(CLI Procedure\) on page 4471](#)
- [Verifying PoE Configuration and Status \(CLI Procedure\) on page 4474](#)

### Monitoring PoE

#### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to view real-time data of the power consumed by each PoE interface, and to enable and configure telemetry values. When telemetry is enabled, the software measures the power consumed by each interface and stores the data for future reference.



**NOTE:** If you are configuring a Virtual Chassis, the PoE monitoring option is displayed if any member of the Virtual Chassis supports PoE, even if the Virtual Chassis master does not support PoE.

**Action** To monitor PoE using the J-Web interface, select **Monitor > Power over Ethernet**.

To monitor PoE power consumption with CLI commands in the CLI Terminal in the J-Web interface:

1. Select **Troubleshoot > CLI Terminal**.
2. Type a CLI command:

- [show poe controller](#)
- [show poe interface](#)
- [show poe telemetries](#)

For detailed information about using these CLI commands to monitor PoE power consumption, see [Monitoring PoE Power Consumption \(CLI Procedure\)](#) in the EX Series documentation at <http://www.juniper.net/techpubs>.

**Meaning** In the J-Web interface the PoE Monitoring screen is divided into two parts. The top half of the screen displays real-time data of the power consumed by each PoE-capable interface and a list of ports that utilize maximum power.

Select a particular interface to view a graph of the power consumed by the selected interface.

The bottom half of the screen displays telemetry information for interfaces. The Telemetry Status field displays whether telemetry has been enabled on the interface. Click the **Show Graph** button to view a graph of the telemetries. The graph can be based on power or voltage. To modify telemetry values, click **Edit**. Specify Interval in minutes, Duration in hours, and select **Log Telemetries** to enable telemetry on the selected interface.

#### Related Documentation

- [Configuring PoE on EX Series Switches \(CLI Procedure\) on page 4440](#)
- [Configuring PoE \(J-Web Procedure\) on page 4447](#)
- [Example: Configuring PoE Interfaces on an EX Series Switch on page 4433](#)
- [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)
- [Monitoring PoE Power Consumption \(CLI Procedure\) on page 4471](#)
- [Verifying PoE Configuration and Status \(CLI Procedure\) on page 4474](#)

## Monitoring PoE Power Consumption (CLI Procedure)

You can monitor Power over Ethernet (PoE) power consumption, both for the switch as a whole and for individual PoE interfaces.

This topic describes how to monitor:

- [PoE Power Consumption on a Switch on page 4471](#)
- [Current Power Consumption for PoE Interfaces on page 4471](#)
- [Power Consumption for PoE Interfaces over Time on page 4473](#)

### PoE Power Consumption on a Switch

**Purpose** Determine the current PoE power consumption on a switch.

**Action** Enter the following command:

```
user@switch> show poe controller
Controller  Maximum   Power      Guard    Management  Status  Lldp
index      power     consumption band      Class       AT_MODE  Priority
0          405.00W   130.00W    0W       Class       AT_MODE  Disabled
```

**Meaning** At the time the command was executed, the PoE interfaces on the switch were consuming 130 W out of the PoE power budget of 405 W.

### Current Power Consumption for PoE Interfaces

**Purpose** Determine the current power consumption for individual PoE interfaces.

**Action** To monitor the power consumption of all PoE interfaces on the switch, use the following command:

```
user@switch> show poe interface
```

| Interface | Admin status | Oper status | Max power | Priority | Power consumption | Class          |
|-----------|--------------|-------------|-----------|----------|-------------------|----------------|
| ge-0/0/0  | Enabled      | ON          | 15.4W     | Low      | 7.4W              | 0              |
| ge-0/0/1  | Enabled      | ON          | 15.4W     | High     | 12.0W             | 0              |
| ge-0/0/2  | Enabled      | ON          | 15.4W     | Low      | 12.4W             | 0              |
| ge-0/0/3  | Enabled      | ON          | 7.0W      | Low      | 5.3W              | 2              |
| ge-0/0/4  | Enabled      | ON          | 4.0W      | Low      | 4.0W              | 1              |
| ge-0/0/5  | Enabled      | ON          | 7.0W      | Low      | 6.1W              | 2              |
| ge-0/0/6  | Enabled      | ON          | 15.4W     | Low      | 12.3W             | 3              |
| ge-0/0/7  | Disabled     | Disabled    | 0.0W      | Low      | 0.0W              | not-applicable |

To monitor the power consumption of the PoE interfaces on a specific EX6200 or EX8200 line card, use the following command:

```
user@switch> show poe interface fpc-slot 3
```

| Interface | Admin status | Oper status | Max power | Priority | Power consumption | Class |
|-----------|--------------|-------------|-----------|----------|-------------------|-------|
| ge-3/0/0  | Enabled      | ON          | 30.0W     | Low      | 20.3W             | 4     |
| ge-3/0/1  | Enabled      | ON          | 30.0W     | Low      | 17.8W             | 4     |
| ge-3/0/2  | Enabled      | ON          | 30.0W     | High     | 16.3W             | 4     |
| ge-3/0/3  | Enabled      | ON          | 30.0W     | High     | 16.2W             | 4     |
| ge-3/0/4  | Enabled      | ON          | 30.0W     | Low      | 25.9W             | 4     |
| ge-3/0/5  | Enabled      | ON          | 30.0W     | Low      | 10.1W             | 4     |
| ge-3/0/6  | Enabled      | ON          | 30.0W     | Low      | 16.2W             | 4     |
| ge-3/0/7  | Enabled      | ON          | 30.0W     | Low      | 6.4W              | 4     |
| ge-3/0/8  | Enabled      | ON          | 30.0W     | Low      | 5.2W              | 4     |
| ge-3/0/9  | Enabled      | ON          | 30.0W     | Low      | 5.2W              | 4     |
| ge-3/0/10 | Enabled      | ON          | 30.0W     | Low      | 21.5W             | 4     |
| ge-3/0/11 | Enabled      | ON          | 30.0W     | Low      | 21.7W             | 4     |
| ge-3/0/12 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/13 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/14 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/15 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/16 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/17 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/18 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/19 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/20 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/21 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/22 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/23 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/24 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/25 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/26 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/27 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/28 | Enabled      | ON          | 15.4W     | Low      | 7.0W              | 0     |
| ge-3/0/29 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/30 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/31 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/32 | Enabled      | ON          | 15.4W     | Low      | 2.0W              | 1     |
| ge-3/0/33 | Enabled      | ON          | 15.4W     | Low      | 2.0W              | 1     |
| ge-3/0/34 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/35 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/36 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/37 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/38 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/39 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |



To monitor the power consumption of an individual PoE interface (for example, ge-0/0/3), use the following command:

```
user@switch> show poe interface ge-0/0/3
PoE interface status:
PoE interface           : ge-0/0/3
Administrative status    : Enabled
Operational status      : ON
Power limit on the interface : 7.0W
Priority                 : Low
Power consumed           : 5.3W
Class of power device    : 2
PoE Mode                 : 802.3at
```

**Meaning** At the time the command was executed, the individual PoE ports were consuming the amount of power shown. For example, interface ge-0/0/3 was consuming 5.3 W at the time the command was executed.

### Power Consumption for PoE Interfaces over Time

**Purpose** Monitor the power consumption of a PoE interface over a period of time. The records collected remain available for future viewing.

You can specify the intervals at which power consumption data is collected, from once every minute to once every 30 minutes. The default is once every 5 minutes. You can also specify the duration over which the records are collected, from 1 hour (default) to 24 hours.

**Action** To collect historical records of PoE interface power consumption and display those records:

1. Add the **telemetries** statement to the PoE interface configuration:

```
[edit]
user@switch# set poe interface ge-0/0/5 telemetries interval 10
When you commit the configuration, record collection begins.
```

2. Display the collected records:

```
user@switch> show poe telemetries interface ge-0/0/5 count all
SI No    Timestamp                Power    Voltage
1        03-19-2010 13:00:07 UTC   3.9W    50.9V
2        03-19-2010 12:50:07 UTC   3.9W    50.9V
3        03-19-2010 12:40:07 UTC   3.9W    50.9V
4        03-19-2010 12:30:07 UTC   3.9W    50.9V
5        03-19-2010 12:20:07 UTC   3.9W    50.9V
6        03-19-2010 12:10:07 UTC   3.9W    50.9V
```

To start another session of record collection on the interface, you must delete the existing telemetries configuration on the interface and then reconfigure telemetries. Deleting the telemetries configuration also clears the power consumption history data.

To clear the history of PoE power consumption without deleting the telemetries configuration, use the command **clear poe telemetries interface**.

**Meaning** Over the hour in which the PoE power consumption data on ge-0/0/5 was collected, the connected powered device consistently consumed 3.9 W.

- Related Documentation**
- [Configuring PoE on EX Series Switches \(CLI Procedure\) on page 4440](#)
  - [Example: Configuring PoE Interfaces on an EX Series Switch on page 4433](#)
  - [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)
  - [Example: Configuring PoE on an EX6200 or EX8200 Switch](#)
  - [Verifying PoE Configuration and Status \(CLI Procedure\) on page 4474](#)

## Verifying PoE Configuration and Status (CLI Procedure)

You can verify the Power over Ethernet (PoE) configuration and status on an EX Series switch.

This topic describes how to verify:

- [PoE Controller Configuration and Status on page 4474](#)
- [PoE Interface Configuration and Status on page 4475](#)
- [PoE SNMP Trap Generation Status on page 4476](#)
- [PoE Line Card Configuration and Status on page 4477](#)

### PoE Controller Configuration and Status

---

**Purpose** Verify the PoE controller configuration and status, such as the PoE power budget, total PoE power consumption, power management mode, and the supported PoE standard.

**Action** Enter the following command:

```
user@switch> show poe controller
```

Example output for an EX2200 switch:

| Controller index | Maximum power | Power consumption | Guard band | Management | Status  | Lldp Priority |
|------------------|---------------|-------------------|------------|------------|---------|---------------|
| 0                | 405.00W       | 130.00W           | 19W        | Class      | AT_MODE | Disabled      |

Example output for an EX8200 switch:

| Controller index | Maximum power | Power consumption | Guard band | Management | Status      | Lldp Priority |
|------------------|---------------|-------------------|------------|------------|-------------|---------------|
| 3                | 540.00W       | 435.25W           | 0W         | Static     | AT/AF COMBO | Disabled      |
| 4                | 915.00W       | 627.01W           | 15W        | Class      | AT/AF COMBO | Disabled      |

**Meaning**

- For the EX2200 switch—The switch has a PoE power budget of 405 W, of which 130 W were being used by the PoE ports at the time the command was executed. The **Guard band** field shows that 19 W is reserved out of the PoE power budget to protect against spikes in power demand. The power management mode is class. The PoE ports on the switch support PoE+ (IEEE 802.3at).

- For the EX8200 switch—Line card 3 has a PoE power budget of 540 W, of which 435.25 W were being used by the PoE ports on the line card at the time the command was executed. The management mode for line card 3 is static and the line card has a mix of PoE (IEEE 802.3af) and PoE+ (IEEE 802.3at) ports.

Line card 4 has a PoE power budget of 915 W, of which 627.01 W were being used by the PoE ports on the line card at the time the command was executed. The **Guard band** field shows that 15 W is reserved out of the PoE power budget to protect against spikes in power demand. The management mode for line card 4 is class and the line card has a mix of PoE (IEEE 802.3af) and PoE+ (IEEE 802.3at) ports.

### PoE Interface Configuration and Status

**Purpose** Verify that PoE interfaces are enabled and set to the correct maximum power and priority settings. Also verify current operational status and power consumption.

**Action** To view configuration and status for all PoE interfaces, enter:

```
user@switch> show poe interface
```

| Interface | Admin status | Oper status | Max power | Priority | Power consumption | Class          |
|-----------|--------------|-------------|-----------|----------|-------------------|----------------|
| ge-0/0/0  | Enabled      | ON          | 15.4W     | Low      | 7.9W              | 3              |
| ge-0/0/1  | Enabled      | ON          | 30.0W     | High     | 4.8W              | 0              |
| ge-0/0/2  | Enabled      | ON          | 30.0W     | High     | 4.8W              | 0              |
| ge-0/0/3  | Enabled      | ON          | 7.0W      | High     | 3.3W              | 2              |
| ge-0/0/4  | Enabled      | ON          | 7.0W      | Low      | 3.3W              | 2              |
| ge-0/0/5  | Enabled      | ON          | 7.0W      | Low      | 3.2W              | 2              |
| ge-0/0/6  | Enabled      | ON          | 7.0W      | Low      | 3.3W              | 2              |
| ge-0/0/7  | Enabled      | OFF         | 30.0W     | Low      | 0.0W              | not-applicable |

To view the configuration and status for the PoE interfaces on an EX6200 or EX8200 line card:

```
user@switch> show poe interface fpc-slot 3
```

| Interface | Admin status | Oper status | Max power | Priority | Power consumption | Class |
|-----------|--------------|-------------|-----------|----------|-------------------|-------|
| ge-3/0/0  | Enabled      | ON          | 30.0W     | Low      | 20.3W             | 4     |
| ge-3/0/1  | Enabled      | ON          | 30.0W     | Low      | 17.8W             | 4     |
| ge-3/0/2  | Enabled      | ON          | 30.0W     | High     | 16.3W             | 4     |
| ge-3/0/3  | Enabled      | ON          | 30.0W     | High     | 16.2W             | 4     |
| ge-3/0/4  | Enabled      | ON          | 30.0W     | Low      | 25.9W             | 4     |
| ge-3/0/5  | Enabled      | ON          | 30.0W     | Low      | 10.1W             | 4     |
| ge-3/0/6  | Enabled      | ON          | 30.0W     | Low      | 16.2W             | 4     |
| ge-3/0/7  | Enabled      | ON          | 30.0W     | Low      | 6.4W              | 4     |
| ge-3/0/8  | Enabled      | ON          | 30.0W     | Low      | 5.2W              | 4     |
| ge-3/0/9  | Enabled      | ON          | 30.0W     | Low      | 5.2W              | 4     |
| ge-3/0/10 | Enabled      | ON          | 30.0W     | Low      | 21.5W             | 4     |
| ge-3/0/11 | Enabled      | ON          | 30.0W     | Low      | 21.7W             | 4     |
| ge-3/0/12 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/13 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/14 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/15 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/16 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/17 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/18 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/19 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/20 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |

|           |         |    |       |     |      |   |
|-----------|---------|----|-------|-----|------|---|
| ge-3/0/21 | Enabled | ON | 15.4W | Low | 9.4W | 0 |
| ge-3/0/22 | Enabled | ON | 15.4W | Low | 9.4W | 0 |
| ge-3/0/23 | Enabled | ON | 15.4W | Low | 9.4W | 0 |
| ge-3/0/24 | Enabled | ON | 15.4W | Low | 9.4W | 0 |
| ge-3/0/25 | Enabled | ON | 15.4W | Low | 9.4W | 0 |
| ge-3/0/26 | Enabled | ON | 15.4W | Low | 9.4W | 0 |
| ge-3/0/27 | Enabled | ON | 15.4W | Low | 9.4W | 0 |
| ge-3/0/28 | Enabled | ON | 15.4W | Low | 7.0W | 0 |
| ge-3/0/29 | Enabled | ON | 15.4W | Low | 2.2W | 1 |
| ge-3/0/30 | Enabled | ON | 15.4W | Low | 2.2W | 1 |
| ge-3/0/31 | Enabled | ON | 15.4W | Low | 2.2W | 1 |
| ge-3/0/32 | Enabled | ON | 15.4W | Low | 2.0W | 1 |
| ge-3/0/33 | Enabled | ON | 15.4W | Low | 2.0W | 1 |
| ge-3/0/34 | Enabled | ON | 15.4W | Low | 2.2W | 1 |
| ge-3/0/35 | Enabled | ON | 15.4W | Low | 2.2W | 1 |
| ge-3/0/36 | Enabled | ON | 15.4W | Low | 2.2W | 1 |
| ge-3/0/37 | Enabled | ON | 15.4W | Low | 2.2W | 1 |
| ge-3/0/38 | Enabled | ON | 15.4W | Low | 2.2W | 1 |
| ge-3/0/39 | Enabled | ON | 15.4W | Low | 2.2W | 1 |

To view configuration and status for a single PoE interface, enter:

```

user@switch> show poe interface ge-0/0/3
PoE interface status:
PoE interface           : ge-0/0/3
Administrative status   : Enabled
Operational status      : ON
Power limit on the interface : 7.0W
Priority                 : High
Power consumed          : 3.3W
Class of power device    : 2
PoE Mode                : 802.3at

```

**Meaning** The command output shows the status and configuration of interfaces. For example, the interface ge-0/0/3 is administratively enabled. Its operational status is **ON**; that is, the interface is currently delivering power to a connected powered device. The maximum power allocated to the interface is 7.0 W. The interface has a high power priority. At the time the command was executed, the powered device was consuming 3.3 W. The IEEE 802.3af class of the powered device is class 2. If the PoE power management mode is class, the class of the powered device determines the maximum power allocated to the interface, which is 7 W in the case of class 2 devices.

The PoE Mode field indicates that the interface supports IEEE 802.3at.

### PoE SNMP Trap Generation Status

**Purpose** Verify the status of the **notification-control** option, which determines whether or not PoE SNMP traps are enabled.

**Action** Enter the following command:

```

user@switch> show poe notification-control
FPC slot      Notification-control-status
0              OFF

```

**Meaning** PoE SNMP traps are not enabled.

## PoE Line Card Configuration and Status

**Purpose** Verify the PoE configuration and status for line cards on an EX6200 or EX8200 switch, such as the PoE power allocation and priority for each line card.

**Action** Enter the following command:

```
user@switch> show chassis power-budget-statistics
```

Example output for an EX6200 switch:

```

PSU 0      (EX6200-PWR-AC2500)           :   2500 W   Online
PSU 1      (EX6200-PWR-AC2500)           :   2500 W   Online
PSU 2      (EX6200-PWR-AC2500)           :   2500 W   Online
PSU 3      (EX6200-PWR-AC2500)           :   2500 W   Online
Total Power supplied by all Online PSUs : 10000 W
Power Redundancy Configuration          :   N+1
Power Reserved for the Chassis           :    500 W

Fan Tray Statistics
FTC 0      Base power   Power Used
          :    300 W    43.04 W

FPC Statistics
          Base power   Power Used   PoE power   Priority

FPC 1      (EX6200-48P)           :    220 W    49.47 W    1440 W     1
FPC 2      (EX6200-48P)           :    220 W    47.20 W     800 W     2
FPC 3      (EX6200-48P)           :    220 W   1493.57 W    1440 W     0
FPC 4      (EX6200-SRE64-4XS)     :    100 W    51.38 W     0 W       0
FPC 5      (EX6200-SRE64-4XS)     :    100 W    50.28 W     0 W       0
FPC 6      (EX6200-48P)           :    220 W    49.38 W     800 W     6
FPC 8      (EX6200-48P)           :    220 W    61.41 W    1440 W     9
FPC 9      (EX6200-48T)           :    150 W    12.49 W     0 W       9

Total (non-PoE) Power allocated          :    1750 W
Total Power allocated for PoE            :    5920 W
Power Available (Redundant case)         :    5750 W
Total Power Available                    :    2515 W

```

Example output for an EX8200 switch:

```

PSU 0      (EX8200-AC2K)           :    2000 W   Online
PSU 1      (EX8200-AC2K)           :    2000 W   Online
PSU 2      (EX8200-AC2K)           :    2000 W   Online
PSU 3      (EX8200-AC2K)           :    2000 W   online
PSU 4      (EX8200-AC2K)           :    2000 W   Online
Total Power supplied by all Online PSUs : 10000 W
Power Redundancy Configuration          :   N+1
Power Reserved for the Chassis           :    2400 W

FPC Statistics
          Base power   PoE power   Priority

FPC 1      (EX8200-48T)           :    350 W     0 W     15
FPC 5      (EX8200-2XS-40P)       :    387 W    792 W     0
FPC 9      (EX8200-48PL)          :    267 W    915 W    15
FPC 10     (EX8200-2XS-40T)       :    350 W     0 W     1
FPC 12     (EX8200-48T)           :    350 W     0 W    15

Total (non-PoE) Power allocated          :    4104 W
Total Power allocated for PoE            :    1707 W
Power Available (Redundant case)         :    3896 W
Total Power Available                    :    4263 W

```



- Meaning**
- For the EX6200 switch—The total of the PoE power budgets allocated to the line cards in the switch is 5920 W. This figure includes the 37 W of PoE power always included in the base allocation for each line card that supports PoE. For line cards with PoE ports, the **PoE power** field shows the PoE power budget allocated to each line card, along with the line card priority.
  - For the EX8200 switch—The total of the PoE power budgets allocated to the line cards in the switch is 1707 W. This figure includes the 37 W of PoE power always included in the base allocation for each line card that supports PoE. For line cards with PoE ports, the **PoE power** field shows the PoE power budget allocated to each line card, along with the line card priority.
- Related Documentation**
- [Configuring PoE on EX Series Switches \(CLI Procedure\) on page 4440](#)
  - [Example: Configuring PoE Interfaces on an EX Series Switch on page 4433](#)
  - [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)
  - [Monitoring PoE Power Consumption \(CLI Procedure\) on page 4471](#)

---

## Operational Commands

- `request system firmware upgrade poe`
- `show poe controller`
- `show poe interface`
- `show poe notification-control`
- `show poe telemetry`

## request system firmware upgrade poe

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | request system firmware upgrade poe fpc-slot ( <i>number</i>   all-members)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Upgrade the PoE controller software on switches and line cards.</p> <p>The Junos OS image running on the switch contains a copy of the PoE controller software. This command compares the Junos OS version with the version of the software running on the PoE controller. If the Junos OS version is a more recent version, the command downloads the more recent version to the controller.</p> <p>For all Virtual Chassis except EX8200 Virtual Chassis, execute this command on the master. The master itself need not support PoE for this command to work—for example, you can execute this command on the master of a mixed EX4200 and EX4500 Virtual Chassis when the master is an EX4500 switch, which does not support PoE. On an EX8200 Virtual Chassis, you must execute this command on the member switch, not the master XRE200 External Routing Engine.</p> <p>We recommend that all members of a Virtual Chassis run the same version of the PoE controller software.</p> <p>Upgrading the controller software can take up to 10 minutes. Use the <a href="#">show poe controller</a> command to monitor the progress of the software download.</p> <p>You cannot downgrade the PoE controller software.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> When you enter the <code>request system firmware upgrade poe</code> command, a message advises you that the controller software upgrade has started and that it will take about 10 minutes to complete. This message appears even if the FPC you have specified does not have a PoE controller or if the PoE controller software is up-to-date. To determine whether or not the controller software upgrade has actually started, use the <code>show poe controller</code> command.</p> </div> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> While the upgrade is in progress, power to the powered devices is not guaranteed. We recommend that you upgrade the controller software during a regularly scheduled maintenance window.</p> </div> |
| <b>Options</b>                  | fpc-slot ( <i>number</i>   all-members)—Upgrade the PoE controller firmware for the Virtual Chassis member or line card specified by <i>number</i> , or for all Virtual Chassis members and line cards, specified by <b>all-members</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

- Related Documentation**
- [show poe controller on page 4481](#)
  - [Upgrading the PoE Controller Software on page 4467](#)

**List of Sample Output**    [request system firmware upgrade poe \(Specific FPC Slot\) on page 4480](#)

**Output Fields**    When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### [request system firmware upgrade poe \(Specific FPC Slot\)](#)

```
user@switch> request system firmware upgrade poe fpc-slot 8
Firmware upgrade initiated. Poe Upgrade takes about 10 minutes
Use 'show poe controller' to get the download status
```



## show poe controller

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show poe controller</code>                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.2 for ACX2000 routers.                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Display configuration and status of the PoE controllers.                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show poe interface on page 4484</a></li> <li>• <a href="#">request system firmware upgrade poe on page 4479</a></li> <li>• <a href="#">Verifying PoE Configuration and Status (CLI Procedure) on page 4474</a></li> <li>• <a href="#">Monitoring PoE Power Consumption (CLI Procedure) on page 4471</a></li> <li>• <a href="#">Upgrading the PoE Controller Software on page 4467</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show poe controller (EX3200 Switch) on page 4482</a><br><a href="#">show poe controller (EX8200 Switch) on page 4482</a><br><a href="#">show poe controller (Controller Software Upgrade in Progress) on page 4482</a><br><a href="#">show poe controller (ACX2000 Router) on page 4483</a>                                                                                                                                               |
| <b>Output Fields</b>            | <a href="#">Table 490 on page 4481</a> lists the output fields for the <b>show poe controller</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                    |

**Table 490: show poe controller Output Fields**

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Controller index</b>  | PoE controller number: <ul style="list-style-type: none"> <li>• 0 for EX2200, EX3200, standalone EX3300, standalone EX4200 switches, standalone EX4300 switches, and ACX2000 routers.</li> <li>• Member ID for switches in an EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4300 Virtual Chassis, a mixed EX4200 and EX4500 Virtual Chassis, or a mixed EX4300 and EX4600 Virtual Chassis.</li> <li>• Slot number for line cards with a PoE controller in an EX6200 or EX8200 switch.</li> </ul> |
| <b>Maximum power</b>     | The PoE power budget for the switch or line card. The PoE controller allocates power to the PoE ports from this budget.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Power consumption</b> | Total amount of power being used by the PoE ports at the time the command is executed.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Guard Band</b>        | Amount of power that has been placed in reserve for power demand spikes and that cannot be allocated to a PoE interface.                                                                                                                                                                                                                                                                                                                                                                             |

Table 490: show poe controller Output Fields (*continued*)

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Management    | Power management mode: <b>class</b> or <b>static</b> or <b>high-power</b> .<br><br><b>NOTE:</b> The mode <b>high-power</b> is available on only ACX2000 routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Status        | Status of the PoE controller: <ul style="list-style-type: none"> <li>• <b>AF_ENHANCE</b>—Controller supports enhanced PoE. The maximum power per PoE port is 18.6 W in static mode (15.4 W in class mode).</li> <li>• <b>DEVICE FAIL</b>—Software download to the controller has failed or the PoE controller is not initialized because of a hardware failure.</li> <li>• <b>DOWNLOAD_INIT</b>—Software download to the controller is in the initial phase.</li> <li>• <b>AF_MODE</b>—Controller supports standard IEEE 802.3af. The maximum power per PoE port is 15.4 W.</li> <li>• <b>AT/AF COMBO</b>—Controller supports a mix of standard IEEE 802.3af and IEEE 802.3at (PoE+) ports. The maximum power per port is 30 W for IEEE 802.3at (PoE+) ports and 15.4 W for the IEEE 802.3af ports.</li> <li>• <b>AT_MODE</b>—Controller supports IEEE 802.3at (PoE+). The maximum power per PoE port is 30 W.</li> <li>• <b>SW_DOWNLOAD (n%)</b>—Software download to the controller is in progress.</li> </ul> |
| Lldp Priority | Link Layer Discovery Protocol (LLDP) priority operating state. The state can be <b>Enabled</b> or <b>Disabled</b> .<br><br>LLDP priority enables the PoE controller to assign interfaces the power priority provided by the connected powered device by using LLDP power negotiation rather than the power priority configured on the switch interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## Sample Output

### show poe controller (EX3200 Switch)

```

user@switch> show poe controller
Controller Maximum Power Guard Management Status Lldp
index      power  consumption band    ment     Status  Priority
0          130.00W  81.20W    10W     Static   AF_ENHANCE Disabled

```

### show poe controller (EX8200 Switch)

```

user@switch> show poe controller
Controller Maximum Power Guard Management Status Lldp
index      power  consumption band    ment     Status  Priority
0          792.00W  603.50W    0W      Class    AT/AF COMBO Disabled
4          915.00W  781.00W    0W      Class    AT/AF COMBO Disabled
7          915.00W   0.00W     0W      Class    AT/AF COMBO Disabled

```

### show poe controller (Controller Software Upgrade in Progress)

```

user@switch> show poe controller
Controller Maximum Power Guard Management Status Lldp
index      power  consumption band    ment     Status  Priority
0          130.00W   0.00W     0W     Static   AF_ENHANCE Disabled

```

```
8**      130.00W  0.00W      0W   Static      SW_DOWNLOAD(10%) Disabled
```

```
**New PoE software upgrade available.
```

```
Use 'request system firmware upgrade poe fpc-slot <slot>'
```

```
This procedure will take around 10 minutes (recommended to be performed during  
maintenance)
```

#### show poe controller (ACX2000 Router)

```
user@host> show poe controller
```

| Controller<br>index | Maximum<br>power | Power<br>consumption | Guard<br>band | Management | Status | Lldp<br>Priority |
|---------------------|------------------|----------------------|---------------|------------|--------|------------------|
| 0                   | 130.0 W          | 14.2 W               | 0 W           | high-power | UP     |                  |

## show poe interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show poe interface</b><br><b>&lt;fpc-slot number&gt;</b><br><b>&lt;interface-name&gt;</b>                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.2 for ACX2000 routers.                                                                                                                                                                                                                                                |
| <b>Description</b>              | Display the status of PoE interfaces.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p>none—Display status of all PoE interfaces on the switch or router.</p> <p><b>fpc-slot number</b>—(Optional) (EX6200 or EX8200 switches only) Display the status of the PoE interfaces on the specified line card.</p> <p><b>interface-name</b>—(Optional) Display the status of a specific PoE interface on the switch.</p>                                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show poe controller on page 4481</a></li> <li>• <a href="#">Verifying PoE Configuration and Status (CLI Procedure) on page 4474</a></li> <li>• <a href="#">Monitoring PoE Power Consumption (CLI Procedure) on page 4471</a></li> <li>• <a href="#">Troubleshooting PoE Interfaces on page 4493</a></li> </ul>                   |
| <b>List of Sample Output</b>    | <a href="#">show poe interface on page 4485</a><br><a href="#">show poe interface (with LLDP Negotiation) on page 4485</a><br><a href="#">show poe interface (Specific Interface) on page 4486</a><br><a href="#">show poe interface (Specific FPC Slot) on page 4486</a><br><a href="#">show poe interface (Specific Interface on ACX2000 Universal Access Routers) on page 4487</a> |
| <b>Output Fields</b>            | Table 491 on page 4484 lists the output fields for the <b>show poe interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                     |

**Table 491: show poe interface Output Fields**

| Field Name (All Interfaces Output) | Field Name (Single Interface Output) | Field Description                                                                                                                                                                                       |
|------------------------------------|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b>                   | <b>PoE Interface</b>                 | Interface name.                                                                                                                                                                                         |
| <b>Admin status</b>                | <b>Administrative status</b>         | Administrative state of the PoE interface: <b>Enabled</b> or <b>Disabled</b> . If the PoE interface is disabled, it can provide network connectivity, but it cannot provide power to connected devices. |

Table 491: show poe interface Output Fields (*continued*)

| Field Name (All Interfaces Output) | Field Name (Single Interface Output) | Field Description                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oper status                        | Operational status                   | Operational state of the PoE interface: <ul style="list-style-type: none"> <li>• <b>ON</b>—The interface is currently supplying power to a powered device.</li> <li>• <b>OFF</b>—PoE is enabled on the interface, but the interface is not currently supplying power to a powered device.</li> <li>• <b>Disabled</b>—PoE is disabled on the interface.</li> </ul> |
| Max power                          | Power limit on the interface         | Maximum power that can be provided by the interface.<br>An (L) next to the value indicates that the value on the port was negotiated by LLDP.                                                                                                                                                                                                                     |
| Priority                           | Priority                             | Interface power priority: either <b>High</b> or <b>Low</b> .<br>An (L) next to the value indicates that the value on the port was negotiated by LLDP.                                                                                                                                                                                                             |
| Power consumption                  | Power consumed                       | Amount of power being used by the interface at the time the command is executed.                                                                                                                                                                                                                                                                                  |
| Class                              | Class of power device                | IEEE 802.3af (PoE) or IEEE 802.3at (PoE+) class of the powered device. <b>Class 0</b> is the default class and is used when the class of the powered device is unknown. If no powered device is connected, this field contains <b>not applicable</b> .                                                                                                            |
|                                    | PoE Mode                             | IEEE PoE standard supported by the interface—either <b>802.3af</b> or <b>802.3at</b> .                                                                                                                                                                                                                                                                            |

## Sample Output

### show poe interface

```

user@switch> show poe interface
Interface Admin    Oper    Max    Priority Power    Class
           status    status  power                consumption
ge-0/0/0 Enabled    ON      15.4W  Low     7.9W     0
ge-0/0/1 Enabled    ON      15.4W  Low     3.2W     2
ge-0/0/2 Enabled    ON      15.4W  Low     3.2W     2
ge-0/0/3 Enabled    ON      15.4W  Low     3.2W     2
ge-0/0/4 Enabled    ON      15.4W  Low     3.2W     2
ge-0/0/5 Enabled    ON      15.4W  Low     3.2W     2
ge-0/0/6 Enabled    ON      15.4W  Low     3.2W     2
ge-0/0/7 Enabled    ON      15.4W  Low     3.2W     2

```

### show poe interface (with LLDP Negotiation)

```

user@switch> show poe interface
Interface Admin    Oper    Max    Priority Power    Class
           status    status  power                consumption
ge-0/0/0 Enabled    ON      17.5W(L) Low(L)    16.2W     4
ge-0/0/1 Enabled    ON      17.5W(L) Low(L)    16.0W     4
ge-0/0/2 Enabled    ON      17.5W(L) High(L)   16.0W     4
ge-0/0/3 Enabled    ON      17.5W(L) Low(L)    16.0W     4
ge-0/0/4 Enabled    ON      10.1W(L) Low(L)    10.0W     3

```

```

ge-0/0/5    Enabled    ON    3.5W(L)    High(L)    3.0W    2
(L) LLDP-negotiated value on the port.

```

### show poe interface (Specific Interface)

```

user@switch> show poe interface ge-0/0/3
PoE interface status:
PoE interface           : ge-0/0/3
Administrative status    : Enabled
Operational status      : ON
Power limit on the interface : 7.0W
Priority                 : Low
Power consumed           : 5.3W
Class of power device    : 2
PoE Mode                : 802.3af

```

### show poe interface (Specific FPC Slot)

```

user@switch> show poe interface fpc-slot 3

```

| Interface | Admin status | Oper status | Max power | Priority | Power consumption | Class |
|-----------|--------------|-------------|-----------|----------|-------------------|-------|
| ge-3/0/0  | Enabled      | ON          | 30.0W     | Low      | 20.3W             | 4     |
| ge-3/0/1  | Enabled      | ON          | 30.0W     | Low      | 17.8W             | 4     |
| ge-3/0/2  | Enabled      | ON          | 30.0W     | High     | 16.3W             | 4     |
| ge-3/0/3  | Enabled      | ON          | 30.0W     | High     | 16.2W             | 4     |
| ge-3/0/4  | Enabled      | ON          | 30.0W     | Low      | 25.9W             | 4     |
| ge-3/0/5  | Enabled      | ON          | 30.0W     | Low      | 10.1W             | 4     |
| ge-3/0/6  | Enabled      | ON          | 30.0W     | Low      | 16.2W             | 4     |
| ge-3/0/7  | Enabled      | ON          | 30.0W     | Low      | 6.4W              | 4     |
| ge-3/0/8  | Enabled      | ON          | 30.0W     | Low      | 5.2W              | 4     |
| ge-3/0/9  | Enabled      | ON          | 30.0W     | Low      | 5.2W              | 4     |
| ge-3/0/10 | Enabled      | ON          | 30.0W     | Low      | 21.5W             | 4     |
| ge-3/0/11 | Enabled      | ON          | 30.0W     | Low      | 21.7W             | 4     |
| ge-3/0/12 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/13 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/14 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/15 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/16 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/17 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/18 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/19 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/20 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/21 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/22 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/23 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/24 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/25 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/26 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/27 | Enabled      | ON          | 15.4W     | Low      | 9.4W              | 0     |
| ge-3/0/28 | Enabled      | ON          | 15.4W     | Low      | 7.0W              | 0     |
| ge-3/0/29 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/30 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/31 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/32 | Enabled      | ON          | 15.4W     | Low      | 2.0W              | 1     |
| ge-3/0/33 | Enabled      | ON          | 15.4W     | Low      | 2.0W              | 1     |
| ge-3/0/34 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/35 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/36 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/37 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/38 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |
| ge-3/0/39 | Enabled      | ON          | 15.4W     | Low      | 2.2W              | 1     |

**show poe interface (Specific Interface on ACX2000 Universal Access Routers)**

```
user@host> show poe interface ge-0/1/7
PoE interface status:
PoE interface           : ge-0/1/7
Administrative status   : Enabled
Operational status     : Powered-up
Power limit on the interface : 9.0 W
Priority                 : Low
Power consumed          : 14.2 W
Class of power device   : 4
```

## show poe notification-control

---

|                                 |                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show poe notification-control</b>                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                         |
| <b>Description</b>              | Display the state of the PoE <b>notification-control</b> option, which enables or disables PoE SNMP traps.                                                                                                                                                 |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show poe controller on page 4481</a></li><li>• <a href="#">show poe interface on page 4484</a></li><li>• <a href="#">Verifying PoE Configuration and Status (CLI Procedure) on page 4474</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">show poe notification-control on page 4489</a>                                                                                                                                                                                                 |
| <b>Output Fields</b>            | <a href="#">Table 492 on page 4488</a> lists the output fields for the <b>show poe notification-control</b> command. Output fields are listed in the approximate order in which they appear.                                                               |

**Table 492: show poe notification-control Output Fields**

| Field Name                         | Field Description                                                                                                                                               |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>FPC slot</b>                    | FPC slot number: <ul style="list-style-type: none"><li>• 0 for a standalone switch</li><li>• Member ID for a Virtual Chassis</li></ul>                          |
| <b>Notification-control-status</b> | Status of notification control: <ul style="list-style-type: none"><li>• <b>ON</b>—PoE traps are enabled.</li><li>• <b>OFF</b>—PoE traps are disabled.</li></ul> |



## Sample Output

### show poe notification-control

```
user@switch> show poe notification-control
FPC slot      Notification-control-status
  0              OFF
```

## show poe telemetries

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show poe telemetries<br><count (all   <i>number</i> )><br><interface (all   <i>interface-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display a history of power consumption on the specified interface or on all interfaces.<br><br>Telemetries must be enabled on the interface before you can display a history of power consumption.                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>none</b>—Displays all records for all interfaces that have power consumption history data.</p> <p><b>count (all   <i>number</i>)</b>—(Optional) Specify the number of power consumption records to display. The most recent records are displayed. If you do not specify the count, all available records are displayed.</p> <p><b>interface (all   <i>interface-name</i>)</b>—(Optional) Display power consumption records for the specified PoE interface or for all PoE interfaces. If you do not specify interfaces, all interfaces are displayed.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show poe interface on page 4484</a></li> <li>• <a href="#">show poe controller on page 4481</a></li> <li>• <a href="#">Monitoring PoE Power Consumption (CLI Procedure) on page 4471</a></li> <li>• <a href="#">Verifying PoE Configuration and Status (CLI Procedure) on page 4474</a></li> <li>• <a href="#">Troubleshooting PoE Interfaces on page 4493</a></li> </ul>                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show poe telemetries interface all count on page 4491</a><br><a href="#">show poe telemetries interface (Specific Interface) count all on page 4491</a>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Output Fields</b>            | Table 493 on page 4490 lists the output fields for the <b>show poe telemetries interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                    |

**Table 493: show poe telemetries interface Output Fields**

| Field Name       | Field Description                                                                     |
|------------------|---------------------------------------------------------------------------------------|
| <b>Interface</b> | Name of the interface.                                                                |
| <b>SI No</b>     | Number of the record for the specified interface. Record number 1 is the most recent. |
| <b>Timestamp</b> | Date and time when the power-consumption data was gathered.                           |

Table 493: show poe telemetries interface Output Fields (*continued*)

| Field Name     | Field Description                                                                      |
|----------------|----------------------------------------------------------------------------------------|
| <b>Power</b>   | Amount of power provided by the specified interface at the time the data was gathered. |
| <b>Voltage</b> | Maximum voltage provided by the specified interface at the time the data was gathered. |

## Sample Output

### show poe telemetries interface all count

```

user@switch> show poe telemetries interface all count 2
Interface  S1 No    Timestamp                Power    Voltage
ge-0/0/1   1    03-09-2012 11:52:03 UTC  4.2W    54.9V
           2    03-09-2012 11:47:03 UTC  4.2W    54.8V
ge-0/0/2   1    03-09-2012 11:52:03 UTC  4.2W    54.9V
           2    03-09-2012 11:47:03 UTC  4.1W    54.8V
ge-0/0/3   1    03-09-2012 11:52:03 UTC  4.2W    54.9V
           2    03-09-2012 11:47:03 UTC  4.3W    54.8V
ge-0/0/4   1    03-09-2012 11:52:03 UTC  0.0W    54.9V
           2    03-09-2012 11:47:03 UTC  0.0W    54.8V
ge-0/0/5   1    03-09-2012 11:52:03 UTC  4.2W    54.9V
           2    03-09-2012 11:47:03 UTC  4.2W    54.8V
ge-0/0/6   1    03-09-2012 11:52:03 UTC  4.2W    54.9V
           2    03-09-2012 11:47:03 UTC  4.2W    54.8V
ge-0/0/7   1    03-09-2012 11:52:03 UTC  4.2W    54.9V

```

### show poe telemetries interface (Specific Interface) count all

```

user@switch> show poe telemetries interface ge-0/0/0 count all
S1 No    Timestamp                Power    Voltage
1      01-27-2008 18:19:58 UTC  15.4W    51.6V
2      01-27-2008 18:18:58 UTC  15.4W    51.6V
3      01-27-2008 18:17:58 UTC  15.4W    51.6V
4      01-27-2008 18:16:58 UTC  15.4W    51.6V
5      01-27-2008 18:15:58 UTC  15.4W    51.6V
6      01-27-2008 18:14:58 UTC  15.4W    51.6V
7      01-27-2008 18:13:58 UTC  15.4W    51.6V
8      01-27-2008 18:12:57 UTC  15.4W    51.6V
9      01-27-2008 18:11:57 UTC  15.4W    51.6V
10     01-27-2008 18:10:57 UTC  15.4W    51.6V
11     01-27-2008 18:09:57 UTC  15.4W    51.6V
12     01-27-2008 18:08:57 UTC  15.4W    51.6V
13     01-27-2008 18:07:57 UTC  15.4W    51.6V
14     01-27-2008 18:06:57 UTC  15.4W    51.6V
15     01-27-2008 18:05:57 UTC  15.4W    51.6V
16     01-27-2008 18:04:56 UTC  15.4W    51.6V
17     01-27-2008 18:03:56 UTC  15.4W    51.6V
18     01-27-2008 18:02:56 UTC  15.4W    51.6V
19     01-27-2008 18:01:56 UTC  15.4W    51.6V
20     01-27-2008 18:00:56 UTC  15.4W    51.6V
21     01-27-2008 17:59:56 UTC  15.4W    51.6V

```



# Troubleshooting Procedures

- [Troubleshooting PoE Interfaces on page 4493](#)

## Troubleshooting PoE Interfaces

**Problem**    **Description:** A Power over Ethernet (PoE) interface is not supplying power to the powered device.

**Solution**    Check for the items shown in [Table 494 on page 4493](#).

**Table 494: Troubleshooting a PoE Interface**

| Items to Check                                                                                                        | Explanation                                                                                                       |
|-----------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| Is the switch a full PoE model or a partial PoE model?                                                                | If you are using a partial PoE model, only interfaces ge-0/0/0 through ge-0/0/7 can function as PoE ports.        |
| Has PoE capability been disabled for that interface?                                                                  | Use the <a href="#">show poe interface</a> command to check PoE interface status.                                 |
| Is the cable properly seated in the port socket?                                                                      | Check the hardware.                                                                                               |
| Has the PoE power budget been exceeded for the switch?                                                                | Use the <a href="#">show poe controller</a> command to check the PoE power budget and consumption for the switch. |
| Does the powered device require more power than is available on the interface?                                        | Use the <a href="#">show poe interface</a> command to check the maximum power provided by the interface.          |
| If the <a href="#">telemetries</a> option has been enabled for the interface, check the history of power consumption. | Use the <a href="#">show poe telemetries</a> command to display the history of power consumption.                 |

- Related Documentation**
- [Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 4435](#)
  - [Verifying PoE Configuration and Status \(CLI Procedure\) on page 4474](#)
  - [Monitoring PoE Power Consumption \(CLI Procedure\) on page 4471](#)

- [Configuring PoE on EX Series Switches \(CLI Procedure\) on page 4440](#)

## PART 23

# Port Security

- [Overview on page 4497](#)
- [Configuration on page 4529](#)
- [Administration on page 4657](#)





## CHAPTER 74

# Overview

- [Port Security Overview on page 4497](#)

### Port Security Overview

---

- [Understanding Port Security on page 4497](#)
- [Understanding DAI for Port Security on page 4500](#)
- [Understanding DHCP Option 82 for Port Security on Switching Devices on page 4503](#)
- [Understanding DHCP Snooping for Port Security on page 4507](#)
- [Understanding IP Source Guard for Port Security on EX Series Switches on page 4514](#)
- [Understanding IPv6 Neighbor Discovery Inspection on page 4516](#)
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches on page 4518](#)
- [Understanding Media Access Control Security \(MACsec\) on page 4520](#)
- [Understanding Persistent MAC Learning \(Sticky MAC\) on page 4526](#)
- [Understanding Trusted DHCP Servers for Port Security on page 4527](#)

### Understanding Port Security

Ethernet LANs are vulnerable to attacks such as address spoofing (forging) and Layer 2 denial of service (DoS) on network devices. Port security features help protect the access ports on your device against the loss of information and productivity that such attacks can cause.

The Juniper Networks Junos operating system (Junos OS) provides features to help secure ports on a device. Ports can be categorized as either trusted or untrusted. You apply policies appropriate to each category to protect ports against various types of attacks.

Basic port security features are enabled in the device's default configuration. You can configure additional features with minimal configuration steps.

Depending on the particular feature, you can configure the feature either on VLANs or bridge domain interfaces.

Port security features supported on switching devices are:

- DHCP snooping—Filters and blocks ingress Dynamic Host Configuration Protocol (DHCP) server messages on untrusted ports; builds and maintains a database of DHCP lease information, which is called the DHCP snooping database.



**NOTE:** DHCP snooping is not enabled in the default configuration of the switching device. DHCP snooping is enabled on a VLAN or bridge domain. The details of enabling DHCP snooping depend on the particular device.

- DHCPv6 snooping—DHCP snooping for IPv6.
- DHCP option 82—Also known as the DHCP Relay Agent Information option. This DHCPv4 feature helps protect the switching device against attacks such as spoofing of IP addresses and MAC addresses and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.
- DHCPv6 option 37—Option 37 is the remote ID option for DHCPv6 and is used to insert information about the network location of the remote host into DHCPv6 packets. You enable option 37 on a VLAN.



**NOTE:** DHCPv6 snooping with option 37 is not supported on the MX Series.

- DHCPv6 option 18—Option 18 is the circuit ID option for DHCPv6 and is used to insert information about the client port into DHCPv6 packets. This option includes other details that can be optionally configured, such as the prefix and the interface description.
- DHCPv6 option 16—Option 16 is the vendor ID option for DHCPv6 and is used to insert information about the vendor of the client hardware into DHCPv6 packets.
- Dynamic ARP inspection (DAI)—Prevents Address Resolution Protocol (ARP) spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made on the basis of the results of those comparisons. You enable DAI on a VLAN.
- IPv6 neighbor discovery inspection—Prevents IPv6 address spoofing attacks. Neighbor discovery requests and replies are compared against entries in the DHCPv6 snooping database, and filtering decisions are made on the basis of the results of those comparisons. You enable neighbor discovery inspection on a VLAN.
- IP source guard—Mitigates the effects of IP address spoofing attacks on the Ethernet LAN. With IP source guard enabled, the source IP address in the packet sent from an untrusted access interface is validated against the DHCP snooping database. If the packet cannot be validated, it is discarded. You enable IP source guard on a VLAN or bridge domain.



**NOTE:** IP source guard is not supported on the QFX Series.

- IPv6 source guard—IP source guard for IPv6.



**NOTE:** IPv6 source guard is not supported on the QFX Series.

- MAC limiting—Protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You can enable MAC limiting on an interface.
- MAC move limiting—(Not supported on EX9200) Tracks MAC movement and detects MAC spoofing on access ports. You enable this feature on a VLAN or bridge domain.
- Persistent MAC learning—Also known as sticky MAC. Persistent MAC learning enables interfaces to retain dynamically learned MAC addresses across switch reboots. You enable this feature on an interface.
- Trusted DHCP server—Configuring the DHCP server on a trusted port protects against rogue DHCP servers sending leases. You enable this feature on an interface (port). By default, access ports are untrusted, and trunk ports are trusted. (Access ports are the switch ports that connect to Ethernet endpoints such as user PCs and laptops, servers, and printers. Trunk ports are the switch ports that connect an Ethernet switch to other switches or to routers.)

#### Related Documentation

- [Security Features for EX Series Switches Overview on page 4693](#)
- [Understanding DHCP Snooping for Port Security](#)
- [Understanding DHCP Snooping for Port Security on page 4507](#)
- [Understanding IPv6 Neighbor Discovery Inspection on page 4516](#)
- [Understanding DAI for Port Security on page 4500](#)
- [Understanding IP Source Guard for Port Security on EX Series Switches on page 4514](#)
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches on page 4518](#)
- [Understanding DHCP Option 82 for Port Security on Switching Devices on page 4503](#)

## Understanding DAI for Port Security

Dynamic ARP inspection (DAI) protects switching devices against ARP spoofing.

DAI inspects Address Resolution Protocol (ARP) packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP spoofing (also known as ARP poisoning or ARP cache poisoning). ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons. When an attacker tries to use a forged ARP packet to spoof an address, the switch compares the address with entries in the database. If the media access control (MAC) address or IP address in the ARP packet does not match a valid entry in the DHCP snooping database, the packet is dropped.

ARP packets are sent to the Routing Engine and are rate-limited to protect the switching device from CPU overload.

- [Address Resolution Protocol on page 4500](#)
- [ARP Spoofing on page 4500](#)
- [Dynamic ARP Inspection on page 4501](#)
- [Prioritizing Inspected Packets on page 4502](#)

---

### Address Resolution Protocol

Sending IP packets on a multi-access network requires mapping an IP address to an Ethernet MAC address.

Ethernet LANs use ARP to map MAC addresses to IP addresses.

The switching device maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

---

### ARP Spoofing

ARP spoofing is one way to initiate man-in-the-middle attacks. The attacker sends an ARP packet that spoofs the MAC address of another device on the LAN. Instead of the switching device sending traffic to the proper network device, it sends the traffic to the device with the spoofed address that is impersonating the proper device. If the impersonating device is the attacker's machine, the attacker receives all the traffic from the switch that must have gone to another device. The result is that traffic from the switching device is misdirected and cannot reach its proper destination.

One type of ARP spoofing is gratuitous ARP, which is when a network device sends an ARP request to resolve its own IP address. In normal LAN operation, gratuitous ARP messages indicate that two devices have the same MAC address. They are also broadcast when a network interface card (NIC) in a device is changed and the device is rebooted, so that other devices on the LAN update their ARP caches. In malicious situations, an attacker can poison the ARP cache of a network device by sending an ARP response to

the device that directs all packets destined for a certain IP address to go to a different MAC address instead.

To prevent MAC spoofing through gratuitous ARP and through other types of spoofing, the switches examine ARP responses through DAI.

### Dynamic ARP Inspection

DAI examines ARP requests and responses on the LAN and validates ARP packets. The switch intercepts ARP packets from an access port and validates them against the DHCP snooping database. If no IP-MAC entry in the database corresponds to the information in the ARP packet, DAI drops the ARP packet and the local ARP cache is not updated with the information in that packet. DAI also drops ARP packets when the IP address in the packet is invalid. ARP probe packets are not subjected to dynamic ARP inspection. The switch always forwards such packets.

Junos OS for EX Series switches and the QFX Series uses DAI for ARP packets received on access ports because these ports are untrusted by default. Trunk ports are trusted by default, and therefore ARP packets bypass DAI on them.

You configure DAI for each VLAN, not for each interface (port). By default, DAI is disabled for all VLANs.

If you set an interface to be a DHCP trusted port, it is also trusted for ARP packets.



#### NOTE:

- If your switching device is an EX Series switch and uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style, see [“Enabling a Trusted DHCP Server \(CLI Procedure\)” on page 4573](#) for information about configuring an access interface to be a DHCP trusted port.
- If your switching device is an EX Series switch and is *not* using Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style, see *Enabling a Trusted DHCP Server (CLI Procedure)* for information about configuring an access interface to be a DHCP trusted port.

For packets directed to the switching device to which a network device is connected, ARP queries are broadcast on the VLAN. The ARP responses to those queries are subjected to the DAI check.

For DAI, all ARP packets are trapped to the Packet Forwarding Engine. To prevent CPU overloading, ARP packets destined for the Routing Engine are rate-limited.

If the DHCP server goes down and the lease time for an IP-MAC entry for a previously valid ARP packet runs out, that packet is blocked.

## Prioritizing Inspected Packets

---



**NOTE:** Prioritizing inspected packets is not supported on the QFX Series and the EX4600 switch.

---

You can use class-of-service (CoS) forwarding classes and queues to prioritize DAI packets for a specified VLAN. This type of configuration places inspected packets for that VLAN in the egress queue, that you specify, ensuring that the security procedure does not interfere with the transmission of high-priority traffic.

### Related Documentation

- [Understanding Port Security on page 4497](#)
- *Understanding DHCP Snooping for Port Security*
- *Example: Configuring Basic Port Security Features*
- *Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch*
- *Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks*
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529](#)
- *Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic*
- *Enabling Dynamic ARP Inspection (CLI Procedure)*
- [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 4572](#)
- [Enabling Dynamic ARP Inspection \(J-Web Procedure\) on page 4572](#)

## Understanding DHCP Option 82 for Port Security on Switching Devices

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect Juniper Networks EX Series Ethernet Switches and MX Series 3D Universal Edge Routers against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Hosts on untrusted access interfaces on an Ethernet LAN switching device send requests for IP addresses to access the Internet. The switching device forwards or relays these requests to DHCP servers, and the servers send offers for IP address leases in response. Attackers can use these messages to penetrate the network by address spoofing.

Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client. The Junos OS implementation of DHCP option 82 supports RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.

This topic covers:

- [DHCP Option 82 Processing on page 4503](#)
- [Suboption Components of Option 82 on page 4504](#)
- [Switching Device Configurations That Support Option 82 on page 4505](#)
- [DHCPv6 Options on page 4506](#)

### DHCP Option 82 Processing

If DHCP option 82 is enabled on a VLAN or bridge domain, then when a network device—a DHCP client—that is connected to the VLAN or bridge domain on an untrusted interface sends a DHCP request, the switching device inserts information about the client's network location into the packet header of that request. The switching device then sends the request to the DHCP server. The DHCP server reads the option 82 information in the packet header and uses it to implement the IP address or another parameter for the client. See “[Suboption Components of Option 82](#)” on [page 4504](#) for more information about option 82.



#### NOTE:

- If your switching device is an EX Series switch and uses Junos OS with Enhanced Layer 2 Software (ELS) configuration style, you can enable DHCP option 82 only for a specific VLAN. See “[Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\)](#)” on [page 4577](#).
- If your switching device is an EX Series switch and does *not* use Junos OS with Enhanced Layer 2 Software (ELS) configuration style, you can enable DHCP option 82 either for a specific VLAN or for all VLANs. See *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*.

If option 82 is enabled on a VLAN or bridge domain, the following sequence of events occurs when a DHCP client sends a DHCP request:

1. The switching device receives the request and inserts the option 82 information in the packet header.
2. The switching device forwards (or relays) the request to the DHCP server.
3. The server uses the DHCP option 82 information to formulate its reply and sends a response to the switching device. It does not alter the option 82 information.
4. The switching device strips the option 82 information from the response packet.
5. The switching device forwards the response packet to the client.



**NOTE:** To use the DHCP option 82 feature, you must ensure that the DHCP server is configured to accept option 82. If the DHCP server is not configured to accept option 82, then when it receives requests containing option 82 information, it does not use the information for setting parameters and it does not echo the information in its response message.

---

### Suboption Components of Option 82

---

Option 82 as implemented on a switching device comprises the suboptions circuit ID, remote ID, and vendor ID. These suboptions are fields in the packet header:

- circuit ID—Identifies the circuit (interface or VLAN) on the switching device on which the request was received. The circuit ID contains the interface name and VLAN name, with the two elements separated by a colon—for example, `ge-0/0/10:vlan1`, where `ge-0/0/10` is the interface name and `vlan1` is the VLAN name. If the request packet is received on a Layer 3 interface, the circuit ID is just the interface name—for example, `ge-0/0/10`.

Use the prefix option to add an optional prefix to the circuit ID. If you enable the prefix option, the hostname for the switching device is used as the prefix; for example, `device1:ge-0/0/10:vlan1`, where `device1` is the hostname.

You can also specify that the interface description be used rather than the interface name or that the VLAN ID be used rather than the VLAN name.

- remote ID—Identifies the remote host. See [remote-id](#) for details.
- vendor ID—Identifies the vendor of the host. If you specify the **vendor-id** option but do not enter a value, the default value Juniper is used. To specify a value, you type a character string.



## Switching Device Configurations That Support Option 82

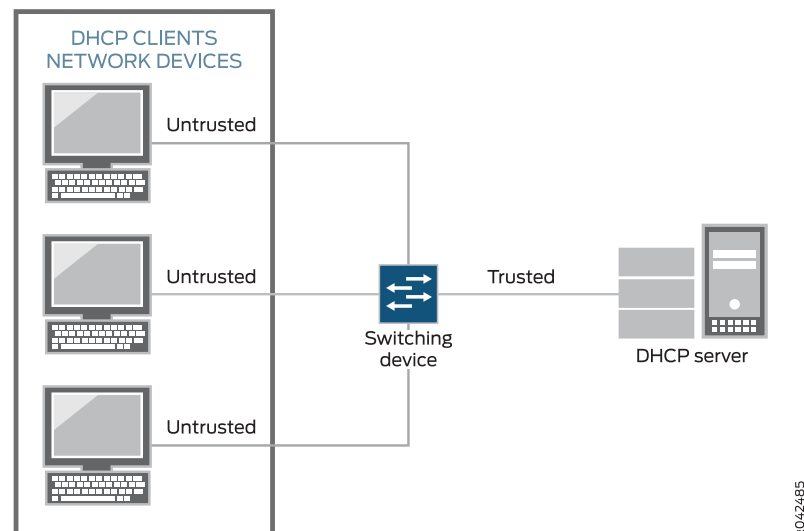
Switching device configurations that support option 82 are:

- [Switching Device, DHCP Clients, and the DHCP Server Are on the Same VLAN or Bridge Domain on page 4505](#)
- [Switching Device Acts as a Relay Agent on page 4505](#)

### ***Switching Device, DHCP Clients, and the DHCP Server Are on the Same VLAN or Bridge Domain***

If the switching device, the DHCP clients, and the DHCP server are all on the same VLAN or bridge domain, the switching device forwards the requests from the clients on untrusted access interfaces to the server on a trusted interface. See [Figure 58 on page 4505](#).

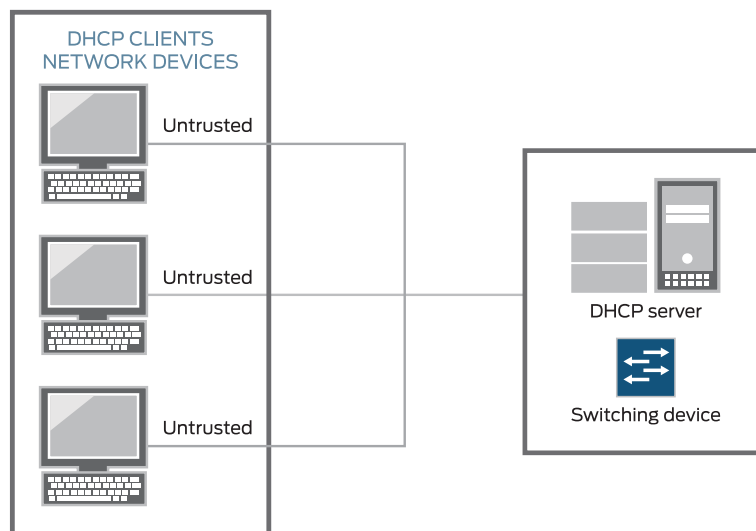
**Figure 58: DHCP Clients, Switching Device, and the DHCP Server Are All on the Same VLAN or Bridge Domain**



### ***Switching Device Acts as a Relay Agent***

The switching device functions as a relay agent (extended relay server) when the DHCP clients or the DHCP server is connected to the switching device through a Layer 3 interface. On the switching device, these interfaces are configured as routed VLAN interfaces (RVIs). [Figure 59 on page 4506](#) illustrates a scenario for the switching device acting as an extended relay server; in this instance, the switching device relays requests to the server.

Figure 59: Switching Device Acting as an Extended Relay Server



8042486

## DHCPv6 Options



**NOTE:** MX Series routers do not support DHCPv6.

DHCPv6 provides several options that can be used to insert information into the DHCPv6 request packets that are relayed to a server from a client. These options are equivalent to the sub-options of DHCP option 82.

- Option 37—Identifies the remote host. Option 37 is equivalent to the **remote-id** sub-option of DHCP option 82.
- Option 18—Identifies the interface on which the DHCP request packet was received from the client. Option 18 is equivalent to the **circuit-id** sub-option of DHCP option 82.
- Option 16—Identifies the vendor of the hardware on which the client is hosted. Option 16 is equivalent to the **vendor-id** sub-option of DHCP option 82.

DHCPv6 options are not enabled automatically when DHCPv6 snooping is enabled on a VLAN. They must be configured using the **dhcpv6-options** statement.

### Related Documentation

- [Setting Up DHCP Option 82 on an MX Series Router \(CLI Procedure\)](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4577](#)

## Understanding DHCP Snooping for Port Security



**NOTE:** This topic includes information about enabling DHCP snooping when using Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Understanding DHCP Snooping for Port Security*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

DHCP snooping enables the switch to monitor and control DHCP messages received from untrusted devices connected to the switch. When DHCP snooping is enabled on a VLAN, the system snoops the DHCP messages to view DHCP lease information and build and maintain a database of valid IP address to MAC address (IP-MAC) bindings called the DHCP snooping database. Only clients with valid bindings are allowed access to the network.

- [DHCP Snooping Basics on page 4507](#)
- [Enabling DHCP Snooping on page 4508](#)
- [DHCP Snooping Process on page 4508](#)
- [DHCPv6 Snooping on page 4509](#)
- [Rapid Commit for DHCPv6 on page 4509](#)
- [DHCP Server Access on page 4510](#)
- [Static IP Address Additions to the DHCP Snooping Database on page 4513](#)
- [Snooping DHCP Packets That Have Invalid IP Addresses on page 4513](#)

### DHCP Snooping Basics

Dynamic Host Configuration Protocol (DHCP) allocates IP addresses dynamically, *leasing* addresses to devices so that the addresses can be reused when no longer needed. Hosts and end devices that require IP addresses obtained through DHCP must communicate with a DHCP server across the LAN.

DHCP snooping acts as a guardian of network security by keeping track of valid IP addresses assigned to downstream network devices by a trusted DHCP server (the server is connected to a trusted network port).

By default, all trunk ports on the switch are trusted and all access ports are untrusted for DHCP snooping.

When DHCP snooping is enabled, the lease information from the server is used to create the DHCP snooping table, also known as the binding table. The table shows current IP-MAC bindings, as well as lease time, type of binding, names of associated VLANs, and associated interface.

Entries in the DHCP snooping table are updated in these events:

- When a DHCP client releases an IP address (sends a DHCPRELEASE message). In this event, the associated mapping entry is deleted from the database.

- If you move a network device from one VLAN to another. In this event, typically the device needs to acquire a new IP address. Therefore, its entry in the database, including the VLAN ID, is updated.
- When the lease time (timeout value) assigned by the DHCP server expires. In this event, the associated entry is deleted from the database.



**TIP:** By default, the IP-MAC bindings are lost when the switch is rebooted and DHCP clients (the network devices, or hosts) must reacquire bindings. However, you can configure the bindings to persist by setting the `dhcp-snooping-file` statement to store the database file either locally or remotely.

You can configure the switch to snoop DHCP server responses only from specific VLANs. Doing this prevents spoofing of DHCP server messages.

### Enabling DHCP Snooping

---

DHCP snooping is not enabled in the default switch configuration. DHCP snooping is enabled automatically by Junos OS software when you configure any port security features at the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level. You enable DHCP snooping per VLAN, not per interface (port). For additional information about enabling DHCP snooping, see [“Configuring Port Security \(CLI Procedure\)” on page 4539](#).

### DHCP Snooping Process

---

The basic process of DHCP snooping consists of the following steps:



**NOTE:** When DHCP snooping is enabled for a VLAN, all DHCP packets sent from that network devices in that VLAN are subjected to DHCP snooping. The final IP-MAC binding occurs when the DHCP server sends a DHCPACK packet to the DHCP client.

1. The network device sends a DHCPDISCOVER packet to request an IP address.
2. The switch forwards the packet to the DHCP server.
3. The server sends a DHCPOFFER packet to offer an address. If the DHCPOFFER packet is from a trusted interface, the switch forwards the packet to the DHCP client.
4. The network device sends a DHCPREQUEST packet to accept the IP address. The switch adds an IP-MAC placeholder binding to the database. The entry is considered a placeholder until a DHCPACK packet is received from the server. Until then, the IP address could still be assigned to some other host.
5. The server sends a DHCPACK packet to assign the IP address or a DHCPNAK packet to deny the address request.
6. The switch updates the DHCP database in accordance with the type of packet received:

- Upon receipt of a DHCPACK packet, the switch updates lease information for the IP-MAC binding in its database.
- Upon receipt of a DHCPNACK packet, the switch deletes the placeholder.



**NOTE:** The DHCP database is updated only after the DHCPREQUEST packet has been sent.

For general information about the messages that the DHCP client and DHCP server exchange during the assignment of an IP address for the client, see the [Junos OS System Basics Configuration Guide](#).

### DHCPv6 Snooping

DHCP snooping is also supported for IPv6 packets. The process for DHCPv6 snooping is similar to that for DHCP snooping, but uses different names for the messages exchanged between the client and server to assign IPv6 addresses. [Table 495 on page 4509](#) shows DHCPv6 messages and their DHCPv4 equivalents.

**Table 495: DHCPv6 Messages and DHCPv4 Equivalent Messages**

| Sent by | DHCPv6 Messages         | DHCPv4 Equivalent Messages |
|---------|-------------------------|----------------------------|
| Client  | SOLICIT                 | DHCPDISCOVER               |
| Server  | ADVERTISE               | DHCPOFFER                  |
| Client  | REQUEST, RENEW, REBIND  | DHCPREQUEST                |
| Server  | REPLY                   | DHCPACK/DHCPNACK           |
| Client  | RELEASE                 | DHCPRELEASE                |
| Client  | INFORMATION-REQUEST     | DHCPINFORM                 |
| Client  | DECLINE                 | DHCPDECLINE                |
| Client  | CONFIRM                 | none                       |
| Server  | RECONFIGURE             | DHCPFORCERENEW             |
| Client  | RELAY-FORW, RELAY-REPLY | none                       |

### Rapid Commit for DHCPv6

DHCPv6 provides for a Rapid Commit option, which, when supported by the server and set by the client, shortens the exchange from a four-way relay to a two-message handshake. For more information about enabling Rapid Commit, see [“Enabling DHCPv6 Rapid Commit Support” on page 1428](#).

In the Rapid Commit process:

1. The DHCPv6 client sends out a SOLICIT message that contains a Rapid Commit option, requesting that rapid assignment of address/prefix and other configuration parameters are preferred.
2. If the DHCPv6 server supports rapid assignment, it responds with a REPLY message containing the assigned IPv6 address and prefix and other configuration parameters.

## DHCP Server Access

---

A switch's access to the DHCP server can be configured in three ways:

- [Switch, DHCP Clients, and DHCP Server Are All on the Same VLAN on page 4510](#)
- [Switch Acts as DHCP Server on page 4511](#)
- [Switch Acts as Relay Agent on page 4512](#)

### **Switch, DHCP Clients, and DHCP Server Are All on the Same VLAN**

When the switch, DHCP clients, and DHCP server are *all members of the same VLAN*, the DHCP server can be connected to the switch in one of two ways:



**NOTE:** To enable DHCP snooping on the VLAN, set `[edit vlans vlan-name forwarding-options] dhcp-security`.

- (See [Figure 60 on page 4511](#).) The server is directly connected to the same switch as the one connected to the DHCP clients (the hosts, or network devices, that are requesting IP addresses from the server). The VLAN is enabled for DHCP snooping to protect the untrusted access ports. The trunk port is configured by default as a trusted port.
- (See [Figure 61 on page 4511](#).) The server is connected to an intermediary switch (Switch 2) that is connected through a trunk port to the switch (Switch 1) that the DHCP clients are connected to. Switch 2 is being used as a transit switch. The VLAN is enabled for DHCP snooping to protect the untrusted access ports of Switch 1. The trunk port is configured by default as a trusted port. In [Figure 61 on page 4511](#)—, ge-0/0/11 is a trusted trunk port.

Figure 60: DHCP Server Connected Directly to Switch






Figure 61: DHCP Server Connected Directly to Switch 2, with Switch 2 Connected to Switch 1 Through a Trusted Trunk Port

**Switch Acts as DHCP Server**

You can configure DHCP local server options on the switch, which enables the switch to function as an extended DHCP local server. In [Figure 62 on page 4512](#), the DHCP clients are connected to the extended DHCP local server through untrusted access ports..

Figure 62: Switch Is the DHCP Server

**Switch Acts as Relay Agent**


The switch functions as a relay agent when the DHCP clients or the DHCP server is connected to the switch through a Layer 3 interface. The Layer 3 interfaces on the switch are configured as routed VLAN interfaces (RVIs)—also called integrated routing and bridging (IRB) interfaces. The trunk interfaces are trusted by default.

The switch can act as relay agent in these two scenarios:

- The DHCP server and clients are in different VLANs.
- The switch is connected to a router that is in turn connected to the DHCP server. See [Figure 63 on page 4513](#).



Figure 63: Switch Acting as Relay Agent Through Router to DHCP Server



### Static IP Address Additions to the DHCP Snooping Database

You can add specific static IP addresses to the database as well as have the addresses dynamically assigned through DHCP snooping. To add static IP addresses, you provide the IP address, the MAC address of the device, the interface on which the device is connected, and the VLAN with which the interface is associated. You do not assign a lease time to the entry. The statically configured entry never expires.

### Snooping DHCP Packets That Have Invalid IP Addresses

If you enable DHCP snooping on a VLAN and then devices on that VLAN send DHCP packets that request invalid IP addresses, these invalid IP addresses are stored in the DHCP snooping database until they are deleted when their default timeout is reached. To eliminate this unnecessary consumption of space in the DHCP snooping database, the switch drops the DHCP packets that request invalid IP addresses. The invalid IP addresses are:

- 0.0.0.0
- 128.0.x.x
- 191.255.x.x
- 192.0.0.x
- 223.255.255.x
- 224.x.x.x
- 240.x.x.x to 255.255.255.255

**Related Documentation**

- [Understanding Port Security on page 4497](#)
- [Configuring Port Security \(CLI Procedure\) on page 4539](#)
- [Understanding Trusted DHCP Servers for Port Security on page 4527](#)
- [Understanding DHCP Option 82 for Port Security on Switching Devices on page 4503](#)
- [Understanding DHCP Services for Switches](#)
- [DHCP/BOOTP Relay for Switches Overview](#)
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 4573](#)
- [Enabling DHCP Snooping \(J-Web Procedure\)](#)
- [Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Performance \(CLI Procedure\) on page 4575](#)

## Understanding IP Source Guard for Port Security on EX Series Switches

Ethernet LAN switches are vulnerable to attacks that involve spoofing (forging) of source IP addresses or source MAC addresses. You can use the IP source guard access port security feature on Juniper Networks EX Series Ethernet Switches to mitigate the effects of these attacks.

- [IP Address Spoofing on page 4514](#)
- [How IP Source Guard Works on page 4514](#)
- [IPv6 Source Guard on page 4515](#)
- [The DHCP Snooping Table on page 4515](#)
- [Typical Uses of Other Junos OS Features with IP Source Guard on page 4516](#)

### IP Address Spoofing

---

Hosts on access interfaces can spoof source IP addresses and source MAC addresses by flooding the switch with packets containing invalid addresses. Such attacks combined with other techniques such as TCP SYN flood attacks can cause denial-of-service (DoS) attacks. With source IP address or source MAC address spoofing, the system administrator cannot identify the source of the attack. The attacker can spoof addresses on the same subnet or on a different subnet.

### How IP Source Guard Works

---

IP source guard checks the IP source address and MAC source address in a packet sent from a host attached to an untrusted access interface on the switch against entries stored in the DHCP snooping database. If IP source guard determines that the packet header contains an invalid source IP address or source MAC address, it ensures that the switch does not forward the packet—that is, the packet is discarded.

**NOTE:**

- If your switch uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, DHCP snooping is enabled automatically when you enable IP source guard on a VLAN. See [“Configuring IP Source Guard \(CLI Procedure\)”](#) on page 4546.
- If your switch is *not* using Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style and you enable IP source guard on a VLAN, you must also explicitly enable DHCP snooping on that VLAN. Otherwise, the default value of no DHCP snooping applies to the VLAN. See *Configuring IP Source Guard (CLI Procedure)*.

IP source guard applies its checking rules to packets sent from untrusted access interfaces on those VLANs. By default, on EX Series switches, access interfaces are untrusted and trunk interfaces are trusted. IP source guard does not check packets that have been sent to the switch by devices connected to either trunk interfaces or to trusted access interfaces so that a DHCP server can be connected to that interface to provide dynamic IP addresses.



**NOTE:** IP source guard is not supported on trunk interfaces regardless of whether the trunk interface is trusted or untrusted.

## IPv6 Source Guard

IPv6 source guard is available on switches that support DHCPv6 snooping. To determine whether your switch supports DHCPv6 snooping, see the *EX Series Switch Software Features Overview*.

## The DHCP Snooping Table

IP source guard obtains information about IP address to MAC address bindings (IP-MAC binding) from the DHCP snooping table, also known as the DHCP binding table. The DHCP snooping table is populated either through dynamic DHCP snooping or through configuration of specific static IP address to MAC address bindings. For more information about the DHCP snooping table, see [“Understanding DHCP Snooping for Port Security”](#) on page 4507.

To display the DHCP snooping table, issue the operational mode command that appears in the switch CLI.

For DHCP snooping:

- (For non-ELS switches) **show ip-source-guard**
- (EX4300 switches only) **show dhcp-security binding ip-source-guard**

For DHCPv6 snooping:

- (For non-ELS switches) **show dhcpv6 snooping binding**

- (EX4300 switches only) [show dhcp-security ipv6 binding](#)

### Typical Uses of Other Junos OS Features with IP Source Guard

---

You can configure IP source guard with various other features on the EX Series switch to provide access port security, including:

- VLAN tagging (used for voice VLANs)
- GRES (graceful Routing Engine switchover)
- Virtual Chassis configurations (See *EX Series Switch Software Features Overview* for list of models that support IP Source Guard.)
- Link aggregation groups (LAGs)
- 802.1X user authentication in single supplicant, single-secure supplicant, or multiple supplicant mode.



**NOTE:** While implementing 801.X user authentication in single-secure supplicant or multiple supplicant mode, use the following configuration guidelines:

- If the 802.1X interface is part of an untagged MAC-based VLAN and you want to enable IP source guard and DHCP snooping on that VLAN, you must enable IP source guard and DHCP snooping on all dynamic VLANs in which the interface has untagged membership. This also applies to IPv6 source guard and DHCPv6 snooping.
- If the 802.1X interface is part of a tagged MAC-based VLAN and you want to enable IP source guard and DHCP snooping on that VLAN, you must enable IP source guard and DHCP snooping on all dynamic VLANs in which the interface has tagged membership. This also applies to IPv6 source guard and DHCPv6 snooping.

#### Related Documentation

- [Understanding DHCP Snooping for Port Security](#)
- [Configuring IP Source Guard \(CLI Procedure\) on page 4546](#)
- [Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN](#)
- [Example: Configuring IP Source Guard with Other EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces](#)
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529](#)

### Understanding IPv6 Neighbor Discovery Inspection

IPv6 nodes (hosts and routers) use Neighbor Discovery Protocol (NDP) to discover the presence and link-layer addresses of other nodes residing on the same link. Hosts use

NDP to find neighboring routers that are willing to forward packets on their behalf, while routers use it to advertise their presence. Nodes also use NDP to maintain reachability information about the paths to active neighbors. When a router or the path to a router fails, a host can search for alternate paths.

IPv6 nodes use NDP exchange neighbor solicitation and advertisement messages to learn the link-layer addresses of their neighbors. This process makes NDP susceptible to attacks that involve the spoofing (or forging) of link-layer addresses. An attacking node can cause packets for legitimate nodes to be sent to some other link-layer address by either sending a neighbor solicitation message with a spoofed source MAC address, or by sending a neighbor advertisement address with a spoofed target MAC address. The spoofed MAC address is then associated with a legitimate network IPv6 address by the other nodes.

IPv6 neighbor discovery inspection prevents NDP security vulnerabilities by inspecting neighbor discovery messages and verifying them against the DHCPv6 binding table. The DHCPv6 binding table is a database of valid matches, or bindings, of IP addresses to MAC addresses, also known as the DHCPv6 snooping table. With neighbor discovery inspection, the source IP address and source MAC address of the ICMPv6 packet carrying the neighbor discovery message are compared against the entries in the binding table. If no match is found, the packet is dropped.

The neighbor discovery process uses five types of ICMPv6 packets for the purposes of advertisement, solicitation, or redirection: Neighbor Solicit, Neighbor Advertise, Router Solicit, Router Advertise, and Router Redirect. Neighbor discovery inspection checks all Neighbor or Router Solicit messages and Neighbor or Router Advertise messages for their source IPv6 address and MAC address, and also checks that Router Redirect messages are sent only by trusted routers. These checks can prevent the following types of attacks:

Cache poisoning attacks—Neighbor Discovery cache poisoning is the IPv6 equivalent of ARP spoofing, in which an attacker sends an unsolicited advertisement to other hosts on the network with a forged address, to associate its own MAC address with a legitimate network IP address. These bindings between IPv6 addresses and MAC addresses are stored by each node in its neighbor cache. When the caches are updated with the malicious binding, the attacker can initiate a man-in-the-middle attack, intercepting traffic that was intended for a legitimate host.

Routing denial-of-service (DoS) attacks—An attacker could cause a host to disable its first-hop router by spoofing the address of a router and sending a neighbor advertisement message with the *router* flag cleared. The victim host assumes that the device that used to be its first-hop router is no longer a router.

Redirect attacks—Routers use ICMPv6 redirect requests to inform a host of a more efficient route to a destination. Hosts can be redirected to a better first-hop router but can also be informed by a Router Redirect message that the destination is in fact a neighbor. An attacker using this provision can achieve an effect similar to cache poisoning and intercept all traffic from the victim host.

**Related Documentation** • *IPv6 Neighbor Discovery Protocol Overview*

- [Enabling IPv6 Neighbor Discovery Inspection on page 4547](#)
- [Understanding Port Security on page 4497](#)
- [Configuring Port Security \(CLI Procedure\)](#)
- [Understanding DHCP Snooping for Port Security on page 4507](#)

## Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches

MAC limiting for port security protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You enable this feature on interfaces (ports).

MAC move limiting detects MAC movement and MAC spoofing on access interfaces. You enable this feature on VLANs.

This topic describes the various method of MAC limiting and MAC move limiting for port security:

- [MAC Limiting for Port Security by Limiting the Number of MAC Addresses That Can be Learned on Interfaces on page 4518](#)
- [MAC Limiting for Port Security by Specifying MAC Addresses That Are Allowed to Access Interfaces on page 4519](#)
- [MAC Move Limiting for Port Security by Monitoring MAC Address Moves within VLANs on page 4519](#)

### MAC Limiting for Port Security by Limiting the Number of MAC Addresses That Can be Learned on Interfaces

One method to enhance port security is to set the maximum number of MAC addresses that can be learned (added to the Ethernet switching table) on any of the following:

- A specific access interface (port)
- All access interfaces
- A specific access interface on the basis of its membership within a specific virtual LAN (VLAN membership MAC limit)



**NOTE:** Static MAC addresses do not count toward the limit you specify for dynamic MAC addresses.

---

When you are configuring the maximum MAC limit for an interface, you can choose the action that occurs on incoming packets when the MAC limit is exceeded. For additional information about configuring MAC limit for an interface, see *Configuring MAC Limiting (CLI Procedure)* or [“Configuring MAC Limiting \(CLI Procedure\)” on page 2361](#).

### MAC Limiting for Port Security by Specifying MAC Addresses That Are Allowed to Access Interfaces

Another method to enhance port security is to configure specific MAC addresses as *allowed MAC addresses* for specific access interfaces. Any MAC address that is not in the list of the configured addresses is not learned and the switch logs a message.

Allowed MAC binds MAC addresses to a VLAN so that the address does not get registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.

### MAC Move Limiting for Port Security by Monitoring MAC Address Moves within VLANs

MAC move limiting causes the switch to limit and track the frequency with which a MAC address can move to a new interface (port). It can help prevent MAC spoofing, and it can also detect and prevent loops.



**NOTE:** MAC move limiting is not supported on EX9200.

If a MAC address moves more than the configured number of times within one second, the switch performs the configured action. You can configure MAC move limiting to apply to all VLANs or to a specific VLAN.

#### Related Documentation

- [Understanding Port Security on page 4497](#)
- [Configuring MAC Limiting \(J-Web Procedure\) on page 4566](#)
- [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\)](#)
- [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\) on page 2206](#)
- [Adding a Static MAC Address Entry to the Ethernet Switching Table \(CLI Procedure\)](#)
- [Adding a Static MAC Address Entry to the Ethernet Switching Table \(CLI Procedure\) on page 2361](#)

## Understanding Media Access Control Security (MACsec)

Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication for all traffic on Ethernet links. MACsec provides point-to-point security on Ethernet links between directly connected nodes and is capable of identifying and preventing most security threats, including denial of service, intrusion, man-in-the-middle, masquerading, passive wiretapping, and playback attacks. MACsec is standardized in IEEE 802.1AE.

MACsec allows you to secure an Ethernet link for almost all traffic, including frames from the Link Layer Discovery Protocol (LLDP), Link Aggregation Control Protocol (LACP), Dynamic Host Configuration Protocol (DHCP), Address Resolution Protocol (ARP), and other protocols that are not typically secured on an Ethernet link because of limitations with other security solutions. MACsec can be used in combination with other security protocols such as IP Security (IPsec) and Secure Sockets Layer (SSL) to provide end-to-end network security.

This topic contains the following sections:

- [How MACsec Works on page 4520](#)
- [Understanding Connectivity Associations and Secure Channels on page 4521](#)
- [Understanding MACsec Security Modes on page 4521](#)
- [Understanding the Requirements to Enable MACsec on a Switch-to-Host Link on page 4523](#)
- [Understanding MACsec Hardware Requirements for EX Series and QFX Series Switches on page 4524](#)
- [Understanding MACsec Software Requirements for EX Series and QFX Series Switches on page 4524](#)
- [Understanding the MACsec Feature License Requirement on page 4525](#)
- [MACsec Limitations on page 4525](#)

---

### How MACsec Works

MACsec provides industry-standard security through the use of secured point-to-point Ethernet links. The point-to-point links are secured after matching security keys—a user-configured pre-shared key when you enable MACsec using static connectivity association key (CAK) security mode, a user-configured static secure association key when you enable MACsec using static secure association key (SAK) security mode, or a dynamic key included as part of the AAA handshake with the RADIUS server when you enable MACsec using dynamic security mode—are exchanged and verified between the interfaces at each end of the point-to-point Ethernet link. Other user-configurable parameters, such as MAC address or port, must also match on the interfaces on each side of the link to enable MACsec. See [“Configuring Media Access Control Security \(MACsec\)” on page 4548](#).

Once MACsec is enabled on a point-to-point Ethernet link, all traffic traversing the link is MACsec-secured through the use of data integrity checks and, if configured, encryption.



The data integrity checks verify the integrity of the data. MACsec appends an 8-byte header and a 16-byte tail to all Ethernet frames traversing the MACsec-secured point-to-point Ethernet link, and the header and tail are checked by the receiving interface to ensure that the data was not compromised while traversing the link. If the data integrity check detects anything irregular about the traffic, the traffic is dropped.

MACsec can also be used to encrypt all traffic on the Ethernet link. The encryption used by MACsec ensures that the data in the Ethernet frame cannot be viewed by anybody monitoring traffic on the link. MACsec encryption is optional and user-configurable; you can enable MACsec to ensure the data integrity checks are performed while still sending unencrypted data “in the clear” over the MACsec-secured link, if desired.

MACsec is configured on point-to-point Ethernet links between MACsec-capable interfaces. If you want to enable MACsec on multiple Ethernet links, you must configure MACsec individually on each point-to-point Ethernet link.

### Understanding Connectivity Associations and Secure Channels

MACsec is configured in connectivity associations. MACsec is enabled when a connectivity association is assigned to an interface.

When you are configuring MACsec using static secure association key (SAK) security mode, you must configure secure channels within a connectivity association. The secure channels are responsible for transmitting and receiving data on the MACsec-enabled link, and also responsible for transmitting SAKs across the link to enable and maintain MACsec. A single secure channel is uni-directional—it can only be used to apply MACsec to inbound or outbound traffic. A typical connectivity association when MACsec is enabled using SAK security mode contains two secure channels—one secure channel for inbound traffic and another secure channel for outbound traffic.

When you enable MACsec using static CAK or dynamic security mode, you have to create and configure a connectivity association. Two secure channels—one secure channel for inbound traffic and another secure channel for outbound traffic—are automatically created. The automatically-created secure channels do not have any user-configurable parameters; all configuration is done in the connectivity association outside of the secure channels.

### Understanding MACsec Security Modes

#### ***Understanding Static Connectivity Association Key Security Mode (Recommended Security Mode for Switch-to-Switch Links)***

When you enable MACsec using static connectivity association key (CAK) security mode, two security keys—a connectivity association key (CAK) that secures control plane traffic and a randomly-generated secure association key (SAK) that secures data plane traffic—are used to secure the point-to-point Ethernet link. Both keys are regularly exchanged between both devices on each end of the point-to-point Ethernet link to ensure link security.

You initially establish a MACsec-secured link using a pre-shared key when you are using static CAK security mode to enable MACsec. A pre-shared key includes a connectivity association name (CKN) and its own connectivity association key (CAK). The CKN and

CAK are configured by the user in the connectivity association and must match on both ends of the link to initially enable MACsec.

Once matching pre-shared keys are successfully exchanged, the MACsec Key Agreement (MKA) protocol is enabled. The MKA protocol is responsible for maintaining MACsec on the link, and decides which switch on the point-to-point link becomes the key server. The key server then creates an SAK that is shared with the switch at the other end of the point-to-point link only, and that SAK is used to secure all data traffic traversing the link. The key server will continue to periodically create and share a randomly-created SAK over the point-to-point link for as long as MACsec is enabled.

You enable MACsec using static CAK security mode by configuring a connectivity association on both ends of the link. All configuration is done within the connectivity association but outside of the secure channel. Two secure channels—one for inbound traffic and one for outbound traffic—are automatically created when using static CAK security mode. The automatically-created secure channels do not have any user-configurable parameters that cannot already be configured in the connectivity association.

We recommend enabling MACsec on switch-to-switch links using static CAK security mode. Static CAK security mode ensures security by frequently refreshing to a new random security key and by only sharing the security key between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available when you enable MACsec using static CAK security mode.

See [“Configuring Media Access Control Security \(MACsec\)” on page 4548](#) for step-by-step instructions on enabling MACsec using static CAK security mode.

### ***Understanding Dynamic Secure Association Key Security Mode (Switch-to-Host Links)***

Dynamic secure association key security mode is used to enable MACsec on a switch-to-host link.

To enable MACsec on a link connecting an endpoint device—such as a server, phone, or personal computer—to a switch, the endpoint device must support MACsec and must be running software that allows it to enable a MACsec-secured connection. When configuring MACsec on a switch-to-host link, the MACsec Key Agreement (MKA) keys, which are included as part of 802.1X authentication, are retrieved from a RADIUS server as part of the AAA handshake. A master key is passed from the RADIUS server to the switch and from the RADIUS server to the host in independent authentication transactions. The master key is then passed between the switch and the host to create a MACsec-secured connection.

A secure association using dynamic secure association security mode must be configured on the switch's Ethernet interface that connects to the host in order for the switch to create a MACsec-secured connection after receiving the MKA keys from the RADIUS server.

The RADIUS server must be using Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) in order to support MACsec. The RADIUS servers that support other widely-used authentication frameworks, such as password-only or md5, cannot be used

to support MACsec. In order to enable MACsec on a switch to secure a connection to a host, you must be using 802.1X authentication on the RADIUS server. MACsec must be configured into dynamic mode. MACsec is still enabled using connectivity associations when enabled on a switch-to-host link, as it is on a switch-to-switch link.

### ***Understanding Static Secure Association Key Security Mode (Supported for Switch-to-Switch Links)***

When you enable MACsec using static secure association key (SAK) security mode, one of up to two manually configured SAKs is used to secure data traffic on the point-to-point Ethernet link. All SAK names and values are configured by the user; there is no key server or other tool that creates SAKs. Security is maintained on the point-to-point Ethernet link by periodically rotating between the two security keys. Each security key name and value must have a corresponding matching value on the interface at the other end of the point-to-point Ethernet link to maintain MACsec on the link.

You configure SAKs within secure channels when you enable MACsec using static SAK security mode. You configure secure channels within connectivity associations. A typical connectivity association for MACsec using static SAK security mode contains two secure channels—one for inbound traffic and one for outbound traffic—that have each been configured with two manually-configured SAKs. You must attach the connectivity association with the secure channel configurations to an interface to enable MACsec using static SAK security mode.

We recommend enabling MACsec using static CAK security mode. You should only use static SAK security mode if you have a compelling reason to use it instead of static CAK security mode.

See [“Configuring Media Access Control Security \(MACsec\)” on page 4548](#) for step-by-step instructions on enabling MACsec using SAKs.

### **Understanding the Requirements to Enable MACsec on a Switch-to-Host Link**

When configuring MACsec on a switch-to-host link, the MACsec Key Agreement (MKA) keys, which are included as part of 802.1X authentication, are retrieved from a RADIUS server as part of the AAA handshake. A master key is passed from the RADIUS server to the switch and from the RADIUS server to the host in independent authentication transactions. The master key is then passed between the switch and the host to create a MACsec-secured connection.

The following requirements must be met in order to enable MACsec on a link connecting a host device to a switch.

The host device:

- must support MACsec and must be running software that allows it to enable a MACsec-secured connection with the switch.

The switch:

- must be an EX4200, EX4300, or EX4550 switch running Junos OS Release 14.1X53-D10 or later

- must be configured into dynamic secure association key security mode.
- must be using 802.1X authentication to communicate with the RADIUS server.

The RADIUS server:

- must be using the Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication framework.



**NOTE:** RADIUS servers that support other widely-used authentication frameworks, such as password-only or md5, cannot be used to support MACsec.

- must be using 802.1X authentication.
- can be multiple hops from the switch and the host device.

---

### Understanding MACsec Hardware Requirements for EX Series and QFX Series Switches

---

MACsec is currently supported on the following EX Series and QFX Series switch interfaces:

- The uplink port connections on the SFP+ MACsec uplink module that can be installed on EX4200 series switches.
- All access and uplink ports on EX4300 switches.
- All EX4550 optical interfaces that use the LC connection type. See *Pluggable Transceivers Supported on EX4550 Switches*.
- All twenty-four fixed SFP+ interfaces on an EX4600 switch.
- All eight SFP+ interfaces on the EX4600-EM-8F expansion module, when installed in an EX4600 or QFX5100-24Q switch.

MACsec can be configured on supported switch interfaces when those switches are configured in a Virtual Chassis or Virtual Chassis Fabric (VCF), including when MACsec-supported interfaces are on member switches in a mixed Virtual Chassis or VCF that includes switch interfaces that do not support MACsec. MACsec, however, cannot be enabled on Virtual Chassis ports (VCPs) to secure traffic travelling between member switches in a Virtual Chassis or VCF.

---

### Understanding MACsec Software Requirements for EX Series and QFX Series Switches

---

MACsec was initially released on EX4200, EX4300, and EX4550 switches in Junos OS Release 13.2X50-D15.

MACsec support for dynamic security mode, which allows MACsec to be configured on switch-to-host links, for EX4200, EX4300, and EX4550 switches was introduced in Junos OS Release 14.1X53-D10.

MACsec support for EX4600 switches and QFX5100-24Q switches was introduced in Junos OS Release 14.1X53-D15. The EX4600 and QFX5100-24Q switches supports MACsec on switch-to-switch links only.

The switches on each end of a MACsec-secured switch-to-switch link must either both be using Junos OS Release 14.1X53-D10 or later, or must both be using an earlier version of Junos, in order to establish a MACsec-secured connection when using static CAK security mode.

You must download the controlled version of your Junos OS software to enable MACsec. MACsec software support is not available in the domestic version of your Junos OS software. The controlled version of Junos OS software includes all features and functionality available in the domestic version of Junos OS, while also supporting MACsec. The domestic version of Junos OS software is shipped on all switches that support MACsec, so you must download and install a controlled version of Junos OS software for your switch before you can enable MACsec.

The controlled version of Junos OS software contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of the controlled version of Junos OS software is strictly controlled under United States export laws. The export, import, and use of the controlled version of Junos OS software is also subject to controls imposed under the laws of other countries. If you have questions about acquiring the controlled version of your Junos OS software, contact Juniper Networks Trade Compliance group at [compliance\\_helpdesk@juniper.net](mailto:compliance_helpdesk@juniper.net).

The process for installing a controlled version of Junos OS software on your switch is identical to installing the domestic version. See *Downloading Software Packages from Juniper Networks*.

### Understanding the MACsec Feature License Requirement

A feature license is required to configure MACsec on a switch.

To purchase a feature license for MACsec, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with a feature license file and a license key. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show chassis hardware** command.

The MACsec feature license is an independent feature license; the enhanced feature licenses (EFLs) or advanced feature licenses (AFLs) that must be purchased to enable some features on your switches cannot be purchased to enable MACsec.

### MACsec Limitations

All types of Spanning Tree Protocol frames cannot currently be encrypted using MACsec.

#### Related Documentation

- [Configuring Media Access Control Security \(MACsec\) on page 4548](#)

## Understanding Persistent MAC Learning (Sticky MAC)

Persistent MAC learning, also known as sticky MAC, is a port security feature that enables an interface to retain dynamically learned MAC addresses when the switch is restarted or if the interface goes down and is brought back online.

Persistent MAC address learning is disabled by default. You can enable persistent MAC address learning in conjunction with MAC limiting to restrict the number of persistent MAC addresses. You enable this feature on interfaces.

Configure persistent MAC learning on an interface to:

- Prevent traffic losses for trusted workstations and servers because the interface does not have to relearn the addresses from ingress traffic after a restart.
- Protect the switch against security attacks. Use persistent MAC learning in combination with MAC limiting to protect against attacks, such as Layer 2 denial-of-service (DoS) attacks, overflow attacks on the Ethernet switching table, and DHCP starvation attacks, by limiting the MAC addresses allowed while still allowing the interface to dynamically learn a specified number of MAC addresses. The interface is secured because after the limit has been reached, additional devices cannot connect to the port.

By configuring persistent MAC learning along with MAC limiting, you enable interfaces to learn MAC addresses of trusted workstations and servers from the time when you connect the interface to your network until the limit for MAC addresses is reached, and ensure that after this limit is reached, new devices will not be allowed to connect to the interface even if the switch restarts. As an alternative to using persistent MAC learning with MAC limiting, you can statically configure each MAC address on each port or allow the port to continuously learn new MAC addresses after restarts or interface-down events. Allowing the port to continuously learn MAC addresses represents a security risk.



**NOTE:** While a switch is restarting or an interface is coming back up, there might be a short delay before the interface can learn more MAC addresses. This delay occurs while the system re-enters previously learned persistent MAC addresses into the forwarding database for the interface.



**TIP:** If you move a device within your network that has a persistent MAC address entry on the switch, use the `clear ethernet-switching table persistent-mac` command to clear the persistent MAC address entry from the interface. If you move the device and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address of the device and the device will not be able to connect.

If the original port is down when you move the device, then the new port will learn the MAC address and the device can connect. However, if you do not clear the persistent MAC address on the original port, then when the port restarts, the system reinstalls the persistent MAC address in the forwarding

table for that port. If this occurs, the persistent MAC address is removed from the new port and the device loses connectivity.

Consider the following configuration guidelines when configuring persistent MAC learning:

- Interfaces must be configured in access mode (use the **port-mode** configuration statement or, for switches operating on the Enhanced Layer 2 Software (ELS) configuration style, the **interface-mode** configuration statement).
- You cannot enable persistent MAC learning on an interface on which 802.1x authentication is configured.
- You cannot enable persistent MAC learning on an interface that is part of a redundant trunk group.
- You cannot enable persistent MAC learning on an interface on which **no-mac-learning** is enabled.

**Related  
Documentation**

- [Understanding Port Security on page 4497](#)
- *Configuring Persistent MAC Learning (CLI Procedure)*
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 4569 \(ELS\)](#)

## Understanding Trusted DHCP Servers for Port Security

Any interface on the switching device that connects to a DHCP server can be configured as a trusted port. Configuring a DHCP server on a trusted port protects against rogue DHCP servers sending leases.

Ensure that the DHCP server interface is physically secure—that is, that access to the server is monitored and controlled at the site—before you configure the port as trusted.

**Related  
Documentation**

- *Understanding DHCP Snooping for Port Security*
- *Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks*
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529](#)
- *Enabling a Trusted DHCP Server (CLI Procedure)*
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 4573](#)





# Configuration

- [Configuration Examples on page 4529](#)
- [Configuration Tasks on page 4539](#)
- [Configuration Statements on page 4579](#)

## Configuration Examples

---

- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529](#)
- [Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect the Switch from IPv6 Address Spoofing on page 4534](#)

### Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing



NOTE: This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 3.



NOTE: On EX9200 switches, DHCP snooping, DAI, and IP source guard are not supported in an MC-LAG scenario.

This example describes how to enable IP source guard and Dynamic ARP Inspection (DAI) on a specified VLAN to protect the switch against spoofed IP/MAC addresses and ARP spoofing attacks. When you enable either IP source guard or DAI, the configuration automatically enables DHCP snooping for the same VLAN.

- [Requirements on page 4530](#)
- [Overview and Topology on page 4530](#)
- [Configuration on page 4532](#)
- [Verification on page 4532](#)

## Requirements

---

This example uses the following hardware and software components:

- One EX4300 switch or EX9200 switch
- Junos OS Release 13.2X50-D10 or later for EX Series switches
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure IP source guard to prevent IP/MAC spoofing or DAI to mitigate ARP spoofing attacks, be sure you have:

- Connected the DHCP server to the switch.
- Configured the VLAN to which you are adding DHCP security features. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 2337](#).

## Overview and Topology

---

Ethernet LAN switches are vulnerable to attacks on security that involve spoofing (forging) of source MAC addresses or source IP addresses. These spoofed packets are sent from hosts connected to untrusted access interfaces on the switch. These spoofed packets are sent from hosts connected to untrusted access interfaces on the switch. IP source guard checks the IP source address and MAC source address in a packet sent from a host attached to an untrusted access interface on the switch against entries stored in the DHCP snooping database. If IP source guard determines that the packet header contains an invalid source IP address or source MAC address, it ensures that the switch does not forward the packet—that is, the packet is discarded.

Another type of security attack is ARP spoofing (also known as ARP poisoning or ARP cache poisoning). ARP-spoofing is a way to initiate man-in-the-middle attacks. The attacker sends an ARP packet that spoofs the MAC address of another device on the LAN. Instead of the switch sending traffic to the proper network device, it sends it to the device with the spoofed address that is impersonating the proper device. If the impersonating device is the attacker's machine, the attacker receives all the traffic from the switch that should have gone to another device. The result is that traffic from the switch is misdirected and cannot reach its proper destination.




**NOTE:** When dynamic ARP inspection (DAI) is enabled, the switch logs the number of invalid ARP packets that it receives on each interface, along with the sender's IP and MAC addresses. You can use these log messages to discover ARP spoofing on the network.

This example shows how to configure these important port security features on a switch that is connected to a DHCP server. The setup for this example includes the VLAN **employee-vlan** on the switch. [Figure 64 on page 4531](#) illustrates the topology for this example.

**NOTE:**

The trunk interface connecting to the DHCP server interface is a trusted port by default.

**Figure 64: Network Topology for Basic Port Security**



The components of the topology for this example are shown in [Table 496 on page 4531](#).

**Table 496: Components of the Port Security Topology**

| Properties                          | Settings                                                                                       |
|-------------------------------------|------------------------------------------------------------------------------------------------|
| Switchhardware                      | One EX4300 or EX9200 switch                                                                    |
| VLAN name and ID                    | <b>employee-vlan</b> , tag 20                                                                  |
| VLAN subnets                        | 192.0.2.16/28<br>192.0.2.17 through 192.0.2.30<br>192.0.2.31 is the subnet's broadcast address |
| Interfaces in <b>employee-vlan</b>  | ge-0/0/1, ge-0/0/2, ge-0/0/3, ge-0/0/8                                                         |
| Interface connecting to DHCP server | ge-0/0/8                                                                                       |

In this example, the switch has already been configured as follows:

- All access ports are untrusted, which is the default setting.
- The trunk port (**ge-0/0/8**) is trusted, which is the default setting.

- The VLAN (**employee-vlan**) has been configured to include the specified interfaces.

---

### Configuration

To configure IP source guard and DAI (and thereby, also automatically configure DHCP snooping) to protect the switch against IP spoofing and ARP attacks:

**CLI Quick Configuration** To quickly configure IP source guard and DAI (and thereby, also automatically configure DHCP snooping), copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans employee-vlan forwarding-options dhcp-security ip-source-guard
set vlans employee-vlan forwarding-options dhcp-security arp-inspection
```

**Step-by-Step Procedure** Configure IP source guard and DAI (and thereby, also automatically configure DHCP snooping) on the VLAN:

1. Configure IP source guard on the VLAN:  

```
[edit vlans employee-vlan forwarding-options dhcp-security]
user@switch# set ip-source-guard
```
2. Enable DAI on the VLAN:  

```
[edit vlans employee-vlan forwarding-options dhcp-security]
user@switch# set arp-inspection
```

**Results** Check the results of the configuration:

```
user@switch> show vlans employee-vlan forwarding-options
employee-vlan {
  forwarding-options {
    dhcp-security {
      arp-inspection;
      ip-source-guard;
    }
  }
}
```

---

### Verification

Confirm that the configuration is working properly.

- [Verifying That DHCP Snooping Is Working Correctly on the Switch on page 4532](#)
- [Verifying That IP Source Guard is Working on the VLAN on page 4533](#)
- [Verifying That DAI Is Working Correctly on the Switch on page 4533](#)

#### ***Verifying That DHCP Snooping Is Working Correctly on the Switch***

**Purpose** Verify that DHCP snooping is working on the switch.

**Action** Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the port on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

user@switch> **show dhcp-security binding**

| IP Address | MAC Address       | Vlan          | Expires | State | Interface  |
|------------|-------------------|---------------|---------|-------|------------|
| 192.0.2.17 | 00:05:85:3A:82:77 | employee-vlan | 86265   | BOUND | ge-0/0/1.0 |
| 192.0.2.18 | 00:05:85:3A:82:79 | employee-vlan | 86265   | BOUND | ge-0/0/1.0 |
| 192.0.2.19 | 00:05:85:3A:82:80 | employee-vlan | 86287   | BOUND | ge-0/0/2.0 |
| 192.0.2.20 | 00:05:85:3A:82:81 | employee-vlan | 86287   | BOUND | ge-0/0/2.0 |
| 192.0.2.21 | 00:05:85:3A:82:83 | employee-vlan | 86287   | BOUND | ge-0/0/2.0 |
| 192.0.2.22 | 00:05:85:27:32:88 | employee-vlan | 86254   | BOUND | ge-0/0/3.0 |

**Meaning** When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for the assigned IP address, the device's MAC address, the VLAN name, and the time, in seconds, remaining before the lease expires.

#### *Verifying That IP Source Guard is Working on the VLAN*

**Purpose** Verify that IP source guard is enabled and working on the VLAN.

**Action** Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch. View the IP source guard information for the data VLAN.

user@switch> **show dhcp-security binding ip-source-guard**

| IP Address | MAC Address       | Vlan          | Expires | State | Interface  |
|------------|-------------------|---------------|---------|-------|------------|
| 192.0.2.17 | 00:05:85:3A:82:77 | employee-vlan | 86265   | BOUND | ge-0/0/1.0 |
| 192.0.2.18 | 00:05:85:3A:82:79 | employee-vlan | 86265   | BOUND | ge-0/0/1.0 |
| 192.0.2.19 | 00:05:85:3A:82:80 | employee-vlan | 86287   | BOUND | ge-0/0/2.0 |
| 192.0.2.20 | 00:05:85:3A:82:81 | employee-vlan | 86287   | BOUND | ge-0/0/2.0 |
| 192.0.2.21 | 00:05:85:3A:82:83 | employee-vlan | 86287   | BOUND | ge-0/0/2.0 |
| 192.0.2.22 | 00:05:85:27:32:88 | employee-vlan | 86254   | BOUND | ge-0/0/3.0 |

**Meaning** The IP source guard database table contains the VLANs enabled for IP source guard.

#### *Verifying That DAI Is Working Correctly on the Switch*

**Purpose** Verify that DAI is working on the switch.

**Action** Send some ARP requests from network devices connected to the switch.

Display the DAI information:

```
user@switch> show dhcp-security arp inspection statistics
```

ARP inspection statistics:

| Interface  | Packets received | ARP inspection pass | ARP inspection failed |
|------------|------------------|---------------------|-----------------------|
| ge-0/0/1.0 | 7                | 5                   | 2                     |
| ge-0/0/2.0 | 10               | 10                  | 0                     |
| ge-0/0/3.0 | 12               | 12                  | 0                     |

**Meaning** The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

- Related Documentation**
- [Configuring IP Source Guard \(CLI Procedure\) on page 4546](#)
  - [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 4572](#)
  - [Enabling Dynamic ARP Inspection \(J-Web Procedure\) on page 4572](#)

## Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect the Switch from IPv6 Address Spoofing



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

This example describes how to enable IPv6 source guard and Neighbor Discovery inspection on a specified VLAN to protect the switch against IPv6 address spoofing attacks. When you enable either IPv6 source guard or Neighbor Discovery inspection, DHCPv6 snooping is automatically enabled on the same VLAN.

- [Requirements on page 4534](#)
- [Overview and Topology on page 4535](#)
- [Configuration on page 4537](#)
- [Verification on page 4537](#)

### Requirements

---

This example uses the following hardware and software components:

- One EX4300 switch
- Junos OS Release 13.2X51-D20 or later for EX Series switches

- A DHCPv6 server to provide IPv6 addresses to network devices on the switch

Before you configure IPv6 source guard and Neighbor Discovery inspection to prevent IPv6 address spoofing attacks, be sure you have:

- Connected the DHCPv6 server to the switch.
- Configured the VLAN to which you are adding DHCPv6 security features. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 2337](#).

### Overview and Topology

Ethernet LAN switches are vulnerable to attacks on security that involve spoofing (forging) of source MAC addresses or source IPv6 addresses. These spoofed packets are sent from hosts connected to untrusted access interfaces on the switch. For more information on IPv6 address spoofing attacks, see [“Understanding IPv6 Neighbor Discovery Inspection” on page 4516](#).

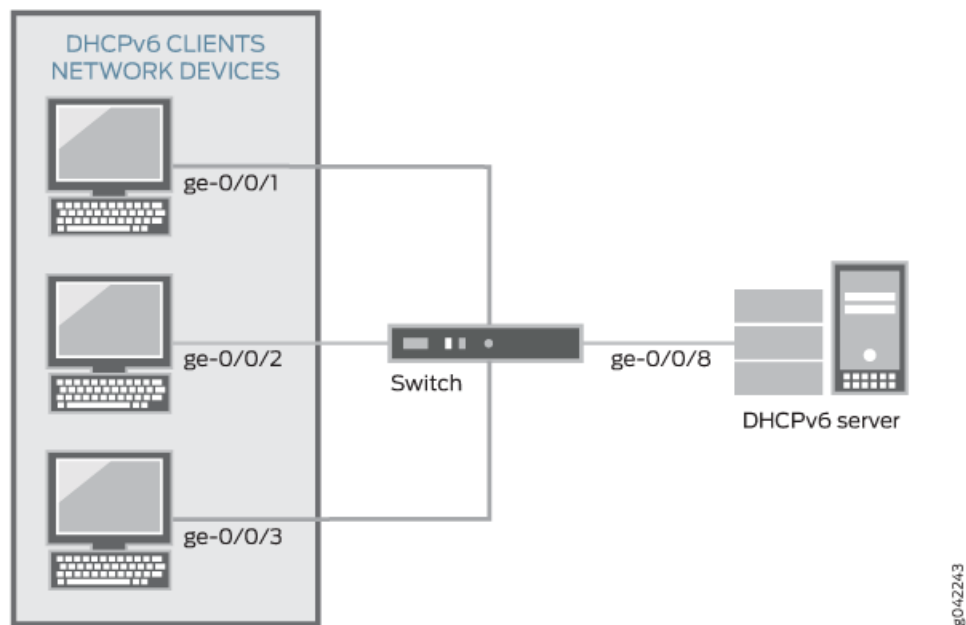
IPv6 source guard and Neighbor Discovery inspection mitigate the risk of IPv6 spoofing attacks by using the DHCPv6 snooping table. Also known as the binding table, the DHCPv6 snooping table contains the valid bindings of IPv6 addresses to MAC addresses. When a packet is sent from a host attached to an untrusted access interface on the switch, IPv6 source guard verifies the source IPv6 address and MAC address of the packet against the DHCPv6 snooping table. If there is no match in the table, the switch does not forward the packet—that is, the packet is discarded. Neighbor Discovery inspection verifies neighbor discovery messages sent between IPv6 nodes on the same network link against the DHCPv6 snooping table, and also discards the packet if no match is found.

This example shows how to configure these important port security features on a switch that is connected to a DHCPv6 server. The setup for this example includes the VLAN **sales** on the switch. [Figure 64 on page 4531](#) illustrates the topology for this example.



**NOTE:** The trunk interface connecting to the DHCPv6 server interface is a trusted port by default.

Figure 65: Network Topology for Basic Port Security



The components of the topology for this example are shown in [Table 496 on page 4531](#).

Table 497: Components of the Port Security Topology

| Properties                            | Settings                                                                                       |
|---------------------------------------|------------------------------------------------------------------------------------------------|
| Switch hardware                       | One EX4300 switch                                                                              |
| VLAN name and ID                      | <b>sales</b> , tag 20                                                                          |
| VLAN subnets                          | 192.0.2.16/28<br>192.0.2.17 through 192.0.2.30<br>192.0.2.31 is the subnet's broadcast address |
| Interfaces in <b>sales</b>            | <code>ge-0/0/1</code> , <code>ge-0/0/2</code> , <code>ge-0/0/3</code> , <code>ge-0/0/8</code>  |
| Interface connecting to DHCPv6 server | <code>ge-0/0/8</code>                                                                          |

In this example, the switch has already been configured as follows:

- All access ports are untrusted, which is the default setting.
- The trunk port (`ge-0/0/8`) is trusted, which is the default setting.
- The VLAN (**sales**) has been configured to include the specified interfaces.



## Configuration

**CLI Quick Configuration** To quickly configure IPv6 source guard and Neighbor Discovery inspection (and thereby, also automatically configure DHCPv6 snooping), copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans sales forwarding-options dhcp-security ipv6-source-guard
set vlans sales forwarding-options dhcp-security nd-inspection
```

**Step-by-Step Procedure** Configure IPv6 source guard and Neighbor Discovery inspection (and thereby, also automatically configure DHCPv6 snooping) on the VLAN:

1. Configure IPv6 source guard on the VLAN:

```
[edit vlans sales forwarding-options dhcp-security]
user@switch# set ipv6-source-guard
```

2. Enable Neighbor Discovery inspection on the VLAN:

```
[edit vlans sales forwarding-options dhcp-security]
user@switch# set neighbor-discovery-inspection
```

**Results** Check the results of the configuration:

```
user@switch> show vlans sales forwarding-options
dhcp-security {
  nd-inspection;
  ipv6-source-guard;
}
```

## Verification

Confirm that the configuration is working properly.

- [Verifying That DHCPv6 Snooping Is Working Correctly on the Switch on page 4537](#)
- [Verifying That Neighbor Discovery Inspection Is Working Correctly on the Switch on page 4538](#)

### *Verifying That DHCPv6 Snooping Is Working Correctly on the Switch*

**Purpose** Verify that DHCPv6 snooping is working on the switch.

**Action** Send DHCPv6 requests from network devices (in this example, they are DHCPv6 clients) connected to the switch.

Display the DHCPv6 snooping information when the port on which the DHCPv6 server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IPv6 addresses and leases:

```
user@switch> show dhcp-security ipv6 binding
```

| IPv6 address             | MAC address       | Vlan  | Expires | State | Interface  |
|--------------------------|-------------------|-------|---------|-------|------------|
| 2001::10:10:0:1          | 00:10:94:00:55:0b | sales | 86265   | BOUND | ge-0/0/1.0 |
| fe80::210:94ff:fe00:550b | 00:10:94:00:55:0b | sales | 86265   | BOUND | ge-0/0/1.0 |
| 2000::10:10:0:2          | 00:10:94:00:00:34 | sales | 86265   | BOUND | ge-0/0/2.0 |
| fe80::210:94ff:fe00:0034 | 00:10:94:00:00:34 | sales | 86265   | BOUND | ge-0/0/2.0 |
| 2000::10:10:0:3          | 00:10:94:00:00:55 | sales | 86265   | BOUND | ge-0/0/3.0 |
| fe80::210:94ff:fe00:0055 | 00:10:94:00:00:55 | sales | 86265   | BOUND | ge-0/0/3.0 |

**Meaning** The output shows the assigned IPv6 addresses, the MAC address, the VLAN name, and the time, in seconds, remaining before the lease expires. Because IPv6 hosts usually have more than one IPv6 address assigned to each of their IPv6-enabled network interfaces, there are two entries added for each client: one with the link-local IPv6 address, which is used by the client for DHCP transactions, and another with the IPv6 address assigned by the server. The link-local address always has the prefix **fe80::/10**.

#### *Verifying That Neighbor Discovery Inspection Is Working Correctly on the Switch*

**Purpose** Verify that Neighbor Discovery inspection is working on the switch.

**Action** Send Neighbor Discovery packets from network devices connected to the switch.

Display the Neighbor Discovery information:

```
user@switch> show dhcp-security nd-inspection statistics
```

ND inspection statistics:

| Interface  | ND Packets received | ND inspection pass | ND inspection failed |
|------------|---------------------|--------------------|----------------------|
| ge-0/0/1.0 | 7                   | 5                  | 2                    |
| ge-0/0/2.0 | 10                  | 10                 | 0                    |
| ge-0/0/3.0 | 12                  | 12                 | 0                    |

**Meaning** The sample output shows the number of Neighbor Discovery packets received and inspected per interface, with a list of how many packets passed and how many failed the inspection on each interface. The switch compares the Neighbor Discovery requests and replies against the entries in the DHCPv6 snooping database. If a MAC address or IPv6 address in the Neighbor Discovery packet does not match a valid entry in the database, the packet is dropped.

- Related Documentation**
- [Configuring IP Source Guard \(CLI Procedure\) on page 4546](#)
  - [Enabling IPv6 Neighbor Discovery Inspection on page 4547](#)
  - [Configuring Port Security \(CLI Procedure\) on page 4539](#)

## Configuration Tasks

- [Configuring Port Security \(CLI Procedure\) on page 4539](#)
- [Configuring Port Security \(J-Web Procedure\) on page 4541](#)
- [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\) on page 4545](#)
- [Configuring IP Source Guard \(CLI Procedure\) on page 4546](#)
- [Enabling IPv6 Neighbor Discovery Inspection on page 4547](#)
- [Configuring Media Access Control Security \(MACsec\) on page 4548](#)
- [Configuring MAC Limiting \(CLI Procedure\) on page 4564](#)
- [Configuring MAC Limiting \(J-Web Procedure\) on page 4566](#)
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 4568](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 4569](#)
- [Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports \(CLI Procedure\) on page 4571](#)
- [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 4572](#)
- [Enabling Dynamic ARP Inspection \(J-Web Procedure\) on page 4572](#)
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 4573](#)
- [Enabling a Trusted DHCP Server \(J-Web Procedure\) on page 4574](#)
- [Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Performance \(CLI Procedure\) on page 4575](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4577](#)
- [Enabling DHCPv6 Rapid Commit Support on page 4579](#)

## Configuring Port Security (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Port Security (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

Ethernet LANs are vulnerable to attacks such as address spoofing and Layer 2 denial of service (DoS) on network devices. DHCP port security features help protect the access

ports on the switch against the loss of information and productivity that can result from such attacks.

The following port security features are supported for DHCPv4:

- DHCP snooping
- DAI (dynamic ARP inspection)
- IP source guard
- DHCP option 82

The following port security features are supported for DHCPv6:

- DHCPv6 snooping
- Neighbor Discovery inspection
- IPv6 source guard
- DHCPv6 option 37

DHCP snooping for DHCPv4 and DHCPv6 is disabled in the default configuration. There is no explicit configuration for enabling DHCP snooping. If you configure any other port security features for a VLAN at the **[edit vlans *vlan-name* forwarding-options dhcp-security]** hierarchy level, then DHCP snooping and DHCPv6 snooping are automatically enabled on that VLAN.

DAI, Neighbor Discovery inspection, IP source guard and IPv6 source guard, and DHCP option 82 are configured per VLAN. You must configure a VLAN before configuring these DHCP port security features. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 2337](#).

The DHCP port security features that you specify for the VLAN apply to all the interfaces included within that VLAN. However, you can create a specific group of access interfaces within the VLAN to have different attributes, such as:

- Specifying an interface to have a static IP-MAC address (**static-ip** or **static-ipv6**)
- Specifying an access interface to act as a trusted interface to a DHCP server (**trusted**)
- Specifying an interface not to transmit DHCP Relay Option (**no-option-82** or **no-option37**)



NOTE:

- If you configure any of these DHCP port security features—including configuring a group of access interfaces—for a specific VLAN, the switch automatically enables DHCP snooping for that VLAN.



NOTE: Trunk interfaces are trusted by default. However, on an EX9200 switch, you can override this default behavior and set a trunk interface as **untrusted**.

---

For additional details, see:

- [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 4572](#)
- [Enabling IPv6 Neighbor Discovery Inspection on page 4547](#)
- [Configuring IP Source Guard \(CLI Procedure\) on page 4546](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4577](#)

You can override the general port security settings for the VLAN by configuring a group of access interfaces within that VLAN. For details, see:

- [Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports \(CLI Procedure\) on page 4571](#)
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 4573](#)

#### Related Documentation

- [Understanding Port Security on page 4497](#)
- [Understanding DHCP Snooping for Port Security](#)

## Configuring Port Security (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

To configure port security on an EX Series switch using the J-Web interface:

1. Select **Configure > Security > Port Security**.

The VLAN List table lists all the VLAN names, VLAN identifiers, port members, and port security VLAN features.

The Interface List table lists all the ports and indicates whether security features have been enabled on the ports.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Edit**—Click this option to modify the security features for the selected port or VLAN.

Enter information as specified in [Table 498 on page 4542](#) to modify port security settings on VLANs.

Enter information as specified in [Table 499 on page 4544](#) to modify port security settings on interfaces.

- **Activate/Deactivate**—Click this option to enable or disable security on the switch.



**NOTE:** This option is not supported on EX4300 switches.

- **Delete**—Click this option to delete the security features of the selected port or VLAN.



**NOTE:** This option is supported only on EX4300 switches.

**Table 498: Port Security Settings on VLANs**

| Field                                                                                                                                                                              | Function                                                                                                                                                                                                                                                          | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| General tab                                                                                                                                                                        |                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Enable DHCP Snooping on VLAN<br><br><b>NOTE:</b> On EX4300 switches, DHCP snooping is enabled implicitly for all VLANs if you configure <b>dhcp-security</b> on one or more VLANs. | Allows the switch to monitor and control DHCP messages received from untrusted devices connected to the switch. Builds and maintains a database of valid IP addresses/MAC address bindings. (By default, access ports are untrusted and trunk ports are trusted.) | Select to enable DHCP snooping on a specified VLAN or all VLANs.<br><br><b>TIP:</b> For private VLANs (P-VLANs), enable DHCP snooping on the primary VLAN. If you enable DHCP snooping only on a community VLAN, DHCP messages coming from P-VLAN trunk ports are not snooped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Enable ARP Inspection on VLAN                                                                                                                                                      | Uses information in the DHCP snooping database to validate ARP packets on the LAN and protect against ARP cache poisoning.                                                                                                                                        | Select to enable ARP inspection on a specified VLAN or all VLANs. (Configure any port on which you do not want ARP inspection to occur as a trusted DHCP server port.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| MAC movement                                                                                                                                                                       | Number of MAC movements allowed on the given VLAN.                                                                                                                                                                                                                | Enter a number. The default is unlimited.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| MAC movement action                                                                                                                                                                | Specifies the action to be taken if the MAC movement limit is exceeded.                                                                                                                                                                                           | Select one of the following options: <ul style="list-style-type: none"> <li>• <b>log</b>—Generate a system log entry, an SNMP trap, or an alarm.</li> <li>• <b>drop</b>—Drop the packets and generate a system log entry, an SNMP trap, or an alarm (default).</li> <li>• <b>shutdown</b>—Shut down the VLAN and generate an alarm. You can mitigate the effect of this option by configuring autorecovery from the disabled state and specifying a <b>disable timeout</b> value. See <i>Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)</i>.</li> <li>• <b>none</b>—Take no action.</li> </ul> EX4300 switches have an additional option: <ul style="list-style-type: none"> <li>• <b>drop-and-log</b>—Drop the packet and generate an alarm, an SNMP trap, or a system log entry.</li> </ul> |
| DHCP Groups                                                                                                                                                                        |                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 498: Port Security Settings on VLANs (*continued*)

|                                                                                        |                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Group Name<br><br><b>NOTE:</b> This option is supported only on EX4300 switches.       | Specifies the DHCP name of the group.                                                                             | Enter a name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Trusted<br><br><b>NOTE:</b> This option is supported only on EX4300 switches.          | Specifies trusting DHCP packets on the selected interface. By default, trunk ports are <b>dhcp-trusted</b> .      | To enable this option, select the check box.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| No Option-82<br><br><b>NOTE:</b> This option is supported only on EX4300 switches.     | Enable or disable the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server. | To enable this option, select the check box.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Interfaces<br><br><b>NOTE:</b> This option is supported only on EX4300 switches.       | Specifies the DHCP interface.                                                                                     | Select the required interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Ports                                                                                  |                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Interface<br><br><b>NOTE:</b> This option is supported only on EX4300 switches.        | Name of the interface.                                                                                            | Click the <b>Edit</b> button of the selected interface, to configure the MAC limit and the MAC limit action.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| MAC Limit<br><br><b>NOTE:</b> This option is supported only on EX4300 switches.        | Maximum number of MAC addresses learned on the interface.                                                         | Enter a number. The default is unlimited.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| MAC Limit Action<br><br><b>NOTE:</b> This option is supported only on EX4300 switches. | Specifies the action to be taken if the MAC move limit is exceeded.                                               | <p>Action to be taken when MAC limit is reached. The options are:</p> <ul style="list-style-type: none"> <li>• <b>drop</b>—Drop the packet and do not learn. Default is forward.</li> <li>• <b>drop-and-log</b>—Drop the packet and generate an alarm, an SNMP trap, or a system log entry.</li> <li>• <b>log</b>—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.</li> <li>• <b>none</b>—Forward the packet.</li> <li>• <b>shutdown</b>—Disable the interface and generate an alarm, an SNMP trap, or a system log entry.</li> </ul> |

Table 499: Port Security on Interfaces

| Field                                                                           | Function                                                                                                                                                                                                    | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Trust DHCP<br><br><b>NOTE:</b> This option is not supported on EX4300 switches. | Specifies trusting DHCP packets on the selected interface. By default, trunk ports are <b>dhcp-trusted</b> .                                                                                                | Select to enable DHCP trust.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| MAC Limit                                                                       | Specifies the number of MAC addresses that can be learned on a single Layer 2 access port. This option is not valid for trunk ports.<br><br><b>NOTE:</b> Trunk ports are supported only on EX4300 switches. | Enter a number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| MAC Limit Action                                                                | Specifies the action to be taken if the MAC limit is exceeded. This option is not valid for trunk ports.                                                                                                    | <p>Select one of the following:</p> <ul style="list-style-type: none"> <li><b>log</b>—Generate a system log entry, an SNMP trap, or an alarm.</li> <li><b>drop</b>—Drop the packets and generate a system log entry, an SNMP trap, or an alarm. (Default)</li> <li><b>shutdown</b>—Shut down the interface and generate an alarm. You can mitigate the effect of this option by configuring autorecovery from the disabled state and specifying a <b>disable timeout</b> value. See <i>Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)</i></li> <li><b>none</b>—Take no action.</li> </ul> <p>EX4300 switches have an additional option:</p> <ul style="list-style-type: none"> <li><b>drop-and-log</b>—Drop the packet and generate an alarm, an SNMP trap, or a system log entry.</li> </ul> |
| Allowed MAC List                                                                | Specifies the MAC addresses that are allowed for the interface.                                                                                                                                             | <p>To add a MAC address:</p> <ol style="list-style-type: none"> <li>Click <b>Add</b>.</li> <li>Enter the MAC address.</li> <li>Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

- Related Documentation**
- *Configuring Port Security (CLI Procedure)*
  - *Example: Configuring Basic Port Security Features*
  - *Monitoring Port Security*
  - [Understanding Port Security on page 4497](#)



## Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)



**NOTE:** This example uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switching device is an EX Series switch and runs software that does not support ELS, see *Understanding Storm Control on EX Series Switches*. If your switching device is an EX Series switch and runs software that does support ELS, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

An Ethernet switching access interface on a switching device might shut down or be disabled as a result of one of the following port-security or storm-control configurations:

- MAC limiting—(Not supported for MX Series routers) The **mac-limit** statement is configured with the **action-shutdown** statement.
- MAC move limiting—(Not supported on EX9200 switches or MX Series routers) The **mac-move-limit** statement is configured with the **action-shutdown** statement.
- Storm control—The **storm-control** statement is configured with the **action-shutdown** statement.

You can configure the switching device to automatically restore the disabled interfaces to service after a specified period of time. The specified time configured in the **recovery-timeout** statement applies to all the interfaces that have been disabled due to MAC limiting, MAC move limiting, or storm control errors.



**NOTE:** To enable autorecovery, specify the recovery timeout value for the interfaces to recover automatically. There is no default recovery timeout. If you do not specify a timeout value, you need to use the [clear ethernet-switching recovery-timeout](#) command for EX Series switches and the [clear bridge recovery-timeout](#) command for MX Series routers to clear the errors and restore the interfaces to service.

To specify the recovery timeout period for the interface:

- Set the **recovery-timeout** statement.

For EX Series switches:

```
[edit interfaces interface-name family unit 0 ethernet-switching]
user@switch# set recovery-timeout seconds
```

For MX Series routers:

```
[edit interfaces interface-name family unit 0 bridge]
user@switch# set recovery-timeout seconds
```

**Related  
Documentation**

- [Configuring MAC Limiting \(CLI Procedure\) on page 2361](#)
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 4568](#)
- [Configuring or Disabling Storm Control \(CLI Procedure\) on page 2207](#)

## Configuring IP Source Guard (CLI Procedure)

---



**NOTE:** This example uses Junos OS for MX Series routers and EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switching device runs software that does not support ELS, see [Configuring IP Source Guard \(CLI Procedure\)](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

---



**NOTE:** On EX9200 switches, IP source guard is not supported in an MC-LAG scenario.

---

You can use the IP source guard access port security feature on EX Series switches to mitigate the effects of source IP address spoofing and source MAC address spoofing. If IP source guard determines that a host connected to an access interface has sent a packet with an invalid source IP address or source MAC address in the packet header, then IP source guard ensures that the switch does not forward the packet—that is, the packet is discarded.

You configure the IP source guard feature on a specific VLAN. When you configure IP source guard on a VLAN, the switch automatically enables DHCP snooping on that VLAN.

IPv6 source guard is supported on switches with support for DHCPv6 snooping. On these switches, configuring IP source guard or IPv6 source guard on a VLAN automatically enables DHCP snooping and DHCPv6 snooping on that VLAN.

IP source guard and IPv6 source guard can be applied only to untrusted interfaces. Access interfaces are untrusted by default.

IP source guard and IPv6 source guard can be used together with 802.1X user authentication in single supplicant, single-secure supplicant, or multiple supplicant mode.

Before you can configure IP source guard or IPv6 source guard on a VLAN, you must configure the VLAN. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 2337](#).

To configure IP source guard on a specific VLAN by using the CLI:

```
[edit vlans vlan-name forwarding-options dhcp-security]
user@switch# set ip-source-guard
```

To configure IPv6 source guard on a specific VLAN by using the CLI:

```
[edit vlans vlan-name forwarding-options dhcp-security]
user@switch# set ipv6-source-guard
```

#### Related Documentation

- [Verifying That IP Source Guard Is Working Correctly](#)
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529](#)
- [Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect the Switch from IPv6 Address Spoofing on page 4534](#)
- [Understanding IP Source Guard for Port Security on EX Series Switches on page 4514](#)

## Enabling IPv6 Neighbor Discovery Inspection



**NOTE:** This procedure uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch uses software that does not support ELS, see [Configuring Port Security \(CLI Procedure\)](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

IPv6 Neighbor Discovery inspection protects switches against IPv6 address spoofing. Neighbor Discovery inspection validates IPv6 packets carrying Neighbor Discovery messages against the DHCPv6 binding table. The source IP address and source MAC address of each packet are checked against the table, and if a valid match is not found, the packet is dropped.

Before you can enable Neighbor Discovery inspection on a VLAN, you must configure the VLAN. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 2337](#).

To enable Neighbor Discovery inspection on a VLAN:

```
[edit vlans vlan-name forwarding-options dhcp-security]
user@switch# set neighbor-discovery-inspection
```



**NOTE:** DHCPv6 snooping is enabled automatically when Neighbor Discovery inspection is configured. There is no explicit configuration required for DHCPv6 snooping.

#### Related Documentation

- [Understanding IPv6 Neighbor Discovery Inspection on page 4516](#)
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529](#)
- [Understanding DHCP Snooping for Port Security on page 4507](#)

## Configuring Media Access Control Security (MACsec)

Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication for almost all types of traffic on Ethernet links. MACsec provides point-to-point security on Ethernet links between directly-connected nodes and is capable of identifying and preventing most security threats, including denial of service, intrusion, man-in-the-middle, masquerading, passive wiretapping, and playback attacks. MACsec is standardized in IEEE 802.1AE.

You can configure MACsec to secure point-to-point Ethernet links connecting EX Series or QFX Series switches, or on Ethernet links connecting a switch to a host device such as a PC, phone, or server. Each point-to-point Ethernet link that you want to secure using MACsec must be configured independently. You can enable MACsec on switch-to-switch links using static secure association key (SAK) security mode or static connectivity association key (CAK) security mode. Both processes are provided in this document.



**BEST PRACTICE:** We recommend enabling MACsec using static CAK security mode on switch-to-switch links. Static CAK security mode ensures security by frequently refreshing to a new random secure association key (SAK) and by only sharing the SAK between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available for MACsec-secured switch-to-switch connections that are enabled using static CAK security mode.

The configuration steps for both processes are provided in this document.

- [Acquiring and Downloading the Junos OS Software on page 4548](#)
- [Acquiring and Downloading the MACsec Feature License on page 4549](#)
- [Configuring the PIC Mode of the MACsec-capable Interfaces \(EX4200 switches only\) on page 4550](#)
- [Configuring MACsec Using Static Connectivity Association Key Security Mode \(Recommended for Enabling MACsec on Switch-to-Switch Links\) on page 4552](#)
- [Configuring MACsec on the Switch Using Dynamic Secure Association Key Security Mode to Secure a Switch-to-Host Link on page 4556](#)
- [Configuring MACsec Using Static Secure Association Key Security Mode to Secure a Switch-to-Switch Link on page 4560](#)

### Acquiring and Downloading the Junos OS Software

---

MACsec was initially released on EX Series switches in Junos OS Release 13.2X50-D15. MACsec was released on EX4600 and QFX5100-24Q switches in Junos OS Release 14.1X53-D15. The switches on each end of a MACsec-secured switch-to-switch link must either both be using Junos OS Release 14.1X51-D10 or later, or must both be using an earlier version of Junos, in order to establish a MACsec-secured connection when using static CAK security mode.

You must download the controlled version of your Junos OS software to enable MACsec. MACsec software support is not available in the domestic version of your Junos OS software. The controlled version of Junos OS software includes all features and functionality available in the domestic version of Junos OS, while also supporting MACsec. The domestic version of Junos OS software is shipped on all EX Series and QFX Series switches, so you must download and install a controlled version of Junos OS software on your switch before you can enable MACsec.

You can identify whether a software package is the controlled or domestic version of Junos OS by viewing the package name. A software package for a controlled version of Junos OS is named using the following format:

***package-name-m.nZx.y-controlled-signed.tgz***

A software package for a domestic version of Junos OS is named using the following format:

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

If you are unsure which version of Junos OS is running on your switch, enter the **show version** command. If the “JUNOS Crypto Software Suite” description appears in the output, you are running the controlled version of Junos OS.

The controlled version of Junos OS software contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of the controlled version of Junos OS software is strictly controlled under United States export laws. The export, import, and use of the controlled version of Junos OS software is also subject to controls imposed under the laws of other countries. If you have questions about acquiring the controlled version of your Junos OS software, contact Juniper Networks Trade Compliance group at [compliance\\_helpdesk@juniper.net](mailto:compliance_helpdesk@juniper.net).

The process for installing the controlled version of Junos OS software on your switch is identical to installing the domestic version of Junos OS software. You must enter the **request system software add** statement to download the Junos OS image, and the **request system reboot** statement to reboot the switch to complete the upgrade procedure. See *Downloading Software Packages from Juniper Networks, Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*, and *Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)* for detailed information about acquiring and installing Junos OS software images for your switches.

### Acquiring and Downloading the MACsec Feature License

A feature license is required to configure MACsec on an EX Series or a QFX Series switch.

The MACsec feature license is an independent feature license; the enhanced feature licenses (EFLs) or advanced feature licenses (AFLs) that must be purchased to enable some features on EX Series or QFX Series switches cannot be purchased to enable MACsec.

To purchase a software license for MACsec, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with a feature license file and a license key. You will be

asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show chassis hardware** command.

For a Virtual Chassis deployment, two MACsec license keys are recommended for redundancy—one for the device in the master role and the other for the device in the backup role

To add one or more new MACsec license keys on the switch, follow this procedure:

1. Add the license key or keys:
  - To add one or more license keys from a file or URL, specify the filename of the file or the URL where the key is located:

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

- To add a license key from the terminal:

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

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

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

A MACsec feature license is installed and maintained like any other switch license. See [“Managing Licenses for the EX Series Switch \(CLI Procedure\)” on page 714](#) or *Adding New Licenses (CLI Procedure)* for more detailed information on configuring and managing your MACsec software license.

### Configuring the PIC Mode of the MACsec-capable Interfaces (EX4200 switches only)

---

To configure MACsec on an EX4200 switch, you must install the SFP+ MACsec uplink module. The interfaces on the SFP+ MACsec uplink module are the only MACsec-capable interfaces available for EX4200 switches. All four ports on the uplink module are MACsec-capable.

The SFP+ MACsec uplink module provides two ports for 10-gigabit small form-factor pluggable (SFP+) transceivers when configured to operate in 10-gigabit mode or four ports for 1-gigabit small form-factor pluggable (SFP) transceivers when configured to operate in 1-gigabit mode.

The PIC mode is set to **10g**, by default. You only need to perform this procedure if you want to operate your uplink in 1-gigabit mode, or if you previously set the uplink module to 1-gigabit mode and would like to return it to 10-gigabit mode.

To configure the PIC mode:

```
[edit chassis]
```

```
user@switch# set fpc fpc-slot-number pic 1 sfpplus pic-mode (1g | 10g)
```

where *fpc-slot-number* is the FPC slot number, *pic-slot-number* is the PIC slot number, and the **[1g | 10g]** option configures the MACsec capability of the four SFP+ ports on the MACsec uplink module.

The **fpc-slot-number** is always 0 on standalone EX4200 switches, and is the member ID of the member switch in an EX4200 Virtual Chassis.

The PIC slot number is always 1 for the uplink module port slot on an EX4200 switch, so **pic 1** is always the specified PIC slot number.

The PIC mode is set to **10g** by default. When the PIC mode is set to **10g**, uplink ports 0 and 2 on the MACsec uplink module support MACsec at 10-Gbps speeds. Ports 1 and 3 cannot be used to send any traffic.

When the PIC mode is set to **1g**, all four SFP+ ports on the MACsec uplink module support MACsec at 1-Gbps speeds.

### Configuring MACsec Using Static Connectivity Association Key Security Mode (Recommended for Enabling MACsec on Switch-to-Switch Links)

---

You can enable MACsec using static connectivity association key (CAK) security mode or static secure association keys (SAK) security mode on a point-to-point Ethernet link connecting switches. This procedure shows you how to configure MACsec using static CAK security mode.



**BEST PRACTICE:** We recommend enabling MACsec using static CAK security mode on switch-to-switch links. Static CAK security mode ensures security by frequently refreshing to a new random secure association key (SAK) and by only sharing the SAK between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available for MACsec-secured switch-to-switch connections that are enabled using static CAK security mode.

When you enable MACsec using static CAK security mode, a pre-shared key is exchanged between the switches on each end of the point-to-point Ethernet link. The pre-shared key includes a connectivity association name (CKN) and a connectivity association key (CAK). The CKN and CAK are configured by the user in the connectivity association and must match on both ends of the link to initially enable MACsec.

After the pre-shared keys are exchanged and verified, the MACsec Key Agreement (MKA) protocol, which enables and maintains MACsec on the link, is enabled. The MKA is responsible for selecting one of the two switches on the point-to-point link as the key server. The key server then creates a randomized security key that is shared only with the other device over the MACsec-secured link. The randomized security key enables and maintains MACsec on the point-to-point link. The key server will continue to periodically create and share a randomly-created security key over the point-to-point link for as long as MACsec is enabled.

You enable MACsec using static CAK security mode by configuring a connectivity association on both ends of the link. All configuration is done within the connectivity association but outside of the secure channel. Two secure channels—one for inbound traffic and one for outbound traffic—are automatically created when using static CAK security mode. The automatically-created secure channels do not have any user-configurable parameters that cannot already be configured in the connectivity association.

To configure MACsec using static CAK security mode to secure a switch-to-switch Ethernet link:



1. Create a connectivity association. You can skip this step if you are configuring an existing connectivity association.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name
```

For instance, to create a connectivity association named **ca1**, enter:

```
[edit security macsec]
user@switch# set connectivity-association ca1
```

2. Configure the MACsec security mode as **static-cak** for the connectivity association:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name security-mode
static-cak
```

For instance, to configure the MACsec security mode to **static-cak** on connectivity association **ca1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 security-mode static-cak
```

3. Create the pre-shared key by configuring the connectivity association key name (CKN) and connectivity association key (CAK):

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name pre-shared-key
ckn hexadecimal-number
user@switch# set connectivity-association connectivity-association-name pre-shared-key
cak hexadecimal-number
```

A pre-shared key is exchanged between directly-connected links to establish a MACsec-secure link. The pre-shared-key includes the CKN and the CAK. The CKN is a 64-digit hexadecimal number and the CAK is a 32-digit hexadecimal number. The CKN and the CAK must match on both ends of a link to create a MACsec-secured link.



**NOTE:** To maximize security, we recommend configuring all 64 digits of a CKN and all 32 digits of a CAK.

If you do not configure all 64 digits of a CKN or all 32 digits of a CAK, however, all remaining digits will be auto-configured to 0.

After the pre-shared keys are successfully exchanged and verified by both ends of the link, the MACsec Key Agreement (MKA) protocol is enabled and manages the secure link. The MKA protocol then elects one of the two directly-connected switches as the key server. The key server then shares a random security with the other device over the MACsec-secure point-to-point link. The key server will continue to periodically create and share a random security key with the other device over the MACsec-secured point-to-point link as long as MACsec is enabled.

To configure a CKN of **37c9c2c45ddd012aa5bc8ef284aa23ff6729ee2e4acb66e91fe34ba2cd9fe311** and CAK of **228ef255aa23ff6729ee664acb66e91f** on connectivity association **ca1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 pre-shared-key ckn
37c9c2c45ddd012aa5bc8ef284aa23ff6729ee2e4acb66e91fe34ba2cd9fe311
user@switch# set connectivity-association ca1 pre-shared-key cak
228ef255aa23ff6729ee664acb66e91f
```



**NOTE:** MACsec is not enabled until a connectivity association is attached to an interface. See the final step of this procedure to attach a connectivity association to an interface.

4. (Required on switches when connecting to EX4300 switches only) Enable SCI tagging:

```
[edit security macsec connectivity-association connectivity-association-name]  
user@switch# set include-sci
```

You must enable SCI tagging on a switch that is enabling MACsec on an Ethernet link connecting to an EX4300 switch.

SCI tags are automatically appended to packets leaving a MACsec-enabled interface on an EX4300 switch. This option is, therefore, not available on EX4300 switches.

You should only use this option when enabling MACsec on a link to an EX4300 switch. SCI tags are eight octets long, so appending an SCI tag to all traffic on the link adds a significant amount of unneeded overhead.

5. (Optional) Set the MKA key server priority:

```
[edit security macsec connectivity-association connectivity-association-name]  
user@switch# set mka key-server-priority priority-number
```

Specifies the key server priority used by the MKA protocol to select the key server. The switch with the lower *priority-number* is selected as the key server.

The default *priority-number* is 16.

If the **key-server-priority** is identical on both sides of the point-to-point link, the MKA protocol selects the interface with the lower MAC address as the key server. Therefore, if this statement is not configured in the connectivity associations at each end of a MACsec-secured point-to-point link, the interface with the lower MAC address becomes the key server.

To change the key server priority to 0 to increase the likelihood that the current device is selected as the key server when MACsec is enabled on the interface using connectivity association **ca1**:

```
[edit security macsec connectivity-association ca1]  
user@switch# set mka key-server-priority 0
```

To change the key server priority to 255 to decrease the likelihood that the current device is selected as the key server in connectivity association **ca1**:

```
[edit security macsec connectivity-association ca1]  
user@switch# set mka key-server-priority 255
```

6. (Optional) Set the MKA transmit interval:

```
[edit security macsec connectivity-association connectivity-association-name]  
user@switch# set mka transmit-interval interval
```

The MKA transmit interval setting sets the frequency for how often the MKA protocol data unit (PDU) is sent to the directly connected device to maintain MACsec connectivity on the link. A lower *interval* increases bandwidth overhead on the link; a higher *interval* optimizes MKA protocol communication.

The default *interval* is 2000ms. We recommend increasing the interval to 6000 ms in high-traffic load environments. The transmit interval settings must be identical on both ends of the link when MACsec using static CAK security mode is enabled.

For instance, if you wanted to increase the MKA transmit interval to 6000 milliseconds when connectivity association *ca1* is attached to an interface:

```
[edit security macsec connectivity-association ca1]
user@switch# set mka transmit-interval 6000
```

7. (Optional) Disable MACsec encryption:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set no-encryption
```

Encryption is enabled for all traffic entering or leaving the interface when MACsec is enabled using static CAK security mode, by default.

When encryption is disabled, traffic is forwarded across the Ethernet link in clear text. You are able to view unencrypted data in the Ethernet frame traversing the link when you are monitoring it. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic sent or received on the link has not been tampered with and does not represent a security threat.

8. (Optional) Set an offset for all packets traversing the link:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set offset (0 | 30 | 50)
```

For instance, if you wanted to set the offset to 30 in the connectivity association named *ca1*:

```
[edit security macsec connectivity-association ca1]
user@switch# set offset 30
```

The default offset is 0. All traffic in the connectivity association is encrypted when encryption is enabled and an **offset** is not set.

When the offset is set to 30, the IPv4 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic. When the offset is set to 50, the IPv6 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

9. (Optional) Enable replay protection.

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set replay-protect replay-window-size number-of-packets
```

When MACsec is enabled on a link, an ID number is assigned to each packet on the MACsec-secured link.

When replay protection is enabled, the receiving interface checks the ID number of all packets that have traversed the MACsec-secured link. If a packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance,

if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

You can require that all packets arrive in order by setting the replay window size to 0.

To enable replay protection with a window size of five on connectivity association **ca1**:

```
[edit security macsec connectivity-association ca1]
user@switch# set replay-protect replay-window-size 5
```

10. (Optional) Exclude a protocol from MACsec:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set exclude-protocol protocol-name
```

For instance, if you did not want Link Level Discovery Protocol (LLDP) to be secured using MACsec:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set exclude-protocol lldp
```

When this option is enabled, MACsec is disabled for all packets of the specified protocol—in this case, LLDP—that are sent or received on the link.

11. Assign the connectivity association to an interface:

```
[edit security macsec]
user@switch# set interfaces interface-names connectivity-association
connectivity-association-name
```

Assigning the connectivity association to an interface is the final configuration step to enabling MACsec on an interface.

For instance, to assign connectivity association **ca1** to interface **xe-0/0/1**:

```
[edit security macsec]
user@switch# set interfaces xe-0/1/0 connectivity-association ca1
```

MACsec using static CAK security mode is not enabled until a connectivity association on the opposite end of the link is also configured, and contains pre-shared keys that match on both ends of the link.

---

### Configuring MACsec on the Switch Using Dynamic Secure Association Key Security Mode to Secure a Switch-to-Host Link

---

Before you begin to enable MACsec on a switch-to-host link:

- Configure a RADIUS server. The RADIUS server:
  - must be configured as the user database for 802.1X authentication.
  - must be using the Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication framework.

- must have connectivity to the switch and to the host. The RADIUS server can be multiple hops from the switch or the host.

See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).

- Enable MACsec on the host device.

The procedures for enabling MACsec on the host device varies by host device, and is beyond the scope of this document.

To configure MACsec using dynamic security mode to secure a switch-to-host Ethernet link:

1. Create a connectivity association. You can skip this step if you are configuring an existing connectivity association.

```
[edit security macsec]
```

```
user@switch# set connectivity-association connectivity-association-name
```

For instance, to create a connectivity association named `ca-dynamic1`, enter:

```
[edit security macsec]
```

```
user@switch# set connectivity-association ca-dynamic1
```

2. Configure the MACsec security mode as dynamic for the connectivity association:

```
[edit security macsec]
```

```
user@switch# set connectivity-association connectivity-association-name security-mode dynamic
```

For instance, to configure the MACsec security mode to dynamic on connectivity association `ca-dynamic1`:

```
[edit security macsec]
```

```
user@switch# set connectivity-association ca-dynamic1 security-mode dynamic
```

3. (Optional) Configure the **must-secure** option:

```
[edit security macsec]
```

```
user@switch# set connectivity-association connectivity-association-name mka must-secure
```

When the **must-secure** option is enabled, all traffic that is not MACsec-secured that is received on the interface is dropped.

When the **must-secure** option is disabled, all traffic from devices that support MACsec is MACsec-secured while traffic received from devices that do not support MACsec is forwarded through the network.

The **must-secure** option is particularly useful in scenarios where multiple devices, such as a phone and a PC, are accessing the network through the same Ethernet interface. If one of the devices supports MACsec while the other device does not support MACsec, the device that doesn't support MACsec can continue to send and receive traffic over the network—provided the **must-secure** option is disabled—while traffic to and from the device that supports MACsec is MACsec-secured. In this scenario, traffic to the device that is not MACsec-secured must be VLAN-tagged.

The **must-secure** option is disabled, by default.

4. (Required only if the host device requires SCI tagging) Enable SCI tagging:

```
[edit security macsec connectivity-association connectivity-association-name]
```

```
user@switch# set include-sci
```

You should only use this option when connecting a switch to a host that requires SCI tags. SCI tags are eight octets long, so appending an SCI tag to all traffic on the link adds a significant amount of unneeded overhead.

5. (Optional) Set the MKA key server priority:

```
[edit security macsec connectivity-association connectivity-association-name]  
user@switch# set mka key-server-priority priority-number
```

Specifies the key server priority used by the MKA protocol to select the key server. The switch with the lower *priority-number* is selected as the key server.

The default *priority-number* is 16. If the **key-server-priority** is identical on both sides of the point-to-point link, the MKA protocol selects the interface with the lower MAC address as the key server. Therefore, if this statement is not configured in the connectivity associations at each end of a MACsec-secured point-to-point link, the interface with the lower MAC address becomes the key server.

To change the key server priority to 0 to increase the likelihood that the current device is selected as the key server when MACsec is enabled on the interface using connectivity association *ca1*:

```
[edit security macsec connectivity-association ca-dynamic1]  
user@switch# set mka key-server-priority 0
```

To change the key server priority to 255 to decrease the likelihood that the current device is selected as the key server in connectivity association *ca-dynamic1*:

```
[edit security macsec connectivity-association ca-dynamic1]  
user@switch# set mka key-server-priority 255
```

6. (Optional) Set the MKA transmit interval:

```
[edit security macsec connectivity-association connectivity-association-name]  
user@switch# set mka transmit-interval interval
```

The MKA transmit interval setting sets the frequency for how often the MKA protocol data unit (PDU) is sent to the directly connected device to maintain MACsec connectivity on the link. A lower interval increases bandwidth overhead on the link; a higher interval optimizes MKA protocol communication.

The default interval is 2000ms. We recommend increasing the interval to 6000 ms in high-traffic load environments. The transmit interval settings must be identical on both ends of the link.

For instance, if you wanted to increase the MKA transmit interval to 6000 milliseconds when connectivity association *ca-dynamic1* is attached to an interface:

```
[edit security macsec connectivity-association ca-dynamic1]  
user@switch# set mka transmit-interval 6000
```

7. (Optional) Disable MACsec encryption:

```
[edit security macsec connectivity-association connectivity-association-name]  
user@switch# set no-encryption
```

Encryption is enabled for all traffic entering or leaving the interface when MACsec is enabled using dynamic security mode, by default. When encryption is disabled, traffic is forwarded across the Ethernet link in clear text. You are able to view unencrypted data in the Ethernet frame traversing the link when you are monitoring it. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are

run on both ends of the link to ensure the traffic sent or received on the link has not been tampered with and does not represent a security threat.

8. (Optional) Set an offset for all packets traversing the link:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set offset (0 | 30 | 50)
```

For instance, if you wanted to set the offset to 30 in the connectivity association named ca-dynamic1:

```
[edit security macsec connectivity-association ca-dynamic1]
user@switch# set offset 30
```

The default offset is 0. All traffic in the connectivity association is encrypted when encryption is enabled and an offset is not set.

When the offset is set to 30, the IPv4 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic. When the offset is set to 50, the IPv6 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

9. (Optional) Enable replay protection.

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set replay-protect replay-window-size number-of-packets
```

When MACsec is enabled on a link, an ID number is assigned to each packet on the MACsec-secured link. When replay protection is enabled, the receiving interface checks the ID number of all packets that have traversed the MACsec-secured link. If a packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

You can require that all packets arrive in order by setting the replay window size to 0.

To enable replay protection with a window size of five on connectivity association ca-dynamic1:

```
[edit security macsec connectivity-association ca-dynamic1]
user@switch# set replay-protect replay-window-size 5
```

10. (Optional) Exclude a protocol from MACsec:

```
[edit security macsec connectivity-association connectivity-association-name]
```

```
user@switch# set exclude-protocol protocol-name
```

For instance, if you did not want Link Level Discovery Protocol (LLDP) to be secured using MACsec:

```
[edit security macsec connectivity-association ca-dynamic1]
user@switch# set exclude-protocol lldp
```

When this option is enabled, MACsec is disabled for all packets of the specified protocol—in this case, LLDP—that are sent or received on the link.

11. Assign the connectivity association to an interface:

```
[edit security macsec]
user@switch# set interfaces interface-names connectivity-association
connectivity-association-name
```

Assigning the connectivity association to an interface is the final configuration step to enabling MACsec on an interface. For instance, to assign connectivity association `ca-dynamic1` to interface `xe-0/0/1`:

```
[edit security macsec]
user@switch# set interfaces xe-0/1/0 connectivity-association ca-dynamic1
```

### Configuring MACsec Using Static Secure Association Key Security Mode to Secure a Switch-to-Switch Link

---

When you enable MACsec using static secure association key (SAK) security mode, one of up to two manually configured security keys is used to secure the point-to-point Ethernet link between the switches. All security key names and values are configured by the user; there is no key server or other tool that creates security keys. Security is maintained on the point-to-point Ethernet link by periodically rotating the security keys. Each security key name and value must have a corresponding matching value on the interface at the other end of the point-to-point Ethernet link to maintain MACsec on the link.

You configure static SAKs within secure channels when you are enabling MACsec using static SAK security mode. You configure secure channels within connectivity associations. A typical connectivity association for MACsec using static SAK security mode contains two secure channels—one for inbound traffic and one for outbound traffic—that have each been configured with two static SAKs. You must attach the connectivity association with the secure channel configurations to an interface to enable MACsec using static SAK security mode.

To configure MACsec on a switch-to-switch Ethernet link using static SAK security mode:

1. Create a connectivity association. You can skip this step if you are configuring an existing connectivity association.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name
```

For instance, to create a connectivity association named `ca1`, enter:

```
[edit security macsec]
user@switch# set connectivity-association ca1
```

2. Configure the MACsec security mode as **static-sak** for the connectivity association:

```
[edit security macsec]
```



```
user@switch# set connectivity-association connectivity-association-name security-mode
static-sak
```

For instance, to configure the MACsec security mode to **static-sak** on connectivity association **ca1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 security-mode static-sak
```

3. Create a secure channel within the connectivity association. You can skip this step if you are configuring an existing secure channel.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name
```

For instance, to create secure channel **sc1** in connectivity association **ca1**, enter:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1
```

4. Define the security associations and the static SAKs for the secure channel:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name security-association number key key-string
```

where the **security-association number** is a number between 0 and 3, and the **key-string** is a 32-digit key defined statically by the network administrator.

The key string is a 32-digit hexadecimal number. The key string and the security association must match on both sides of an Ethernet connection to secure traffic using MACsec.

A secure channel must have at least two security associations with unique key strings. MACsec uses a security associations to establish a secure communications link, and periodically rotates to a new security association to keep the link secure. MACsec, therefore, must have at least one backup security association and key at all times.

To create one secure channel with two security associations and keys, for example:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 security-association 0 key
d183c4002fa6fe3d2d9a852c20ab8412
user@switch# set connectivity-association ca1 secure-channel sc1 security-association 1 key
b976c7494ab6fe2f2d4c432a90fd90a8
```

5. Specify whether the secure channel should be applied to traffic entering or leaving the switch:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name direction [inbound | outbound]
```

where **inbound** applies the secure channel to traffic entering the switch, and **outbound** applies the secure channel to traffic leaving the switch.



**NOTE:** A secure channel can only be applied to traffic entering (inbound) or leaving (outbound) an interface on the switch.

If you need to configure MACsec using SAKs on inbound and outbound traffic on the same interface, you must configure a connectivity association with one secure channel for inbound traffic and a second secure channel for outbound traffic. The connectivity association is assigned to an interface later in this process.

For instance, to configure secure channel **sc1** to apply MACsec to incoming traffic:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 direction inbound
```

To configure secure channel **sc2** to apply MACsec to outgoing traffic:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc2 direction outbound
```

6. Specify a MAC address:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name id mac-address mac-address
```

If you are configuring a MAC address on a secure channel in the outbound direction, you should specify the MAC address of the interface as the **mac-address**.

If you are configuring a MAC address on a secure channel in the inbound direction, you should specify the MAC address of the interface at the other end of the link as the **mac-address**.

The **mac-address** variables must match on the sending and receiving secure channel on each side of a link to enable MACsec using static SAK security mode.



**NOTE:** You can see the MAC address of an interface in the **show interfaces** output.

To configure MACsec to accept frames from MAC address **12:34:56:ab:cd:ef** on secure channel **sc1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 id mac-address
12:34:56:ab:cd:ef
```

7. Specify a port:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name id port-id port-id-number
```

The **port-id-number** variables must match on a sending and receiving secure channel on each side of a link to enable MACsec.



**NOTE:** The only requirement for port numbers in this implementation of MACsec is that they match on the sending and receiving ends of an Ethernet link. When the port numbers match, MACsec is enabled for all traffic on the connection.

To specify port ID 4 on secure channel **sc1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 id port-id 4
```

8. (Optional) Enable encryption:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name encryption
```

You can enable MACsec without enabling encryption. If a secure channel is configured on an interface without encryption, traffic is forwarded across the Ethernet link in clear text, and you will be able to view unencrypted data in the Ethernet frame traversing the link when you are monitoring it. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic on the link does not represent a security threat.

Encryption is disabled by default when you are enabling MACsec using static SAK security mode. To ensure all traffic traversing secure-channel **sc1** is encrypted:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 encryption
```

9. (Optional) Set an offset to send the first 30 or 50 octets in unencrypted plain text when encryption is enabled.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name offset [0 | 30 | 50]
```

When the offset is set to 30, the IPv4 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic. When the offset is set to 50, the IPv6 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

The default offset is 0, so all traffic on the link is encrypted when the **encryption** option is enabled and an **offset** is not set.

To change the offset to 30 for secure channel **sc1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 offset 30
```

10. Assign the connectivity association to an interface:

```
[edit security macsec]
user@switch# set interfaces interface-names connectivity-association
connectivity-association-name
```

Assigning the connectivity association to an interface is the final configuration step to enabling MACsec on an interface.

For instance, to assign connectivity association `ca1` to interface `xe-0/0/1`:

```
[edit security macsec]
user@switch# set interfaces xe-0/1/0 connectivity-association ca1
```

MACsec using static SAK security mode is not enabled until a connectivity association on the opposite end of the link is also configured, and the configuration match on both ends of the link.

**Related  
Documentation**

- [Understanding Media Access Control Security \(MACsec\) on page 4520](#)

## Configuring MAC Limiting (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches and QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring MAC Limiting (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

This topic describes various ways of configuring a limitation on MAC addresses in packets that are received and forwarded by the switch.



**NOTE:** On a QFX Series Virtual Chassis, if you include the `shutdown` option at the `[edit vlans vlan-name switch-options interface interface-name interface-mac-limit packet-action]` hierarchy level and issue the `commit` operation, the system generates a commit error. The system does not generate an error if you include the `shutdown` option at the `[edit switch-options interface interface-name interface-mac-limit packet-action]` hierarchy level.

The different ways of setting a MAC limit are described in the following sections:

- [Limiting the Number of MAC Addresses Learned by an Interface on page 4565](#)
- [Limiting the Number of MAC Addresses Learned by a VLAN on page 4565](#)

### Limiting the Number of MAC Addresses Learned by an Interface

To secure a port, you can set the maximum number of MAC addresses that can be learned by an interface:

- Set the MAC limit on an interface, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit switch-options]
user@switch# set interface interface-name interface-mac-limit limit packet-action action
```

After you set a new MAC limit for the interface, the system clears existing entries in the MAC address forwarding table associated with the interface.

### Limiting the Number of MAC Addresses Learned by a VLAN

To limit the number of MAC addresses learned by a VLAN, perform both of the following steps:

- Set the maximum number of MAC addresses that can be learned by a VLAN, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit vlans]
user@switch# set vlan-name switch-options mac-table-size limit packet-action action
```

- Set the maximum number of MAC addresses that can be learned by one or all interfaces in the VLAN, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit vlans]
user@switch# set vlan-name switch-options interface interface-name
interface-mac-limit limit packet-action action
[edit vlans]
user@switch# set vlan-name switch-options interface-mac-limit limit packet-action action
```



**NOTE:** If you specify a MAC limit and packet action for all interfaces in the VLAN *and* a specific interface in the VLAN, the MAC limit and packet action specified at the specific interface level takes precedence. Also, at the VLAN interface level, only the drop and drop-and-log options are supported.

After you set new MAC limits for a VLAN by using the **mac-table-size** statement or for interfaces associated with a VLAN by using the **interface-mac-limit** statement, the system clears the corresponding existing entries in the MAC address forwarding table.

#### Related Documentation

- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 4569](#)

## Configuring MAC Limiting (J-Web Procedure)

MAC limiting protects against flooding of the Ethernet switching table on an EX Series switch. MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface (port).

Junos OS provides two MAC limiting methods:

- Maximum number of dynamic MAC addresses allowed per interface—If the limit is exceeded, incoming packets with new MAC addresses are dropped.
- Specific “allowed” MAC addresses for the access interface—Any MAC address that is not in the list of configured addresses is not learned.

You configure MAC limiting for each interface, not for each VLAN. You can specify the maximum number of dynamic MAC addresses that can be learned on a single Layer 2 access interface or on all Layer 2 access interfaces. The default action that the switch will take if that maximum number is exceeded is **drop**—drop the packet and generate an alarm, an SNMP trap, or a system log entry.

To enable MAC limiting on one or more interfaces using the J-Web interface:

1. Select **Configure>Security>Port Security**.
2. Select one or more interfaces from the **Interface List**.
3. Click the **Edit** button. If a message appears asking whether you want to enable port security, click **Yes**.
4. To set a dynamic MAC limit:
  1. Type a limit value in the **MAC Limit** box.
  2. Select an action from the **MAC Limit Action** box (optional). The switch takes this action when the MAC limit is exceeded. If you do not select an action, the switch applies the default action, **drop**.
    - Log—Generate a system log entry.
    - Drop—Drop the packets and generate a system log entry. (Default)
    - Shutdown—Shut down the VLAN and generate a system log entry. You can mitigate the effect of this option by configuring the switch for autorecovery from the disabled state and specifying a **disable timeout** value. See *Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*. If you have not configured autorecovery from the disabled state, you can bring up the interfaces by running the **clear ethernet-switching port-error** command.
    - None— No action to be taken.
5. To add allowed MAC addresses:
  1. Click **Add**.
  2. Type the allowed MAC address and click **OK**.

Repeat this step to add more allowed MAC addresses.

6. Click **OK** when you have finished setting MAC limits.
7. Click **OK** after the configuration has been successfully delivered.



**NOTE:** You can enable or disable port security on the switch at any time by clicking the **Activate** or **Deactivate** button on the Port Security Configuration page. If security status is shown as **Disabled** when you try to edit settings for any VLANs or interfaces (ports), a message asking whether you want to enable port security appears.

**Related  
Documentation**

- *Configuring MAC Limiting (CLI Procedure)*
- *Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks*
- *Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks*
- *Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks*
- *Verifying That MAC Limiting Is Working Correctly*
- *Setting the none Action on an Interface to Override a MAC Limit Applied to All Interfaces (CLI Procedure)*
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches on page 4518](#)

## Configuring MAC Move Limiting (CLI Procedure)



**NOTE:** This topic uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring MAC Move Limiting (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

When MAC move limiting is configured, MAC address movements are tracked by the switch and, if a MAC address changes more than the configured number of times within one second, the changes to MAC addresses are dropped, logged or ignored, or the interface is shut down.



**NOTE:** Although you enable this feature on VLANs, the MAC move limitation pertains to the number of movements for each individual MAC address rather than the total number of MAC address moves in the VLAN. For example, if the MAC move limit is set to 1, the switch allows an unlimited number of MAC address movements within the VLAN as long as the same MAC address does not change more than once.

You configure MAC move limiting per VLAN, not per interface (port). In the default configuration, the number of MAC moves permitted is unlimited.

You can choose to have one of the following actions performed when the MAC move limit is exceeded:



**NOTE:** There is no default action.

- **drop**—Drop the packet and do not generate an alarm.
- **drop-and-log**—Drop the packet and generate an alarm, an SNMP trap, or system log entry.
- **log**—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.
- **none**—Take no action.
- **shutdown**—Disable the interfaces in the VLAN and generate an alarm, an SNMP trap, or a system log entry. If you have configured an interface with the [recovery-timeout](#) statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the [clear ethernet-switching recovery-timeout](#) command.



To configure a MAC move limit for MAC addresses within a specific VLAN:

- To limit the number of MAC address movements that can be made by an individual MAC address within the specified VLAN:

```
[edit edit vlans vlan-name switch-options]
user@switch# set mac-move-limit limit
```

- To limit the number of MAC address movements that can be made by an individual MAC address and to specify the action to be taken when the limit is reached:

```
[edit edit vlans vlan-name switch-options]
user@switch# set mac-move-limit limit packet-action action
```

The switch performs the specified action if it tracks that an individual MAC address within the specified VLAN has moved more than the specified number of times within one second.

#### Related Documentation

- [Configuring MAC Move Limiting \(J-Web Procedure\)](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 4569](#)

## Configuring Persistent MAC Learning (CLI Procedure)



**NOTE:** This topic uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Configuring Persistent MAC Learning \(CLI Procedure\)](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

Persistent MAC address learning is disabled by default. You can enable it to:

- Help prevent traffic losses for trusted workstations and servers because the interface does not have to relearn the addresses from ingress traffic after a restart.
- Protect the switch against security attacks—Use persistent MAC learning in combination with MAC limiting to protect against attacks while still avoiding the need to statically configure MAC addresses. When the initial learning of MAC addresses up to the number specified by the MAC limit is done, new addresses will not be allowed even after a restart. The port is secured because after the limit has been reached, additional devices cannot connect to the interface.

To configure persistent MAC learning on an interface and limit the number of allowed MAC addresses:

1. Enable persistent MAC learning on an interface **ge-x/y/z** where x, y, and z are variables:

```
[edit “switch-options” on page 2212]
user@switch# set interface ge-x/y/z persistent-learning
```

2. To configure a MAC move limit for MAC addresses within a specific VLAN:

- To limit the number of MAC address movements that can be made by an individual MAC address within the specified VLAN:

```
[edit vlans vlan-name switch-options]
user@switch# set mac-move-limit limit
```

The action is not specified, so the switch performs the default action **drop** if it tracks that an individual MAC address within the specified VLAN has moved more than the specified number of times within one second.

- To limit the number of MAC address movements that can be made by an individual MAC address and to specify the action to be taken when the limit is reached:

```
[edit vlans vlan-name switch-options]
user@switch# set mac-move-limit limit packet-action action
```

Values for *action* are:

**drop**—Drop packets with new source MAC addresses, and do not learn the new source MAC addresses.

**drop-and-log**—(EX Series switches only) Drop packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

**log**—(EX Series switches only) Hold packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

**none**—(EX Series switches only) Forward packets with new source MAC addresses, and learn the new source MAC address.

**shutdown**—(EX Series switches only) Disable the specified interface, and generate an alarm, an SNMP trap, or a system log entry.



**TIP:** If you move a device within your network that has a persistent MAC address entry on the switch, use the [clear ethernet-switching table persistent-mac](#) command to clear the persistent MAC address entry from the interface. If you move the device and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address of the device and the device will not be able to connect.

If the original port is down when you move the device, then the new port will learn the MAC address and the device can connect. However, if you do not clear the persistent MAC address on the original port, then when the port restarts, the system reinstalls the persistent MAC address in the forwarding table for that port. If this occurs, the persistent MAC address is removed from the new port and the device loses connectivity.

---

**Related  
Documentation**

- [Configuring Persistent MAC Learning \(CLI Procedure\)](#)
- [Configuring MAC Move Limiting \(CLI Procedure\)](#) on page 4568
- [Understanding Persistent MAC Learning \(Sticky MAC\)](#) on page 4526

## Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports (CLI Procedure)



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

You can add static (fixed) IP addresses and bind them to fixed MAC addresses in the DHCP snooping database. These bindings are labeled as *static* in the database, while those bindings that have been added through the process of DHCP snooping are labeled *dynamic*. Static IPv6 address assignment is also available for DHCPv6.

Before you can perform this procedure, you must configure the VLAN. See “[Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)” on page 2337.

To configure a static IP address to MAC address (IP-MAC) binding in the DHCP snooping database, you must first create a group of access interfaces under **[edit vlans *vlan-name* forwarding-options dhcp-security]**. Creating this group automatically enables DHCP snooping, which is a prerequisite for creating the DHCP snooping database. You can then configure a specific interface within the group to have a static IP address that is bound to a fixed MAC address. If you want to have multiple static IP addresses, configure additional interfaces within the same group.



**NOTE:** On switches that support DHCPv6, creating the group of interfaces will automatically enable both DHCP and DHCPv6 snooping.

To configure a static IP-MAC address binding in the DHCP snooping database:

- **[edit vlans *vlan-name* forwarding-options dhcp-security]**  
`user@switch# set group group-name interface interface-name static-ip ip-address mac mac-address`

To configure a static IPv6-MAC address binding in the DHCPv6 snooping database:

- **[edit vlans *vlan-name* forwarding-options dhcp-security]**  
`user@switch# set group group-name interface interface-name static-ipv6 ip-address mac mac-address`

### Related Documentation

- [show dhcp-security binding on page 4668](#)
- [Verifying That DHCP Snooping Is Working Correctly](#)
- [Understanding DHCP Snooping for Port Security](#)

## Enabling Dynamic ARP Inspection (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Enabling Dynamic ARP Inspection (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.



**NOTE:** On EX9200 switches, DHCP snooping, DAI, and IP source guard are not supported in an MC-LAG scenario.

Dynamic ARP inspection (DAI) protects switches against ARP spoofing. DAI inspects ARP packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP cache poisoning.

Before you can enable DAI on a VLAN, you must configure the VLAN. See “[Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)” on page 2337.

To enable DAI on a VLAN by using the CLI:

```
[edit vlans vlan-name forwarding-options dhcp-security]  
user@switch# set arp-inspection
```

### Related Documentation

- [Enabling Dynamic ARP Inspection \(J-Web Procedure\) on page 4572](#)
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529](#)
- [Understanding DAI for Port Security on page 4500](#)

## Enabling Dynamic ARP Inspection (J-Web Procedure)

Dynamic ARP inspection (DAI) protects EX Series switches against ARP spoofing. DAI inspects ARP packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP cache poisoning.

You configure DAI for each VLAN, not for each interface (port). By default, DAI is disabled for all VLANs.

To enable DAI on one or more VLANs by using the J-Web interface:

1. Select **Configure>Security>Port Security**.
2. Select one or more VLANs from the VLAN list.
3. Click the **Edit** button. If a message appears asking if you want to enable port security, click **Yes**.

4. Select the **Enable ARP Inspection on VLAN** check box and then click **OK**.
5. Click **OK** after the command has been successfully delivered.



**NOTE:** You can enable or disable port security on the switch at any time by clicking the **Activate** or **Deactivate** button on the Port Security Configuration page. If security status is shown as **Disabled** when you try to edit settings for any VLANs or interfaces (ports), the message asking if you want to enable port security appears.

#### Related Documentation

- [Enabling Dynamic ARP Inspection \(CLI Procedure\)](#)
- [Example: Configuring Basic Port Security Features](#)
- [Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch](#)
- [Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks](#)
- [Verifying That DAI Is Working Correctly](#)
- [Monitoring Port Security](#)
- [Understanding DAI for Port Security on page 4500](#)

## Enabling a Trusted DHCP Server (CLI Procedure)



**NOTE:** This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Enabling a Trusted DHCP Server \(CLI Procedure\)](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

You can configure any interface on a switch that connects to a DHCP server as a trusted interface (port). Configuring a DHCP server on a trusted interface protects against rogue DHCP servers sending leases.

By default, all access interfaces are untrusted, and all trunk interfaces are trusted. However, you can override the default setting for access interfaces by configuring a group of access interfaces within a VLAN, specifying an interface to belong to that group, and then configuring the group as trusted.

Before you can configure a trusted DHCP server, you must configure a VLAN. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 2337](#).

To configure an untrusted access interface as a trusted interface for a DHCP server by using the CLI :

1. Configure a group within a VLAN with a specific access interface:

```
[edit vlans vlan-name forwarding-options dhcp-security ]  
user@switch# set group group-name interface interface-name
```

2. Configure that group as **trusted** to make the specified interface contained within the group a trusted interface:

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name]  
user@switch# set overrides trusted
```

**Related  
Documentation**

- [Enabling a Trusted DHCP Server \(J-Web Procedure\) on page 4574](#)
- [Understanding Trusted DHCP Servers for Port Security on page 4527](#)

## Enabling a Trusted DHCP Server (J-Web Procedure)

You can configure any interface on the EX Series switch that connects to a DHCP server as a trusted interface (port). Configuring a DHCP server on a trusted interface protects against rogue DHCP servers sending leases.

You configure a trusted DHCP server on an interface, not on a VLAN. By default, all access interfaces are untrusted and all trunk interfaces are trusted.

To enable a trusted DHCP server on one or more interfaces by using the J-Web interface:

1. Select **Configure>Security>Port Security**.
2. Select one or more interfaces from the Port list.
3. Click the **Edit** button. If a message appears asking if you want to enable port security, click **Yes**.
4. Select the **Trust DHCP** check box and then click **OK**.
5. Click **OK** after the command has been successfully delivered.



**NOTE:** You can enable or disable port security on the switch at any time by clicking the **Activate** or **Deactivate** button on the Port Security Configuration page. If security status is shown as **Disabled** when you try to edit settings for any VLANs or interfaces (ports), the message asking if you want to enable port security appears.

---

**Related  
Documentation**

- [Enabling a Trusted DHCP Server \(CLI Procedure\)](#)
- [Example: Configuring Basic Port Security Features](#)
- [Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks](#)
- [Verifying That a Trusted DHCP Server Is Working Correctly](#)
- [Monitoring Port Security](#)
- [Understanding Trusted DHCP Servers for Port Security on page 4527](#)

## Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Performance (CLI Procedure)



**NOTE:** This task uses Junos OS for MX Series routers and EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Making IP-MAC Bindings in the DHCP Snooping Database Persistent (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 3.

By default, IP-MAC address bindings in the DHCP snooping database do not persist through device reboots. You can improve network performance by configuring the IP-MAC address bindings in the DHCP snooping database to persist through reboots so that the table does not need to be rebuilt after rebooting. Do this by configuring a storage location for the DHCP snooping database file, where you must specify how frequently the device writes the database entries into the DHCP snooping database file.



**NOTE:** You can also configure the IPv6-MAC address bindings to persist through reboots on devices that support DHCPv6 snooping.

DHCPv6 is not supported on the MX Series routers.

The DHCP snooping database of IP-MAC address bindings is created when you enable any of the port security features for a specific VLAN or bridge-domain in either hierarchy:

- [edit vlans *vlan-name* forwarding-options dhcp-security]
- [edit bridge-domains *bridge-domain-name* forwarding-options dhcp-security]

On devices that support DHCPv6, enabling any port security features also creates the DHCPv6 snooping database. DHCP snooping and DHCPv6 snooping are not enabled by default.

To configure a *local* storage location for the DHCP snooping database file:

- For DHCP snooping:

```
[edit system processes]
user@device# set dhcp-service dhcp-snooping-file local-pathname write-interval seconds
```

For example:

```
[edit system processes]
user@switch# set dhcp-service dhcp-snooping-file /var/tmp/test.log write-interval 60
```

- For DHCPv6 snooping (not supported on MX Series routers):

```
[edit system processes]
user@device# set dhcp-service dhcpv6-snooping-file local-pathname write-interval seconds
```

For example:

```
[edit system processes]
user@switch# set dhcp-service dhcpv6-snooping-file /var/tmp/test.log write-interval 60
```

To configure a *remote* storage location for IP-MAC bindings, use **tftp://ip-address** or **ftp://hostname/path** as the remote URL or the local pathname for the storage location of the DHCP or DHCPv6 snooping database file:

- For DHCP snooping:

```
[edit system processes]
user@switch# set dhcp-service dhcp-snooping-file remote_url write-interval seconds
```

For example:

```
[edit system processes]
user@switch# set dhcp-service dhcp-snooping-file tftp://test:Test123@14.1.2.1 write-interval 60
```

- For DHCPv6 snooping (not supported on MX Series routers):

```
[edit system processes]
user@switch# set dhcp-service dhcpv6-snooping-file remote_url write-interval seconds
```

For example:

```
[edit system processes]
user@switch# set dhcp-service dhcpv6-snooping-file tftp://test:Test123@14.1.2.1 write-interval 60
```



**NOTE:** Specify any requisite user credentials for the FTP server before you specify the IP address or hostname. In this example, **test** is the username and **Test123** is the password for FTP server 14.1.2.1.

---

**Related  
Documentation**

- *Understanding DHCP Snooping for Port Security*



## Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on [page 3](#).

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

You can configure the DHCP option 82 feature in two topologies:

- The switch, DHCP clients, and DHCP server are all on the same VLAN. The switch forwards the clients' requests to the server and forwards the server's responses to the clients. This topic describes this configuration.
- The switch functions as a relay agent when the DHCP clients or the DHCP server is connected to the switch through a Layer 3 interface. On the switch, these interfaces are configured as integrated routing and bridging (IRB) interfaces. The switch relays the clients' requests to the server and then forwards the server's responses to the clients. This configuration is described in *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*.

Before you configure DHCP option 82 on the switch, perform these tasks:

- Connect and configure the DHCP server.



**NOTE:** Your DHCP server must be configured to accept DHCP option 82. If the server is not configured for DHCP option 82, the server does not use the DHCP option 82 information in the requests sent to it when it formulates its reply messages.

- Configure a VLAN on the switch and associate the interfaces on which the clients and the server connect to the switch with that VLAN. See “[Configuring VLANs for EX Series Switches \(CLI Procedure\)](#)” on [page 2337](#)

To configure DHCP option 82:

1. Specify DHCP option 82 for the VLAN that you configured.

```
[edit vlans vlan-name forwarding-options dhcp-security]  
user@switch# set option-82
```



**NOTE:** If you want to enable DHCP option 82 on all VLANs, you must configure it separately for each specific VLAN.

The remaining steps are optional.

2. Configure the prefix for the circuit ID suboption to include the switch's hostname or the routing instance name for the VLAN:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set circuit-id prefix (host-name | routing-instance-name)
```

3. Specify that the circuit ID suboption value contains the interface description rather than the interface name (the default):

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set circuit-id use-interface-description
```

4. Specify that the circuit ID suboption value contains the VLAN ID rather than the VLAN name (the default):

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set circuit-id use-vlan-id
```

5. Specify that the remote ID suboption is included in the DHCP option 82 information:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set remote-id
```



**NOTE:** If you do not specify a keyword after *remote-id*, the default value for the *remote-id* suboption is the interface name.

6. Specify that the remote ID suboption is the hostname of the switch:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set remote-id host-name
```

7. Specify that the remote ID suboption value contains the interface description:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set remote-id use-interface-description
```

8. Specify that the remote ID suboption value contains a character string:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set remote-id use-string mystring
```

9. Configure a vendor ID suboption:

- To use the default value (the default value is **Juniper**), do not type a character string after the **vendor-id** option keyword:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set vendor-id
```

- To configure that the vendor ID suboption value contains a character string value that you specify rather than **Juniper** (the default):

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]
user@switch# set vendor-id use-string mystring
```

- Related Documentation**
- [Understanding DHCP Option 82 for Port Security on Switching Devices on page 4503](#)
  - [RFC 3046, DHCP Relay Agent Information Option](#), at <http://tools.ietf.org/html/rfc3046>.

## Enabling DHCPv6 Rapid Commit Support

You can configure the extended DHCPv6 local server to support the DHCPv6 Rapid Commit option (DHCPv6 option 14). When rapid commit is enabled on the extended DHCPv6 local server, the server recognizes the Rapid Commit option in Solicit messages sent from the DHCPv6 client. (DHCPv6 clients are configured separately to include the DHCPv6 Rapid Commit option in the Solicit messages.) The server and client then use a two-message exchange (Solicit and Reply) to configure clients, rather than the default four-method exchange (Solicit, Advertise, Request, and Reply). The two-message exchange provides faster client configuration, and is beneficial in environments in which networks are under a heavy load.

You can configure the DHCPv6 local server to support the Rapid Commit option globally, for a specific group, or for a specific interface. By default, rapid commit support is disabled on the DHCPv6 local server.

To configure the DHCPv6 local server to support the DHCPv6 Rapid Commit option:

1. Specify that you want to configure the **overrides** options.

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit overrides
```

2. Enable rapid commit support.

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set rapid-commit
```

- Related Documentation**
- [Overriding Default DHCP Local Server Configuration Settings on page 1423](#)
  - [Deleting DHCP Local Server and DHCP Relay Override Settings on page 1429](#)
  - [Extended DHCP Local Server Overview on page 1356](#)

## Configuration Statements

- [\[edit vlans\] Configuration Statement Hierarchy on EX Series Switches on page 4581](#)
- [arp-inspection on page 4585](#)
- [cak on page 4586](#)
- [circuit-id on page 4587](#)
- [ckn on page 4588](#)
- [connectivity-association on page 4589](#)
- [connectivity-association \(MACsec Interfaces\) on page 4590](#)

- [direction](#) on page 4591
- [dhcp-security](#) on page 4592
- [dhcp-service](#) on page 4594
- [dhcp-snooping-file](#) on page 4595
- [dhcpv6-options](#) on page 4596
- [dhcpv6-snooping-file](#) on page 4597
- [encryption](#) on page 4598
- [exclude-protocol](#) on page 4599
- [group \(DHCP Security\)](#) on page 4600
- [host-name](#) on page 4601
- [id](#) on page 4602
- [interface \(DHCP Security\)](#) on page 4603
- [interface-mac-limit](#) on page 4604
- [interfaces \(MACsec\)](#) on page 4606
- [ip-source-guard](#) on page 4607
- [ipv6-source-guard](#) on page 4609
- [key](#) on page 4610
- [key-server-priority](#) on page 4611
- [mac](#) on page 4612
- [mac-address \(MACsec\)](#) on page 4613
- [mac-move-limit](#) on page 4614
- [macsec](#) on page 4616
- [mka](#) on page 4617
- [must-secure](#) on page 4618
- [neighbor-discovery-inspection](#) on page 4619
- [no-dhcpv6-options](#) on page 4620
- [no-dhcpv6-snooping](#) on page 4620
- [no-encryption](#) on page 4621
- [no-option16](#) on page 4622
- [no-option18](#) on page 4622
- [no-option37](#) on page 4623
- [offset](#) on page 4624
- [option-16 \(DHCPv6 Snooping\)](#) on page 4625
- [option-18 \(DHCPv6 Snooping\)](#) on page 4626
- [option-37 \(DHCPv6 Snooping\)](#) on page 4627
- [overrides \(DHCP Security\)](#) on page 4628
- [packet-action](#) on page 4629

- [persistent-learning on page 4631](#)
- [port-id on page 4632](#)
- [pre-shared-key on page 4633](#)
- [prefix \(DHCPv6 Options\) on page 4634](#)
- [recovery-timeout on page 4635](#)
- [remote-id on page 4637](#)
- [replay-protect on page 4638](#)
- [replay-window-size on page 4639](#)
- [routing-instance-name \(circuit-id\) on page 4640](#)
- [secure-channel on page 4641](#)
- [security-association on page 4642](#)
- [security-mode on page 4643](#)
- [static-ip on page 4644](#)
- [static-ipv6 on page 4645](#)
- [traceoptions \(DHCP\) on page 4646](#)
- [transmit-interval \(MACsec\) on page 4648](#)
- [trusted on page 4649](#)
- [untrusted on page 4649](#)
- [use-interface-description on page 4650](#)
- [use-interface-index on page 4651](#)
- [use-interface-name on page 4652](#)
- [use-string on page 4653](#)
- [use-vlan-id on page 4654](#)
- [vendor-id on page 4655](#)
- [write-interval on page 4656](#)

## [edit vlans] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit vlans]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit vlans\] Hierarchy Level on page 4582](#)
- [Unsupported Statements in the \[edit vlans\] Hierarchy Level on page 4584](#)

### Supported Statements in the [edit vlans] Hierarchy Level

The following hierarchy shows the **[edit vlans]** configuration statements supported on one or more of the EX Series switches:

```
vlans {
  vlan-name {
    description text-description;
    domain-type bridge;
    forwarding-options {
      dhcp-security {
        arp-inspection;
        group group-name {
          interface interface-name {
            static-ip ip-address {
              mac mac-address;
            }
          }
        }
        overrides {
          no-option82;
          trusted;
        }
      }
    }
    ip-source-guard;
    no-dhcp-snooping;
    option-82 {
      circuit-id {
        prefix {
          host-name;
          logical-system-name;
          routing-instance-name;
        }
        use-interface-description (device | logical);
        use-vlan-id;
      }
      remote-id {
        host-name;
        use-interface-description (device | logical);
        use-string string;
      }
      vendor-id {
        use-string string;
      }
    }
  }
  filter {
    input filter-name;
    output filter-name;
  }
  flood {
```

```

        input filter-name;
    }
}
!3-interface irb.logical-unit-number;
multicast-snooping-options {
    flood-groups [group-names];
    forwarding-cache {
        threshold {
            reuse threshold;
            suppress threshold;
        }
    }
    graceful-restart {
        disable;
        restart-duration duration;
    }
    host-outbound-traffic {
        dot1p bits;
        forwarding-class forwarding-class;
    }
    multichassis-lag-replicate-state;
    nexthop-hold-time time;
    options {
        syslog {
            level level;
            mark interval;
            upto level;
        }
    }
}
traceoptions {
    file filename {
        files number;
        no-world-readable;
        size file-size;
        world-readable;
    }
    flag flag {
        disable;
    }
}
}
switch-options {
    interface interface-name {
        interface-mac-limit limit {
            packet-action action;
        }
        static-mac mac-address;
    }
    interface-mac-limit limit {
        packet-action action;
    }
    mac-move-limit limit {
        packet-action action;
    }
    mac-table-size limit {
        packet-action drop;
    }
}

```

```

    }
    no-mac-learning;
  }
  vlan-id number;
  vlan-id-list [vlan-id | vlan-id-vlan-id];
}

```

### Unsupported Statements in the [edit vlans] Hierarchy Level

All statements in the **[edit vlans]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 500: Unsupported [edit vlans] Configuration Statements on EX Series Switches**



| Statement                                                                                         | Hierarchy Level |
|---------------------------------------------------------------------------------------------------|-----------------|
| <b>NOTE:</b> Variables, such as <i>filename</i> , are not shown in the statements or hierarchies. |                 |
| mcae-mac-synchronize                                                                              | [edit vlans]    |
| no-irb-layer-2-copy                                                                               | [edit vlans]    |

**Related Documentation**

- [Example: Connecting Access Switches to a Distribution Switch on page 2291](#)



## arp-inspection

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <pre>arp-inspection {     forwarding-class <i>class-name</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <ul style="list-style-type: none"> <li>For platforms with ELS:           <ul style="list-style-type: none"> <li>[edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security</a>],</li> <li>[edit forwarding-options dhcp-relay ]</li> </ul> </li> <li>For platforms without ELS:           <ul style="list-style-type: none"> <li>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i>)],</li> <li>[edit forwarding-options dhcp-relay ]</li> </ul> </li> </ul> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security</a>] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p>                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <p>Perform dynamic ARP inspection (DAI) on all VLANs or on the specified VLAN.</p> <p>When DAI is enabled, the switch logs invalid ARP packets that it receives on each interface, along with the sender’s IP and MAC addresses. ARP probe packets, which have the sender IP address 0.0.0.0, are validated by DAI.</p>                                                                                                                                                                                   |
| <div>  <p><b>NOTE:</b> If you configure DAI at the [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security</a>] hierarchy level:</p> <ul style="list-style-type: none"> <li>DAI can be configured only for a specific VLAN, not for a list or a range of VLAN IDs.</li> <li>DHCP snooping is automatically enabled on the specified VLAN.</li> <li>The forwarding-class statement is not available at the [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security</a>] hierarchy level.</li> </ul> <p>See <a href="#">“Enabling Dynamic ARP Inspection (CLI Procedure)” on page 4572</a> for more information about this configuration.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <div>  <p><b>NOTE:</b> On EX9200 switches, DAI is not supported in an MC-LAG scenario.</p> </div> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch</i></li><li>• <i>Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks</i></li><li>• <a href="#">Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529</a></li><li>• <i>Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic</i></li><li>• <i>Enabling Dynamic ARP Inspection (CLI Procedure)</i></li><li>• <a href="#">Enabling Dynamic ARP Inspection (J-Web Procedure) on page 4572</a></li></ul> |

---

## cak

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ckn hexadecimal-number;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit security <a href="#">macsec connectivity-association</a> <a href="#">connectivity-association-name</a> <a href="#">pre-shared-key</a> ]                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Specifies the connectivity association key (CAK) for a pre-shared key.</p> <p>A pre-shared key includes a connectivity association key name (CKN) and a CAK. A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security keys. The MACsec Key Agreement (MKA) protocol is enabled once the pre-shared keys are successfully exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link</p> |
| <b>Default</b>                  | No CAK exists, by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>hexadecimal-number</b>—The key name, in hexadecimal format.</p> <p>The key name is 32 hexadecimal characters in length. If you enter a key name that is less than 32 characters long, the remaining characters are set to 0.</p>                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li></ul>                                                                                                                                                                                                                                                                                                                                                     |

## circuit-id

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> circuit-id {   prefix {     host-name;     logical-system-name;     routing-instance-name;   }   use-interface-description (device   logical);   use-vlan-id; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>For platforms with Enhanced Layer 2 Software (ELS):<br/>[edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> option-82 ]</li> <li>For platforms without ELS:<br/>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i>) dhcp-option82],<br/>[edit forwarding-options helpers bootp dhcp-option82] ,<br/>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]</li> <li>For MX Series platforms:<br/>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82]</li> </ul> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b>] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 14.1 for the MX Series.</p>                                                                                                                                           |
| <b>Description</b>              | <p>Configure the <b>circuit-id</b> suboption (suboption 1) of DHCP option 82 (the DHCP relay agent information option) in DHCP packets destined for a DHCP server. This suboption identifies the circuit (the interface, the VLAN, or both) on which the DHCP request arrived.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                |
| <b>Default</b>                  | <p>If DHCP option 82 is enabled on the switch, the circuit ID is supplied by default in the format <i>interface-name:vlan-name</i> or, on a Layer 3 interface, just <i>interface-name</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Setting Up DHCP Option 82 on an MX Series Router (CLI Procedure)</i></li> <li><i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i></li> <li><i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i></li> </ul>                                                                                                                                                                                                                                                        |

- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4577*
- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>

---

## ckn

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ckn hexadecimal-number;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit security <a href="#">macsec connectivity-association</a> <a href="#">connectivity-association-name</a> <a href="#">pre-shared-key</a> ]                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Specifies the connectivity association key name (CKN) for a pre-shared key.</p> <p>A pre-shared key includes a CKN and a connectivity association key (CAK). A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security keys. The MACsec Key Agreement (MKA) protocol is enabled once the pre-shared keys are successfully exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link</p> |
| <b>Default</b>                  | No CKN exists, by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b><i>hexadecimal-number</i></b>—The key name, in hexadecimal format.</p> <p>The key name is 32 hexadecimal characters in length. If you enter a key name that is less than 32 characters long, the remaining characters are set to 0.</p>                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li></ul>                                                                                                                                                                                                                                                                                                                                                     |

## connectivity-association

```
Syntax connectivity-association connectivity-association-name {
    exclude-protocol protocol-name;
    include-sci;
    mka {
        must-secure;
        key-server-priority priority-number;
        transmit-interval interval;
    }
    no-encryption;
    offset (0|30|50);
    pre-shared-key {
        cak hexadecimal-number;
        ckn hexadecimal-number;
    }
    replay-protect{
        replay-window-size number-of-packets;
    }
    secure-channel secure-channel-name {
        direction (inbound | outbound);
        encryption;
        id {
            mac-address mac-address;
            port-id port-id-number;
        }
        offset (0|30|50);
        security-association security-association-number {
            key key-string;
        }
    }
    security-mode security-mode;
}
```

**Hierarchy Level** [edit security *macsec*]

**Release Information** Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.  
Statement introduced in Junos OS Release 14.1X53-D15 for the QFX Series.

**Description** Create or configure a MACsec connectivity association.

A connectivity association is not applying MACsec to traffic until it is associated with an interface. MACsec connectivity associations are associated with interfaces using the *interfaces* statement in the [edit security macsec] hierarchy.

**Default** No connectivity associations are present, by default.

**Options** The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation** • [Configuring Media Access Control Security \(MACsec\) on page 4548](#)

---

## connectivity-association (MACsec Interfaces)

---

|                                 |                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | connectivity-association <i>connectivity-association-name</i> ;                                                                                               |
| <b>Hierarchy Level</b>          | [edit security <a href="#">macsec interfaces</a> <i>interface-name</i> ]                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches. |
| <b>Description</b>              | Applies a connectivity association to an interface, which enables Media Access Control Security (MACsec) on that interface.                                   |
| <b>Default</b>                  | No connectivity associations are associated with any interfaces.                                                                                              |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                               |
| <b>Related Documentation</b>    | • <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a>                                                                             |

## direction

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>direction (inbound   outbound);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit security <code>macsec connectivity-association connectivity-association-name secure-channel secure-channel-name</code> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Configure whether the secure channel applies MACsec security to traffic entering or leaving an interface.</p> <p>If you need to apply MACsec on traffic entering and leaving an interface, you need to create one secure channel to apply MACsec on incoming traffic and another secure channel to apply MACsec on outgoing traffic within the same connectivity association. When you associate the connectivity association with an interface, MACsec is applied on traffic entering and leaving that interface.</p> <p>You only use this configuration option when you are configuring MACsec using static secure association keys (SAK) security mode. When you are configuring MACsec using static connectivity association keys (CAK) security mode, two secure channels that are not user-configurable—one inbound secure channel and one outbound secure channel—are automatically created within the connectivity association.</p> |
| <b>Default</b>                  | <p>This statement does not have a default value.</p> <p>If you have configured a secure channel to enable MACsec using static SAK security mode, you must specify whether the secure channel applies MACsec to traffic entering or leaving an interface. A candidate configuration that contains a secure channel that has not configured a direction cannot be committed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>inbound</b>—Enable MACsec security on traffic entering the interface that has applied the secure channel.</p> <p><b>outbound</b>—Enable MACsec security on traffic leaving the interface that has applied the secure channel.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## dhcp-security

---

**Syntax** `dhcp-security {  
 arp-inspection;  
 group group-name {  
 interface interface-name {  
 static-ip ip-address {  
 mac mac-address;  
 }  
 }  
 }  
 overrides {  
 no-option82;  
 trusted;  
 untrusted;  
 }  
}  
ip-source-guard;  
neighbor-discovery-inspection;  
no-dhcp-snooping;  
option-82 {  
 circuit-id {  
 prefix {  
 host-name;  
 logical-system-name;  
 routing-instance-name;  
 }  
 use-interface-description (device | logical);  
 use-vlan-id;  
 }  
 remote-id {  
 host-name hostname;  
 use-interface-description (device | logical);  
 mac;  
 use-string string;  
 }  
 vendor-id {  
 use-string string;  
 }  
}`

**Hierarchy Level** [edit vlans *vlan-name* forwarding-options]

**Release Information** Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. Support for **static-ipv6**, **nd-inspection**, **ipv6-source-guard**, **no-dhcpv6-snooping**, and **no-option-37** introduced in Junos OS Release 13.2X51-D20 for EX Series switches.

**Description** Configure port security features on the switch. DHCP snooping is enabled automatically if you configure any of the following port security features within this hierarchy:

- Dynamic ARP inspection (DAI)
- IP source guard



- DHCP option 82
- Static IP

For switches that support DHCPv6, both DHCP snooping and DHCPv6 snooping are enabled automatically if you configure any of the features listed above or any of the following IPv6 features:

- IPv6 Neighbor Discovery inspection
- IPv6 source guard
- Static IPv6



**NOTE:** On EX9200 switches, DHCP snooping, DAI, and IP source guard are not supported in an MC-LAG scenario.

The remaining statements are explained separately.


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling Dynamic ARP Inspection (CLI Procedure) on page 4572</a></li> <li>• <a href="#">Configuring IP Source Guard (CLI Procedure) on page 4546</a></li> <li>• <a href="#">Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4577</a></li> <li>• <a href="#">Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports (CLI Procedure) on page 4571</a></li> </ul> |

## dhcp-service

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>dhcp-service {<br/>    <a href="#">dhcp-snooping-file</a> (<i>local_pathname</i>   <i>remote_URL</i>);<br/>    <a href="#">write-interval</a> <i>interval</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit system processes]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1 for the MX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Enable DHCP services on the device. DHCP services automate network-parameter assignment to network devices. The DHCP service process is enabled by default. However, by default, IP-MAC bindings in the DHCP snooping database do not persist through device reboots. You can improve performance after rebooting by configuring the IP-MAC bindings to persist, by configuring a storage location for the DHCP database file. When specifying the location for the DHCP database, you must also specify how frequently the switch writes the database entries into the DHCP snooping database file.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Performance (CLI Procedure)</a> on page 4575</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## dhcp-snooping-file

|                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                 | <pre>dhcp-snooping-file {   (local_pathname   remote_URL);   write-interval seconds; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                        | [edit system processes <a href="#">dhcp-service</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                    | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1 for the MX Series.                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                            | <p>Ensure that IP-MAC bindings persist through the device reboots by specifying a local pathname or a remote URL for the storage location of the DHCP snooping database file. You <i>must</i> specify how frequently the device writes the database entries into the DHCP snooping database file.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                           |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                | <p>The IP-MAC bindings in the DHCP snooping database file are not persistent by default. If the device is rebooted, the bindings are lost, and the table must be rebuilt on reboot.</p> <p>You can set either a local or remote storage location:</p> <ul style="list-style-type: none"> <li>To configure a <i>local</i> storage location for the DHCP snooping database file, use the variable <i>local_pathname</i>.</li> <li>To configure a <i>remote</i> storage location, use <b>tftp://ip-address</b> or <b>ftp://hostname/path</b> as the remote URL.</li> </ul> |
| <div>  <p><b>NOTE:</b> Specify any requisite user credentials for the FTP server before you specify the IP address or hostname. (See example in: “<a href="#">Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Performance (CLI Procedure)</a>” on page 4575)</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                               | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li><a href="#">Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Performance (CLI Procedure)</a> on page 4575</li> <li><a href="#">Understanding DHCP Snooping for Port Security</a></li> </ul>                                                                                                                                                                                                                                                                                                    |

## dhcpv6-options

```
Syntax  dhcpv6-options {
        option-16 {
            use-string string;
        }
        option-18 {
            prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
                vlan-id;
                vlan-name;
            }
            use-interface-mac;
            use-interface-index (device | logical);
            use-interface-description (device | logical);
            use-interface-name (device | logical);
            use-string string;
        }
        option-37 {
            prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
                vlan-id;
                vlan-name;
            }
            use-interface-mac;
            use-interface-index (device | logical);
            use-interface-description (device | logical);
            use-interface-name (device | logical);
            use-string string;
        }
    }
```

**Hierarchy Level** [edit vlans *vlan-name* forwarding-options [dhcp-security](#)]

**Release Information** Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

**Description** Configure optional information to be included in DHCPv6 packets during the DHCPv6 snooping process.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [no-dhcpv6-options on page 4620](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4577](#)

- [Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports \(CLI Procedure\) on page 4571](#)

## dhcpv6-snooping-file

|                                 |                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> dhcpv6-snooping-file (<i>local_pathname</i>   <i>remote_URL</i>);   location <i>local_pathname</i>   <i>remote_URL</i>;   timeout <i>seconds</i>;   write-interval <i>seconds</i>; } </pre>                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit system processes <a href="#">dhcp-service</a>];</p> <p>[edit ethernet-switching-options secure-access-port]</p>                                                                                                                                                 |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches.</p> <p>Support at the <b>[edit ethernet-switching-options secure-access-port]</b> hierarchy level introduced in Junos OS 14.1X53-D10 for EX Series switches.</p>                         |
| <b>Description</b>              | <p>Ensure that IP-MAC address bindings persist through switch reboots by specifying a local pathname or a remote URL for the storage location of the DHCPv6 snooping database file.</p> <p>The remaining statements are explained separately.</p>                        |
| <b>Default</b>                  | <p>The IPv6-MAC address bindings in the DHCPv6 snooping database are not persistent. If the switch is rebooted, the bindings are lost.</p>                                                                                                                               |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Performance (CLI Procedure) on page 4575</a></li> <li>• <a href="#">Understanding DHCP Snooping for Port Security</a></li> </ul> |

## encryption

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | encryption;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit security <b>macsec</b> <b>connectivity-association</b> <i>connectivity-association-name</i> <b>secure-channel</b> <i>secure-channel-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Enable MACsec encryption within a secure channel.</p> <p>You can enable MACsec without enabling encryption. If a connectivity association with a secure channel that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic has not been tampered with and does not represent a security threat.</p> <p>Traffic traversing a MAC-enabled point-to-point Ethernet link traverses the link at the same speed regardless of whether encryption is enabled or disabled. You cannot increase the speed of traffic traversing a MACsec-enabled Ethernet link by disabling encryption.</p> <p>This command is used to enable encryption when MACsec is configured using secure association key (SAK) security mode only. When MACsec is configuring using static connectivity association key (CAK) security mode, the encryption setting is configured outside of the secure channel using the <b>no-encryption</b> configuration statement.</p> |
| <b>Default</b>                  | MACsec encryption is disabled when MACsec is configured using static SAK security mode, by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## exclude-protocol

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>exclude-protocol <i>protocol-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit security <a href="#">macsec connectivity-association</a> <i>connectivity-association-name</i> ]                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Specifies protocols whose packets are not secured using Media Access Control Security (MACsec) when MACsec is enabled on a link using static connectivity association key (CAK) security mode.</p> <p>When this option is enabled in a connectivity association that is attached to an interface, MACsec is not enabled for all packets of the specified protocols that are sent and received on the link.</p> |
| <b>Default</b>                  | <p>Disabled.</p> <p>All packets are secured on a link when MACsec is enabled, with the exception of all types of Spanning Tree Protocol (STP) packets.</p>                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b><i>protocol-name</i></b>—Specifies the name of the protocol that should not be MACsec-secured. Options include:</p> <ul style="list-style-type: none"> <li>• <b>cdp</b>—Cisco Discovery Protocol.</li> <li>• <b>lcp</b>—Link Aggregation Control Protocol.</li> <li>• <b>lldp</b>—Link Level Discovery Protocol.</li> </ul>                                                                                 |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li> </ul>                                                                                                                                                                                                                                                                               |

## group (DHCP Security)

---

**Syntax**    `group group-name {  
          interface interface-name {  
            static-ip ip-address {  
              mac mac-address;  
            }  
            static-ipv6 ip-address {  
              mac mac-address;  
            }  
          }  
          overrides {  
            no-option37;  
            no-option82;  
            trusted;  
            untrusted;  
          }  
        }`

**Hierarchy Level**    [edit vlans *vlan-name* forwarding-options [dhcp-security](#)]

**Release Information**    Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.  
Statement introduced in Junos OS Release 13.2 for the QFX series.  
Support for the **static-ipv6** and **no-option37** statements introduced in Junos OS Release 13.2X51-D20 for EX Series switches.

**Description**    Specify the name of a group of access interfaces that you want to configure for DHCP security attributes that are different from the attributes set for other interfaces in the VLAN. A group must contain at least one interface.

The remaining statements are explained separately.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports \(CLI Procedure\) on page 4571](#)
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 4573](#)
- [Understanding DHCP Snooping for Port Security](#)



## host-name

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | host-name <i>host-name</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level (EX Series)</b> | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> option-82 <b>remote-id</b> ]                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level (MX Series)</b> | [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82 <b>remote-id</b> ]                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>         | Statement introduced in Junos OS Release 13.2X50-D10.<br>Statement introduced in Junos OS Release 14.1 for the MX Series.                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                 | Use the hostname of the switching device as the <b>remote-id</b> suboption of DHCP option 82 (also known as the DHCP relay agent information option) in DHCP request packet headers before forwarding or relaying requests to a DHCP server. This suboption provides a trusted identifier for the host system that has forwarded or relayed requests to the server.                                                                              |
| <b>Required Privilege Level</b>    | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li>• <a href="#">Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4577</a></li> <li>• <a href="#">Setting Up DHCP Option 82 on an MX Series Router (CLI Procedure)</a></li> <li>• RFC 3046, <i>DHCP Relay Agent Information Option</i>, at <a href="http://tools.ietf.org/html/rfc3046">http://tools.ietf.org/html/rfc3046</a></li> </ul> |

## id

---

|                                 |                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>id {<br/>    <code>mac-address</code> <i>mac-address</i>;<br/>    <code>port-id</code> <i>port-id-number</i>;<br/>}</code>                                                       |
| <b>Hierarchy Level</b>          | [edit security <code>macsec</code> <code>connectivity-association</code> <i>connectivity-association-name</i> <code>secure-channel</code> <i>secure-channel-name</i> ]                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                          |
| <b>Description</b>              | Specify a MAC address and a port that traffic on the link must be from to be accepted by the interface when MACsec is enabled using static secure association key (SAK) security mode. |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                     |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li></ul>                                                      |

## interface (DHCP Security)

|                                 |                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interface <i>interface-name</i> {   <b>static-ip</b> <i>ip-address</i> {     <b>mac</b> <i>mac-address</i>;   }   static-ipv6 <i>ip-address</i> {     <b>mac</b> <i>mac-address</i>;   } } </pre>                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security group</b> <i>group-name</i> ]                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Support for the <b>static-ipv6</b> statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches.</p>                                                                    |
| <b>Description</b>              | <p>Configure an interface for a static IPv4 or IPv6 address to MAC address binding (IP-MAC binding) or configure an interface to belong to a group within the VLAN that has DHCP security attributes that are different from the attributes of other interfaces in the VLAN.</p> <p>The remaining statement is explained separately.</p>               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports (CLI Procedure) on page 4571</a></li> <li>• <a href="#">Enabling a Trusted DHCP Server (CLI Procedure) on page 4573</a></li> <li>• <a href="#">Configuring Port Security (CLI Procedure) on page 4539</a></li> </ul> |

## interface-mac-limit

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>interface-mac-limit <i>limit</i> {<br/>    <b>packet-action</b> drop;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>     | <p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> switch-options],<br/>[edit logical-systems <i>logical-system-name</i> switch-options interface <i>interface-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> switch-options],<br/>[edit routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>],<br/>[edit switch-options],<br/>[edit switch-options interface <i>interface-name</i>],<br/>[edit switch-options interface <i>interface-name</i>],<br/>[edit vlans <i>vlan-name</i> switch-options],<br/>[edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | (MX Series routers or EX Series switches only) Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



**NOTE:** For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the `interface-mac-limit` statement or changing the `interface-mac-limit` configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the `interface-mac-limit` statement or use the `commit at configuration` statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers by using the `clear bridge mac-table` command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

**Default** For an access port, the default MAC limit is 1024 MAC addresses. For a trunk port, the default MAC limit is 8192 MAC addresses.

**Options** *limit*—Maximum number of MAC addresses learned from an interface.

**Range:** 1 through 131,071 MAC addresses per interface

The remaining statement is explained separately.

**Required Privilege** routing—To view this statement in the configuration.

**Level** routing-control—To add this statement to the configuration.

**Related  
Documentation**

- *Layer 2 Learning and Forwarding for Bridge Domains Overview*
- *Layer 2 Learning and Forwarding for VLANs Overview*
- *Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

## interfaces (MACsec)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interfaces <i>interface-name</i> {<br/>    connectivity-association <i>connectivity-association-name</i>;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit security <a href="#">macsec</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Applies the specified connectivity association to the specified interface to enable MACsec.</p> <p>One connectivity association can be applied to multiple interfaces.</p> <p>You must always use this statement to apply a connectivity association to an interface to enable MACsec. You must complete this configuration step regardless of whether MACsec is enabled using static connectivity association key (CAK) security mode or static secure association key (SAK) security mode.</p> <p>If you are enabling MACsec using static SAK security mode and need to configure MACsec on inbound and outbound traffic on the same interface, you must configure a connectivity association with one secure channel for inbound traffic and a second secure channel for outbound traffic. The connectivity association is then applied to the interface using this statement to enable MACsec for traffic entering and leaving the interface.</p> |
| <b>Default</b>                  | Interfaces are not associated with any connectivity associations, by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## ip-source-guard

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>ip-source-guard;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | <ul style="list-style-type: none"> <li>For platforms with ELS:<br/>[edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security</a>]</li> <li>For platforms without ELS:<br/>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i>)]</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security</a>] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>         | <p>Perform IP source guard checking on packets sent from access interfaces. Validate source IP addresses and source MAC addresses on all VLANs or on the specified VLAN or VLAN range. Forward packets with valid addresses and drop those with invalid addresses.</p> <ul style="list-style-type: none"> <li><b>ip-source-guard</b>—Enable IP source guard checking.</li> <li><b>no-ip-source-guard</b>—(Not available in [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security</a>]) Disable IP source guard checking.</li> </ul> <p>If you configure IP source guard at the [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security</a>] hierarchy level:</p> <ul style="list-style-type: none"> <li>IP source guard can be configured only for a specific VLAN, not for a list or a range of VLAN IDs.</li> <li>DHCP snooping is automatically enabled.</li> </ul> <p>See <a href="#">“Configuring IP Source Guard (CLI Procedure)” on page 4546</a> for more information about this configuration.</p> <p>If you configure IP source guard at the [edit ethernet-switching-options secure-access-port <i>vlan</i> (all   <i>vlan-name</i>)] hierarchy level:</p> <ul style="list-style-type: none"> <li>You must enable DHCP snooping on all VLANs if you configure IP source guard on all VLANs.</li> <li>You must enable DHCP snooping for the specific VLAN if you configure IP source guard on that specific VLAN. Otherwise, the default behavior of no DHCP snooping applies to that VLAN.</li> </ul> <p>See <a href="#">Enabling DHCP Snooping (CLI Procedure)</a> for more information about this configuration.</p> |



**NOTE:** On EX9200 switches, IP source guard is not supported in an MC-LAG scenario.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | Disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN</i></li><li>• <i>Example: Configuring IP Source Guard with Other EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces</i></li><li>• <a href="#">Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529</a></li><li>• <a href="#">Configuring IP Source Guard (CLI Procedure)</a></li><li>• <a href="#">Configuring IP Source Guard (CLI Procedure) on page 4546</a></li></ul> |



## ipv6-source-guard

|                            |                                                                                                                                                                                                                                                                       |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | ipv6-source-guard;                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> ];<br>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i> )]                                                                                                        |
| <b>Release Information</b> | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches.<br>Support at the [edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i> )] hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches. |
| <b>Description</b>         | Perform IPv6 source guard checking on packets sent from access interfaces. Validate source IP addresses and source MAC addresses on all VLANs or on the specified VLAN. Forward packets with valid addresses and drop those with invalid addresses.                   |



**NOTE:** If you configure the `ipv6-source-guard` statement at the [edit vlans *vlan-name* forwarding-options **dhcp-security**] hierarchy level, DHCPv6 snooping is automatically enabled for the specified VLAN.

If you configure the `ipv6-source-guard` statement at the [edit ethernet-switching-options secure-access-port vlan *vlan-name*] hierarchy level, you must also enable DHCPv6 snooping for the specified VLAN.

|                                 |                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | Disabled.                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect the Switch from IPv6 Address Spoofing on page 4534</a></li> <li>• <a href="#">Configuring IP Source Guard (CLI Procedure) on page 4546</a></li> </ul> |

## key

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>key key-string;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit security <a href="#">macsec connectivity-association connectivity-association-name secure-channel secure-channel-name security-association security-association-number</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Specifies the static security key to exchange to enable MACsec using static secure association key (SAK) security mode.</p> <p>The key string is a 32-digit hexadecimal number. The key string and the security association must match on both sides of an Ethernet connection to secure traffic using MACsec when enabling MACsec using SAK security mode.</p> <p>You must configure at least two security associations with unique security association numbers and key strings to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.</p> |
| <b>Default</b>                  | This statement does not have a default value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <b>key-string</b> —Specifies the key to exchange with the other end of the link on the secure channel. The <i>key-string</i> is a 32-digit hexadecimal string that is created by the user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## key-server-priority

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>key-server-priority <i>priority-number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit security <code>macsec connectivity-association</code> <code>connectivity-association-name mka</code> ]                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Specifies the key server priority used by the MACsec Key Agreement (MKA) protocol to select the key server when MACsec is enabled using static connectivity association key (CAK) security mode.</p> <p>The switch with the lower <i>priority-number</i> is selected as the key server.</p> <p>If the <i>priority-number</i> is identical on both sides of a point-to-point link, the MKA protocol selects the device with the lower MAC address as the key server.</p> |
| <b>Default</b>                  | The default key server priority number is 16.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><i>priority-number</i>—Specifies the MKA server election priority number.</p> <p>The <i>priority-number</i> can be any number between 0 and 255. The lower the number, the higher the priority.</p>                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li> </ul>                                                                                                                                                                                                                                                                                                                                        |

## mac

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mac mac-address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>For platforms with Enhanced Layer 2 Software (ELS):<br/>[edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> group <i>group-name</i> interface <i>interface-name</i> <b>static-ip</b> <i>ip-address</i>]</li><li>For platforms without ELS:<br/>[edit ethernet-switching-options secure-access-port interface (all   <i>interface-name</i>) <b>static-ip</b> <i>ip-address</i> vlan <i>vlan-name</i>]</li><li>For MX Series platforms:<br/>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security group <i>group-name</i> interface <i>interface-name</i> static-ip <i>ip-address</i>]</li></ul> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b>] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)</p> <p>Hierarchy level [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 14.1 for the MX Series.</p>                                                                                                                                                                                 |
| <b>Description</b>              | Configure the media access control (MAC) address or hardware address of the device connected to the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <i>mac-address</i> —Value (in hexadecimal format) of the address assigned to this device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><a href="#">Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure)</a></li><li><a href="#">Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports (CLI Procedure) on page 4571</a></li><li><a href="#">Configuring Static IP Addresses for DHCP Bindings on Access Ports for MX Series Routers (CLI Procedure)</a></li></ul>                                                                                                                                                                                                                                                             |

## mac-address (MACsec)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mac-address <i>mac-address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit security <code>macsec connectivity-association</code> <i>connectivity-association-name</i> <code>secure-channel</code> <i>secure-channel-name</i> <code>id</code> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Specify a MAC address to enable MACsec using static secure association key (SAK) security mode. The <b>mac-address</b> variables must match on the sending and receiving ends of a link to enable MACsec using static SAK security mode.</p> <p>If you are configuring a MAC address on a secure channel in the outbound direction, you should specify the MAC address of the interface as the <b>mac-address</b>.</p> <p>If you are configuring a MAC address on a secure channel in the inbound direction, you should specify the MAC address of the interface at the other end of the link as the <b>mac-address</b>.</p> <p>You only use this configuration option when you are configuring MACsec using static SAK security mode. This option does not need to be specified when you are enabling MACsec using static connectivity association key (CAK) security mode.</p> |
| <b>Default</b>                  | No MAC address is specified in the secure channel, by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>mac-address</b> —The MAC address, in six groups of two hexadecimal digits.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## mac-move-limit

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>mac-move-limit {<br/>    limit;<br/>    &lt;action action   packet-action action&gt;;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | <ul style="list-style-type: none"><li>• For platforms with ELS:<br/>[edit vlans <i>vlan-name</i> switch-options]</li><li>• For platforms without ELS:<br/>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i>)]</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Hierarchy level [edit vlans <i>vlan-name</i> switch-options] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>         | Specify the number of times a MAC address can move to a new interface (port) in one second and the action to be taken by the switch if the MAC address move limit is exceeded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>             | If you do not specify <b>mac-move-limit</b> , the default MAC address move limit is unlimited.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>             | <p><b>limit</b> <i>limit</i>—Maximum number of moves to a new interface per second.</p> <ul style="list-style-type: none"><li>• <b>action</b> <i>action</i>—(Optional) (Available <i>only</i> under the hierarchy level [edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i>) mac-move-limit]) Action to take when the MAC address move limit is reached:<ul style="list-style-type: none"><li>• <b>drop</b>—Drop the packet and generate a system log entry. This is the default.</li><li>• <b>log</b>—Do not drop the packet but generate a system log entry.</li><li>• <b>none</b>—No action.</li><li>• <b>shutdown</b>—Disable the interface and generate a system log entry. If you have configured the switch with the <b>port-error-disable</b> statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the <b>clear ethernet-switching port-error</b> command.</li></ul></li><li>• <b>packet-action</b> <i>action</i>—(Optional) (Available <i>only</i> under the hierarchy level, [edit vlans <i>vlan-name</i> switch-options mac-move-limit]) Action to take when the MAC address move limit is reached:</li></ul> |



**NOTE:** There is no default action.

---

- **drop**—Drop the packet and do not generate an alarm.

- **drop and log**—Drop the packet and generate an alarm, an SNMP trap, or system log entry.
- **log**— Do not drop the packet, but generate an alarm, an SNMP trap, or a system log entry.
- **none**—No action.
- **shutdown**—Disable the interface and generate an alarm or an SNMP trap. If you have configured the interface with the [recovery-timeout](#) statement, the disabled interface recovers automatically upon expiration of the specified timeout. If you have not configured the interface for a recovery timeout, you can bring up the disabled interface by running the operational command [clear ethernet-switching recovery-timeout](#).

**Required Privilege Level** system—To view this statement in the configuration.  
system—control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Basic Port Security Features*
- *Configuring MAC Move Limiting (CLI Procedure)*
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 4568 \(ELS\)](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 4569](#)
- *Configuring MAC Move Limiting (J-Web Procedure)*
- *Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*
- [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\) on page 2206](#)

## macsec

```
Syntax  macsec {
        connectivity-association connectivity-association-name {
            exclude-protocol protocol-name;
            include-sci;
            mka {
                must-secure;
                key-server-priority priority-number;
                transmit-interval interval;
            }
            no-encryption;
            offset (0|30|50);
            pre-shared-key {
                cak hexadecimal-number;
                ckn hexadecimal-number;
            }
            replay-protect {
                replay-window-size number-of-packets;
            }
            secure-channel secure-channel-name {
                direction (inbound | outbound);
                encryption;
                id {
                    mac-address mac-address;
                    port-id port-id-number;
                }
                offset (0|30|50);
                security-association security-association-number {
                    key key-string;
                }
            }
            security-mode security-mode;
        }
        interfaces interface-name {
            connectivity-association connectivity-association-name;
        }
    }
```

**Hierarchy Level** [edit security]

**Release Information** Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.  
Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.

**Description** Configure Media Access Control Security (MACsec)..

**Options** The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Media Access Control Security \(MACsec\) on page 4548](#)



---

## mka

---


|                                 |                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>mka {<br/>    must-secure;<br/>    key-server-priority <i>priority-number</i>;<br/>    transmit-interval <i>interval</i>;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit security <a href="#">macsec connectivity-association</a> <i>connectivity-association-name</i> ]                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15.<br>Statement introduced in Junos OS Release 14.1X53-D15 for the QFX Series.          |
| <b>Description</b>              | Specify parameters for the MACsec Key Agreement (MKA) protocol.                                                                            |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                         |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li></ul>          |

## must-secure

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | must-secure;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit security <b>macsec</b> <b>connectivity-association</b> <i>connectivity-association-name</i> <b>mka</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.<br>Statement introduced in Junos OS Release 14.1X53-D15 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Specifies that all traffic travelling on the MACsec-secured link must be MACsec-secured to be forwarded onward.</p> <p>When the <b>must-secure</b> option is enabled, all traffic that is not MACsec-secured that is received on the interface is dropped.</p> <p>When the <b>must-secure</b> option is disabled, all traffic from devices that support MACsec is MACsec-secured while traffic received from devices that do not support MACsec is forwarded through the network.</p> <p>The <b>must-secure</b> option is particularly useful in scenarios where multiple devices, such as a phone and a PC, are accessing the network through the same Ethernet interface. If one of the devices supports MACsec while the other device does not support MACsec, the device that doesn't support MACsec can continue to send and receive traffic over the network—provided the <b>must-secure</b> option is disabled—while traffic to and from the device that supports MACsec is MACsec-secured. In this scenario, traffic to the device that is not MACsec-secured must be VLAN-tagged.</p> |
| <b>Default</b>                  | The <b>must-secure</b> option is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## neighbor-discovery-inspection

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | neighbor-discovery-inspection;                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> ];<br>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i> )]                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches.<br>Support at the [edit <b>ethernet-switching-options secure-access-port vlan</b> (all   <i>vlan-name</i> )] hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches.                                                                                                                                          |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <p>Perform dynamic IPv6 neighbor discovery inspection on the specified VLAN.</p> <p>When neighbor discovery inspection is configured, the switch inspects IPv6 packets with neighbor discovery messages and validates them against the DHCPv6 binding table. The source IP address and source MAC address of each packet are checked against the table, and if a valid match is not found, the packet is dropped.</p> |
| <div>  <p><b>NOTE:</b> If you configure the <b>neighbor-discovery-inspection</b> statement at the [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b>] hierarchy level, DHCPv6 snooping is automatically enabled for the specified VLAN.</p> <p>See “<a href="#">Enabling IPv6 Neighbor Discovery Inspection</a>” on page 4547 for more information about this configuration.</p> <p>If you configure the <b>neighbor-discovery-inspection</b> statement at the [edit <b>ethernet-switching-options secure-access-port vlan</b> (all   <i>vlan-name</i>)] hierarchy level, you must also enable DHCPv6 snooping for the specified VLAN or VLANs.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Disabled.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Port Security (CLI Procedure) on page 4539</a></li> <li>• <a href="#">Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect the Switch from IPv6 Address Spoofing on page 4534</a></li> <li>• <a href="#">Enabling IPv6 Neighbor Discovery Inspection on page 4547</a></li> </ul>                                      |

## no-dhcpv6-options

---

|                                 |                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-dhcpv6-options;                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security group group-name overrides</a> ]                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure a specific group of one or more access interfaces within the VLAN not to add any DHCPv6 options, even if the VLAN is configured to perform DHCPv6 snooping. DHCPv6 options include option 16, option 18, and option 37.                                                             |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">dhcpv6-options on page 4596</a></li><li>• <a href="#">Understanding DHCP Snooping for Port Security on page 4507</a></li><li>• <a href="#">Understanding DHCP Option 82 for Port Security on Switching Devices on page 4503</a></li></ul> |

## no-dhcpv6-snooping

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-dhcpv6-snooping;                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security</a> ]                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches.                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Disable DHCPv6 snooping for the specified VLAN.                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>                  | DHCPv6 snooping is not enabled by default.<br><br>There is no configuration statement that explicitly enables DHCPv6 snooping. DHCPv6 snooping is enabled automatically by Junos OS if any port security feature, such as IPv6 Neighbor Discovery inspection or IPv6 source guard, is configured at the <b>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security]</b> hierarchy level. |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding DHCP Snooping for Port Security</a></li></ul>                                                                                                                                                                                                                                                                                |

## no-encryption

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-encryption;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit security <b>macsec</b> <b>connectivity-association</b> <i>connectivity-association-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Disables MACsec encryption for a connectivity association that is configured to enable MACsec using static connectivity association key (CAK) or dynamic security mode.</p> <p>You can enable MACsec without enabling encryption. If a connectivity association that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the packet, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic does not represent a security threat.</p> <p>This command is used to disable encryption when MACsec is configured using static CAK or dynamic security mode only. When MACsec is configuring using static secure association key (SAK) security mode, the encryption setting is managed in the secure channel using the <b>encryption</b> configuration statement.</p> |
| <b>Default</b>                  | MACsec encryption is enabled if MACsec is enabled using static CAK or dynamic security mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## no-option16

---

|                                 |                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-option16;                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security group</b> <i>group-name overrides</i> ]                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.                                                                                                                                                                                                             |
| <b>Description</b>              | Configure a specific group of one or more access interfaces within the VLAN not to transmit DHCPv6 option 16 information, even if the VLAN is configured to perform DHCPv6 snooping. Option 16 information that has already been added by a DHCPv6 client will be forwarded as is.       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">option-16 on page 4625</a></li><li>• <a href="#">Understanding DHCP Snooping for Port Security on page 4507</a></li><li>• <a href="#">Understanding DHCP Option 82 for Port Security on Switching Devices on page 4503</a></li></ul> |

## no-option18

---

|                                 |                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-option18;                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security group</b> <i>group-name overrides</i> ]                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.                                                                                                                                                                                                             |
| <b>Description</b>              | Configure a specific group of one or more access interfaces within the VLAN <i>not</i> to transmit DHCP option 18 information, even if the VLAN is configured to perform DHCPv6 snooping.                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">option-18 on page 4626</a></li><li>• <a href="#">Understanding DHCP Option 82 for Port Security on Switching Devices on page 4503</a></li><li>• <a href="#">Understanding DHCP Snooping for Port Security on page 4507</a></li></ul> |

---

## no-option37

---

|                                 |                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-option37;                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security group group-name overrides</a> ]                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches.                                                                                                                                                                                   |
| <b>Description</b>              | Configure a specific group of one or more access interfaces within the VLAN <i>not</i> to transmit DHCP option 37 information, even if the VLAN is configured to perform dhcpv6-snooping.                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">option-82</a></li><li>• <a href="#">Understanding DHCP Option 82 for Port Security on Switching Devices on page 4503</a></li><li>• <a href="#">Understanding DHCP Snooping for Port Security</a></li></ul> |

## offset

---

**Syntax**    offset (0 | 30 | 50);

**Hierarchy Level**    [edit security **macsec connectivity-association** *connectivity-association-name*]  
[edit security **macsec connectivity-association** *connectivity-association-name* **secure-channel** *secure-channel-name*]

**Release Information**    Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.  
Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.

**Description**    Specifies the number of octets in an Ethernet frame that are sent in unencrypted plain-text when encryption is enabled for MACsec.

Setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the remaining traffic. Setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the remaining traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

You configure the **offset** in the [edit security **macsec connectivity-association** *connectivity-association-name*] hierarchy when you are enabling MACsec using static connectivity association key (CAK) or dynamic security mode.

You configure the **offset** in the [edit security **macsec connectivity-association** *connectivity-association-name* **secure-channel** *secure-channel-name*] hierarchy when you are enabling MACsec using static secure association key (SAK) security mode.

**Default**    0

**Options**    **0**—Specifies that no octets are unencrypted. When you set the offset to 0, all traffic on the interface where the connectivity association or secure channel is applied is encrypted.

**30**—Specifies that the first 30 octets of each Ethernet frame are unencrypted.



**NOTE:** In IPv4 traffic, setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the rest of the traffic. An offset of 30, therefore, is typically used when a feature needs this information to perform a task on IPv4 traffic.

---

**50**—Specified that the first 50 octets of each Ethernet frame are unencrypted.





**NOTE:** In IPv6 traffic, setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the rest of the traffic. An offset of 50, therefore, is typically used when a feature needs this information to perform a task on IPv6 traffic.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Media Access Control Security \(MACsec\) on page 4548](#)

## option-16 (DHCPv6 Snooping)

**Syntax**

```
option-16 {
    use-string string;
}
```

**Hierarchy Level** [edit vlans *vlan-name* forwarding-options [dhcp-security dhcpv6-options](#)]

**Release Information** Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

**Description** Configure the DHCPv6 Vendor ID option (option 16) to be included in DHCPv6 requests from clients before forwarding them to a DHCPv6 server. Option 16 provides the server with information about the vendor that manufactured the hardware on which the DHCP client is running. When configured, the switch will overwrite any existing option 16 information sent by clients in the DHCPv6 packets.

Option 16 is the DHCPv6 equivalent of the [vendor-id](#) sub-option of DHCP option 82.

**Options** **use-string *string***—Define a custom string to be used as the DHCPv6 vendor identifier.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 4572](#)
- [Configuring IP Source Guard \(CLI Procedure\) on page 4546](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4577](#)
- [Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports \(CLI Procedure\) on page 4571](#)

## option-18 (DHCPv6 Snooping)

**Syntax**

```
option-18 {
  prefix {
    host-name;
    logical-system-name;
    routing-instance-name;
    vlan-id;
    vlan-name;
  }
  use-interface-index (device | logical);
  use-interface-description (device | logical);
  use-interface-mac;
  use-interface-name (device | logical);
  use-string string;
}
```

**Hierarchy Level** [edit vlans *vlan-name* forwarding-options [dhcp-security dhcpv6-options](#)]

**Release Information** Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

**Description** Configure the DHCPv6 Relay Agent Interface ID option (option 18) to insert information in DHCPv6 requests from clients before forwarding them to a DHCPv6 server. Option 18 provides information about the port on which the request was received, which the server can use to assign IP addresses, prefixes, and other configuration parameters for the client.

When option 18 is configured, a unique interface ID is inserted into the DHCPv6 packet headers. Suboptions can be configured to include a prefix with the interface ID or to change the type of information used to identify the interface. The default fields included in option 18 are the Juniper Enterprise ID, VLAN, and MAC address of the interface.

Option 18 is the DHCPv6 equivalent of the [circuit-id](#) sub-option of DHCP option 82.



**NOTE:** DHCPv6 packets that already contain option 18 information when received from a client are dropped by the switch.

**Options**

**use-interface-mac**—Use the MAC address of the interface in the DHCPv6 interface ID.

**use-string *string***—Use a custom string in the DHCPv6 interface ID.

The remaining statements are explained separately.

**Required Privilege Level**

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

**Related Documentation**

- [no-option18 on page 4622](#)
- [no-dhcpv6-options on page 4620](#)

## option-37 (DHCPv6 Snooping)

```
Syntax  option-37 {
        prefix {
            host-name;
            logical-system-name;
            routing-instance-name;
            vlan-id;
            vlan-name;
        }
        use-interface-index (device | logical);
        use-interface-description (device | logical);
        use-interface-mac;
        use-interface-name (device | logical);
        use-string string;
    }
```

**Hierarchy Level** [edit vlans *vlan-name* forwarding-options [dhcp-security dhcpv6-options](#)]

**Release Information** Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

**Description** Configure the DHCPv6 Relay Agent Remote ID option (option 37) to insert information in DHCPv6 requests from clients before forwarding them to a DHCPv6 server. Option 37 provides information about the remote host, which the server can use to assign IP addresses, prefixes, and other configuration parameters for the client.

When option 37 is configured, a unique remote ID is inserted into the DHCPv6 packet headers. Suboptions can be configured to include a prefix with the remote ID or to change the interface portion of the ID. The default fields included in option 37 are the Juniper Enterprise ID, VLAN, and MAC address of the interface.

Option 37 is the DHCPv6 equivalent of the [remote-id](#) sub-option of DHCP option 82.



**NOTE:** DHCPv6 packets that already contain option 37 information when received from a client are dropped by the switch.

**Options** **use-interface-mac**—Use the MAC address of the interface in the DHCPv6 remote ID.  
**use-string *string***—Use a custom string in the DHCPv6 remote ID.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [no-option37 on page 4623](#)
- [no-dhcpv6-options on page 4620](#)

- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4577](#)
- [Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports \(CLI Procedure\) on page 4571](#)

## overrides (DHCP Security)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>overrides (trusted   untrusted [no-option37   no-option-82]);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> group <i>group-name</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Support for the <b>no-option37</b> option introduced in Junos OS Release 13.2X51-D20 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Modify selected attributes of a specific interface within a group of interfaces that is configured within a specified VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>no-option37</b>—The interface specified in this group does not support DHCPv6 option 37.</p> <p><b>no-option82</b>—The interface specified in this group does not support DHCP option 82.</p> <p><b>trusted</b>—The interface specified in this group is trusted. DHCP snooping does not apply to the trusted interface. Likewise, DAI and IP source guard—even if they are enabled for the VLAN—do not apply to the interface that is configured with the <b>overrides</b> and the <b>trusted</b> options. Access interfaces are untrusted by default.</p> <p><b>untrusted</b>—(Only for EX9200) The interface specified in this group is untrusted. Trunk interface are trusted by default. Access interfaces are untrusted by default.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling a Trusted DHCP Server (CLI Procedure) on page 4573</a></li><li>• <a href="#">Understanding DHCP Snooping for Port Security</a></li><li>• <a href="#">Understanding DHCP Option 82 for Port Security on Switching Devices on page 4503</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## packet-action

**Syntax** `packet-action action;`

**Hierarchy Level** [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* switch-options **interface-mac-limit** *limit*],  
 [edit protocols l2-learning global-mac-limit *limit*],  
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* protocols evpn interface-mac-limit (vpls)],  
 [edit routing-instances *routing-instance-name* protocols evpn interface *interface-name* interface-mac-limit (vpls)],  
 [edit routing-instances *routing-instance-name* protocols evpn mac-table-size *limit*],  
 [edit routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],  
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit switch-options **interface-mac-limit** *limit*],  
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit switch-options **interface-mac-limit** *limit*],  
 [edit switch-options **mac-table-size** *limit*],  
 [edit **switch-options on page 2212** interface *interface-name* **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options **mac-table-size** *limit*]  
 [edit **vlans on page 2370** *vlan-name* switch-options **interface-mac-limit** *limit*],  
 [edit **vlans on page 2370** *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit **vlans on page 2370** *vlan-name* switch-options **mac-table-size** *limit*]

**Release Information** Statement introduced in Junos OS Release 8.4.  
 Support for the **switch-options** statement added in Junos OS Release 9.2.  
 Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy

supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

[edit switch-options interface *interface-name* interface-mac-limit *limit*], [edit switch-options interface-mac-limit *limit*], [edit switch-options mac-table-size *limit*], [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*], [edit vlans *vlan-name* switch-options interface-mac-limit *limit*], and [edit vlans *vlan-name* switch-options mac-table-size *limit*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Support for EVPNs introduced in Junos OS Release 13.2 on MX Series 3D Universal Edge Routers.

Support at the [edit switch-options interface *interface-name* interface-mac-limit *limit*] hierarchy level and hierarchy levels under [edit vlans *vlan-name*] introduced in Junos OS Release 13.2X50-D10 for EX Series switches and Junos OS Release 13.2 for the QFX Series.

**Description** Specify the action taken when packets with new source MAC addresses are received after the MAC address limit is reached. If this statement is not configured, packets with new source MAC addresses are forwarded by default.

**Default**



**NOTE:** On a QFX Series Virtual Chassis, if you include the shutdown option at the [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level and issue the commit operation, the system generates a commit error. The system does not generate an error if you include the shutdown option at the [edit switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level.

Disabled. The default is for packets for new source MAC addresses to be forwarded after the MAC address limit is reached.

- Options**
- drop**—Drop packets with new source MAC addresses, and do not learn the new source MAC addresses.
  - drop-and-log**—(EX Series switches and QFX Series only) Drop packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.
  - log**—(EX Series switches and QFX Series only) Hold packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.
  - none**—(EX Series switches and QFX Series only) Forward packets with new source MAC addresses, and learn the new source MAC address.
  - shutdown**—(EX Series switches and QFX Series only) Disable the specified interface, and generate an alarm, an SNMP trap, or a system log entry.

**Required Privilege Level**

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring EVPN Routing Instances*
  - [Configuring MAC Limiting \(CLI Procedure\) on page 2361](#)
  - [Configuring Persistent MAC Learning \(CLI Procedure\) on page 4569](#)
  - *Layer 2 Learning and Forwarding for Bridge Domains Overview*
  - *Layer 2 Learning and Forwarding for VLANs Overview*
  - *Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
  - *Layer 2 Learning and Forwarding for VLANs Overview*
  - *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

## persistent-learning

|                                 |                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>persistent-learning;</code>                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <a href="#">switch-options on page 2212</a> interface <i>interface-name</i> ]                                                                                                                      |
| <b>Release Information</b>      | Hierarchy level [edit switch-options interface interface-name] introduced in Junos OS Release 13.2X50-D10                                                                                                |
| <b>Description</b>              | Specify that learned MAC addresses persist on the specified interfaces across restarts of the switch and link-down conditions. This feature is also known as sticky MAC.                                 |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system—control—To add this statement to the configuration.                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Basic Port Security Features</i></li> <li>• <a href="#">Configuring Persistent MAC Learning (CLI Procedure) on page 4569</a></li> </ul> |

## port-id

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>port-id port-id-number;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit security <a href="#">macsec connectivity-association</a> <i>connectivity-association-name</i> <a href="#">secure-channel</a> <i>secure-channel-name</i> <b>id</b> ]                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Specify a port ID in a secure channel when enabling MACsec using static secure association key (SAK) security mode. The port IDs must match on a sending and receiving secure channel on each side of a link to enable MACsec.</p> <p>Once the port numbers match, MACsec is enabled for all traffic on the connection.</p> <p>You only use this configuration option when you are configuring MACsec using static SAK security mode. This option does not need to be specified when you are enabling MACsec using static connectivity association key (CAK) security mode.</p> |
| <b>Default</b>                  | No port ID is specified.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>port-id-number</i> —The port ID number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |



## pre-shared-key

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>pre-shared-key {     cak hexadecimal-number;     ckn hexadecimal-number; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit security <b>macsec connectivity-association</b> <i>connectivity-association-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Specifies the pre-shared key used to enable MACsec using static connectivity association key (CAK) security mode.</p> <p>A pre-shared key includes a connectivity association key name (CKN) and a connectivity association key (CAK). A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using static CAK security mode. The MACsec Key Agreement (MKA) protocol is enabled after the pre-shared keys are successfully verified and exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link.</p> |
| <b>Default</b>                  | No pre-shared keys exist, by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## prefix (DHCPv6 Options)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>prefix {<br/>    host-name;<br/>    logical-system-name;<br/>    routing-instance-name;<br/>    vlan-id;<br/>    vlan-name;<br/>}</pre>                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit vlans forwarding-options <a href="#">dhcp-security dhcpv6-options option-18</a> ]<br>[edit vlans forwarding-options <a href="#">dhcp-security dhcpv6-options option-37</a> ]                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure a prefix for DHCPv6 option 18 (Interface ID) or option 37 (Remote ID). When configured, the prefix is inserted into DHCPv6 packets during the DHCPv6 snooping process.                                                                                                                                                                                                       |
| <b>Default</b>                  | If the <b>prefix</b> statement is not explicitly specified, no prefix is inserted in DHCPv6 packets.                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>host-name</b>—Add the host name of the switch to DHCPv6 options.</p> <p><b>logical-system-name</b>—Add the logical system name to the DHCPv6 options.</p> <p><b>routing-instance-name</b>—Add the routing instance name to the DHCPv6 options.</p> <p><b>vlan-id</b>—Add the VLAN ID to the DHCPv6 options.</p> <p><b>vlan-name</b>—Add the VLAN name to the DHCPv6 options.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">option-37 (DHCPv6 Snooping) on page 4627</a></li><li>• <a href="#">option-18 (DHCPv6 Snooping) on page 4626</a></li></ul>                                                                                                                                                                                                          |

## recovery-timeout

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                     | <code>recovery-timeout seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level (EX Series and QFX Series)</b> | [edit interfaces <i>interface-name</i> unit 0 family ethernet-switching]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level (MX Series)</b>                | [edit interfaces <i>interface-name</i> unit 0 family bridge]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                        | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1 for the MX Series routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                                | <p>Disable rather than block an interface when enforcing MAC limiting, MAC move limiting, or rate-limiting configuration options for shutting down the interface, and allow the interface to recover automatically from the error condition after the specified period of time:</p> <ul style="list-style-type: none"> <li>• If you have enabled MAC limiting with the <b>shutdown</b> option and you enable <b>recovery-timeout</b>, the switch disables (rather than shuts down) the interface when the MAC address limit is reached.</li> <li>• If you have enabled MAC move limiting (not supported on EX9200) with the <b>shutdown</b> option and you enable <b>recovery-timeout</b>, the switch disables (rather than shuts down) the interface when the maximum number of moves to a new interface is reached.</li> <li>• If you have enabled storm control with the <b>action-shutdown</b> option and you enable <b>recovery-timeout</b>, the switch disables (rather than shuts down) the interface when applicable traffic exceeds the specified levels. Depending upon the configuration, applicable traffic could include broadcast, unknown unicast, and multicast traffic.</li> </ul> |
|                                                   | <p> <b>NOTE:</b> The <b>recovery-timeout</b> configuration does not apply to pre-existing error conditions. It impacts only error conditions that are detected after the <b>recovery-timeout</b> statement has been enabled and committed. To clear a pre-existing error condition and restore the interface to service, use the operational mode command <b>clear ethernet-switching recovery-timeout</b> for EX Series and QFX Series and <b>clear bridge recovery-timeout</b> for MX Series routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                                    | Not enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                    | <p><b>seconds</b>— Number of seconds that the interface remains in a disabled state due to a port error prior to automatic recovery.</p> <p><b>Range:</b> 10 through 3600</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                   | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

**Related  
Documentation**

- [action-shutdown on page 2215](#)
- [Configuring MAC Limiting \(CLI Procedure\) on page 2361](#)
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 4568](#)
- [Configuring or Disabling Storm Control \(CLI Procedure\) on page 2207](#)

## remote-id

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>remote-id {   host-name <i>host-name</i>;   mac;   prefix ( hostname   mac   none );   use-interface-description ( logical   device );   use-string <i>string</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>     | <ul style="list-style-type: none"> <li>For platforms with Enhanced Level 2 Software (ELS):<br/>[edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security option-82</b>]</li> <li>For platforms without ELS:<br/>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i>) dhcp-option82],<br/>[edit forwarding-options helpers bootp dhcp-option82],<br/>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br/>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br/>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security option-82</b>]<br/>introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>         | <p>Insert the <b>remote-id</b> suboption of DHCP option 82 (also known as the DHCP relay agent information option) in DHCP request packet headers before forwarding or relaying requests to a DHCP server. This suboption provides a trusted identifier for the host system that has forwarded or relayed requests to the server.</p> <p>The remaining statements are explained separately, and their availability depends on the hierarchy level at which the <b>remote-id</b> suboption is specified, as follows:</p> <ul style="list-style-type: none"> <li>The statement <b>prefix</b>, is <i>not</i> supported at the [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security option-82</b>] hierarchy level.</li> <li>The statement <b>host-name</b> is supported <i>only</i> at the [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security option-82</b>] hierarchy level.</li> </ul> |
| <b>Default</b>             | <p>If the <b>remote-id</b> statement is not explicitly set, no remote ID value is inserted in the DHCP request packet header.</p> <p>If the <b>remote-id</b> statement is explicitly set, but is not qualified by a keyword, the following are true:</p> <ul style="list-style-type: none"> <li>At the [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b>] hierarchy level, the default keyword value is <i>interface-name</i>.</li> <li>At all other hierarchy levels, the default value of the <b>remote-id</b> keyword is the MAC address of the switch.</li> </ul>                                                                                                                                                                                                                                                                                                                        |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i></li><li>• <i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i></li><li>• <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li><li>• <a href="#">Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4577</a></li><li>• <i>Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li><li>• RFC 3046, <i>DHCP Relay Agent Information Option</i>, at <a href="http://tools.ietf.org/html/rfc3046">http://tools.ietf.org/html/rfc3046</a></li></ul> |

---

## replay-protect

---

|                                 |                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>replay-protect {<br/>    replay-window-size number-of-packets;<br/>}</pre>                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit security <b>macsec</b> <b>connectivity-association</b> <i>connectivity-association-name</i> ]                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                    |
| <b>Description</b>              | Enable replay protection for MACsec.<br><br>A replay window size specified using the <b>replay-window-size</b> <i>number-of-packets</i> statement must be specified to enable replay protection. |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                               |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li></ul>                                                                |

## replay-window-size

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>replay-window-size <i>number-of-packets</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit security <a href="#">macsec connectivity-association</a> <i>connectivity-association-name</i> replay-protect]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Specifies the size of the replay protection window.</p> <p>This statement has to be configured to enable replay protection.</p> <p>When MACsec is enabled on an Ethernet link, an ID number is assigned to each packet entering the link. The ID number of the packet is checked by the receiving interface after the packet has traversed the MACsec-enabled link.</p> <p>When replay protection is enabled, the sequence of the ID number of received packets are checked. If the packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.</p> <p>Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.</p> <p>Replay protection should not be enabled in cases where packets are expected to arrive out of order.</p> |
| <b>Default</b>                  | Replay protection is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><i>number-of-packets</i>—Specifies the size of the replay protection window, in packets.</p> <p>When this variable is set to 0, all packets that arrive out-of-order are dropped.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## routing-instance-name (circuit-id)

---

|                                 |                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | routing-instance--name;                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> option-82 <b>circuit-id</b> prefix]                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for EX Series switches.                                                                                                                                                                                                                                 |
| <b>Description</b>              | Specify that the routing instance name used by the VLAN is included with the circuit ID suboption in the DHCP option 82 information that is inserted by the switch into the packet header of a DHCP request before it forwards or relays the request to a DHCP server                                 |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</a> on page 4577</li><li>• <a href="#">Understanding DHCP Option 82 for Port Security on Switching Devices</a> on page 4503</li></ul> |



## secure-channel

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>secure-channel <i>secure-channel-name</i> {   <i>direction</i> (inbound   outbound);   encryption;   id {     <i>mac-address</i> <i>mac-address</i>;     <i>port-id</i> <i>port-id-number</i>;   }   <i>offset</i> (0 30 50);   <i>security-association</i> <i>security-association-number</i> {     <i>key</i> <i>key-string</i>;   } }</pre>                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit security <i>macsec</i> <i>connectivity-association</i> <i>connectivity-association-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Create and configure a secure channel to enable and configure MACsec when MACsec is enabled using static secure association key (SAK) security mode.</p> <p>You do not need to use this option to enable MACsec using static connectivity association key (CAK) security mode. All configuration for MACsec using static CAK security mode is done inside of the connectivity association but outside of the secure channel. When MACsec is enabled using static CAK security mode, an inbound and an outbound secure channel—neither of which is user-configurable—is automatically created within the connectivity association.</p> |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## security-association

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>security-association <i>security-association-number</i> {<br/>    <i>key</i> <i>key-string</i>;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit security <i>macsec</i> connectivity-association connectivity-association-name <i>secure-channel</i><br/>    <i>secure-channel-name</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Specifies the number of one of the security associations in the secure channel when MACsec is enabled using static secure association key (SAK) security mode. Because SAKs are created by the key server when MACsec is enabled using static connectivity association key (CAK) security mode, the <b>security-association</b> statement is not used when enabling MACsec using static CAK security mode.</p> <p>You must configure at least two security associations to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.</p> |
| <b>Default</b>                  | No security keys are configured, by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b><i>security-association-number</i></b>—Specifies the security association number and creates the SAK.</p> <p>The security association number is a whole number between 0 and 3. You can configure two security associations in a secure channel when enabling MACsec using static security keys.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## security-mode

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>security-mode <i>security-mode</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit security <a href="#">macsec connectivity-association</a> <i>connectivity-association-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15.<br>The <b>dynamic</b> security mode option was introduced in Junos OS Release 14.1X53-D10.<br>Statement introduced in Junos OS Release 14.1X53-D15 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure the MACsec security mode for the connectivity association.<br><br>We recommend enabling MACsec on switch-to-switch Ethernet links using static connectivity association key (CAK) security mode. Static CAK security mode ensures security by frequently refreshing to a new random secure association key (SAK) and by only sharing the SAK between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available when you enable MACsec using static CAK security mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b><i>security-mode</i></b>—Specifies the MACsec security mode. Options include:</p> <ul style="list-style-type: none"> <li>• <b>dynamic</b>—Dynamic mode.<br/><br/>Dynamic security mode is used to enable MACsec on switch-to-host Ethernet links. In dynamic mode, a master key is retrieved from a RADIUS server by a switch and a host as part of the AAA handshake in separate transactions. The MKA protocol is enabled when the master key is exchanged between the switch and the host.</li> <li>• <b>static-cak</b>—Static connectivity association key (CAK) mode.<br/><br/>Static CAK security mode is used to enable MACsec on switch-to-switch Ethernet links. In <b>static-cak</b> mode, the switch at one end of the point-to-point link acts as the key server and regularly transmits a randomized key using a process that does not transmit any traffic outside of the MACsec-secured point-to-point link.</li> <li>• <b>static-sak</b>—Static secure association key (SAK) mode.<br/><br/>Static SAK security mode is used to enable MACsec on switch-to-switch Ethernet links. In <b>static-sak</b> mode, one of two user-configured security keys is used to secure the point-to-point link. The two security keys are regularly rotated.</li> </ul> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## static-ip

---

|                                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                           | <pre>static-ip <i>ip-addresses</i> {<br/>    vlan <i>vlan-name</i>;<br/>    mac <i>mac-address</i>;<br/>}</pre>                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                  | <ul style="list-style-type: none"><li>For platforms with ELS:<br/>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> interface <i>interface-name</i>]</li><li>For platforms without ELS:<br/>[edit ethernet-switching-options secure-access-port interface (all   <i>interface-name</i>)]</li></ul> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                              | Statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)                         |
| <b>Description</b>                                                                                                                                                                                                                                                                                                      | Configure a static IP address to MAC address (IP-MAC) binding to be added to the DHCP snooping database.                                                                                                                                                                                                                               |
| <div> <b>NOTE:</b> The VLAN is specified at the higher hierarchy level when static-ip is configured at [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> interface <i>interface-name</i>].</div> |                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                                                                                                                                                                                                                                                                                          | <p><i>ip-address</i>—Static IP address assigned to a device connected on the specified interface.</p> <p><i>mac mac-address</i>—Static MAC address assigned to a device connected on the specified interface.</p> <p>The remaining statements are explained separately.</p>                                                            |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                         | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                      |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                            | <ul style="list-style-type: none"><li><a href="#">Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure)</a></li><li><a href="#">Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports (CLI Procedure) on page 4571</a></li></ul>                                                    |

## static-ipv6

|                                 |                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>static-ipv6 <i>ip-address</i> {<br/>    <b>mac</b> <i>mac-address</i>;<br/>}</code>                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> group <i>group-name</i> <b>interface</b> <i>interface-name</i> ];<br>[edit ethernet-switching-options secure-access-port interface <i>interface-name</i> ]                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches.<br>Support at the [edit <b>ethernet-switching-options</b> <b>secure-access-port</b> <b>interface</b> <i>interface-name</i> ] hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches.  |
| <b>Description</b>              | Configure a static IP-MAC binding to be added to the DHCPv6 snooping database.                                                                                                                                                                                                                |
| <b>Options</b>                  | <b><i>ip-address</i></b> —Static IPv6 address assigned to a device connected on the specified interface.<br><br><b><b>mac</b> <i>mac-address</i></b> —Static MAC address assigned to a device connected on the specified interface.<br><br>The remaining statements are explained separately. |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports (CLI Procedure) on page 4571</a></li> </ul>                                                                                                                 |

## traceoptions (DHCP)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i> &gt; &lt;size <i>maximum-file-size</i> &gt;<br/>    &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i>;<br/>    level (all   error   info   notice   verbose   warning);<br/>    no-remote-trace;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | [edit system <a href="#">processes</a> dhcp-service]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | <p>Define global tracing operations for extended DHCP local server and extended DHCP relay agent processes.</p> <p>Replaces deprecated <b>traceoptions</b> statements at the <b>[edit forwarding-options dhcp-relay]</b> and <b>[edit system services dhcp-local-server]</b> hierarchy levels.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements:</p> <ul style="list-style-type: none"><li>• <b>all</b>—Trace all events.</li><li>• <b>auth</b>—Trace authentication events.</li><li>• <b>database</b>—Trace database events.</li><li>• <b>fwd</b>—Trace firewall process events.</li><li>• <b>general</b>—Trace miscellaneous events.</li><li>• <b>ha</b>—Trace high availability-related events.</li><li>• <b>interface</b>—Trace interface operations.</li><li>• <b>io</b>—Trace I/O operations.</li><li>• <b>liveness-detection</b>—Trace liveness detection operations.</li><li>• <b>packet</b>—Trace packet and option decoding operations.</li><li>• <b>performance</b>—Trace performance measurement operations.</li><li>• <b>profile</b>—Trace profile operations.</li></ul> |

- **rpd**—Trace routing protocol process events.
- **rtsock**—Trace routing socket operations.
- **security-persistence**—Trace security persistence events.
- **session-db**—Trace session database events.
- **state**—Trace changes in state.
- **statistics**—Trace baseline statistics.
- **ui**—Trace user interface operations.

**level**—Level of tracing to perform; also known as severity level. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- **all**—Match messages of all levels.
- **error**—Match error messages.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.
- **warning**—Match warning messages.

**Default:** error

**match *regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

**no-remote-trace**—Disable remote tracing.

**no-world-readable**—(Optional) Disable unrestricted file access, allowing only the user **root** and users who have the Junos OS **maintenance** permission to access the trace files.

**size *maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (***maximum-file-sizek***), megabytes (***maximum-file-sizem***), or gigabytes (***maximum-file-sizeg***). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Range:** 10,240 through 1,073,741,824

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

|                           |                                                           |
|---------------------------|-----------------------------------------------------------|
| <b>Required Privilege</b> | trace—To view this statement in the configuration.        |
| <b>Level</b>              | trace-control—To add this statement to the configuration. |

**Related Documentation**    • [Tracing Extended DHCP Operations](#)

---

## transmit-interval (MACsec)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>transmit-interval <i>interval</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit security <a href="#">macsec connectivity-association connectivity-association-name mka</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Specifies the transmit interval for MACsec Key Agreement (MKA) protocol data units (PDUs).</p> <p>The MKA transmit interval setting sets the frequency for how often the MKA PDU is sent to the directly connected device to maintain MACsec on a point-to-point Ethernet link. A lower <i>interval</i> increases bandwidth overhead on the link; a higher <i>interval</i> optimizes the MKA protocol data unit exchange process.</p> <p>The transmit interval settings must be identical on both ends of the link when MACsec using static connectivity association key (CAK) security mode is enabled.</p> <p>We recommend increasing the interval to 6000 ms in high-traffic load environments.</p> |
| <b>Default</b>                  | The default transmit interval is 2000 milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>interval</i> —Specifies the transmit interval, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | • <a href="#">Configuring Media Access Control Security (MACsec) on page 4548</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



## trusted

---

|                                 |                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | trusted;                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security group</b> <i>group-name</i> <b>overrides</b> ]                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.                                                                                                                                                     |
| <b>Description</b>              | Allow DHCP responses from the specified interface. The interface is not subject to DHCP snooping, even if the VLAN is enabled for DHCP snooping.                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling a Trusted DHCP Server (CLI Procedure) on page 4573</a></li> <li>• <a href="#">Understanding Trusted DHCP Servers for Port Security on page 4527</a></li> <li>• <a href="#">Understanding DHCP Snooping for Port Security</a></li> </ul> |

## untrusted

---

|                                 |                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | untrusted;                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security group</b> <i>group-name</i> <b>overrides</b> ]                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.                                                                                                                                                            |
| <b>Description</b>              | Override the default behavior of a trunk interface from trusted to untrusted.                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling a Trusted DHCP Server (CLI Procedure) on page 4573</a></li> <li>• <a href="#">Understanding Trusted DHCP Servers for Port Security on page 4527</a></li> <li>• <a href="#">Understanding DHCP Snooping for Port Security</a></li> </ul> |

## use-interface-description

|                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                             | use-interface-description (device   logical);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>For Platforms with Enhanced Layer 2 Software (ELS)</b> | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> option-82 <b>circuit-id</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>For Platforms Without ELS</b>                          | [edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i> ) dhcp-option82 <b>circuit-id</b> ],<br>[edit forwarding-options helpers bootp dhcp-option82 <b>circuit-id</b> ],<br>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82 <b>circuit-id</b> ],<br>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i> ) dhcp-option82 <b>remote-id</b> ],<br>[edit forwarding-options helpers bootp dhcp-option82 <b>remote-id</b> ],<br>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82 <b>remote-id</b> ] |
| <b>For MX Series Platforms</b>                            | [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82 <b>circuit-id</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>                                | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b>] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)</p> <p>Hierarchy level [edit bridge-domains <i>bridge domain name</i> forwarding-options <b>dhcp-security</b>] introduced in Junos OS Release 14.1 for the MX Series.</p>                                                |
| <b>Description</b>                                        | Use the interface description rather than the interface name (which is the default value) in the circuit ID or remote ID value in the DHCP option 82 information.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                            | <p><b>device</b>—Use the device interface description. Only available for MX Series platform configuration.</p> <p><b>logical</b>—Use the logical interface description. Only available for MX Series platform configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b>                           | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>                              | <ul style="list-style-type: none"> <li>• <i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i></li> <li>• <i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i></li> <li>• <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li> <li>• <i>Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li> </ul>                                                                             |

- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 4577](#)
- [Setting Up DHCP Option 82 on an MX Series Router \(CLI Procedure\)](#)
- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>

## use-interface-index

|                                 |                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>use-interface-index (logical   device);</code>                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security dhcpv6-options <a href="#">option-18</a> ],<br>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security dhcpv6-options <a href="#">option-37</a> ],                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.                                                                                                                                                                                                                           |
| <b>Description</b>              | Use the index number of the interface instead of the interface name in the DHCPv6 option 18 (Relay Agent Interface-ID) or option 37 (Relay Agent Remote-ID). These options are used by a relay agent to insert information in DHCPv6 requests before the relay agent forwards them to a DHCPv6 server. |
| <b>Options</b>                  | <b>logical</b> —Use the textual description that is configured for the logical interface.<br><br><b>device</b> —Use the textual description that is configured for the device interface.                                                                                                               |
| <b>Required Privilege Level</b> | <b>interface</b> —To view this statement in the configuration.<br><b>interface-control</b> —To add this statement to the configuration.                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Including a Textual Description in DHCP Options on page 1468</a></li> <li>• <a href="#">Using DHCP Relay Agent Option 82 Information on page 1464</a></li> <li>• <a href="#">Configuring DHCPv6 Relay Agent Options</a></li> </ul>                |

## use-interface-name

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | use-interface-name (logical   device);                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security dhcpv6-options <a href="#">option-18</a> ],<br>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security dhcpv6-options <a href="#">option-37</a> ],                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.                                                                                                                                                                                                        |
| <b>Description</b>              | Configure DHCPv6 option 18 (Relay Agent Interface-ID) or option 37 (Relay Agent Remote-ID) to use the interface name to identify the port identity of the DHCP client to the DHCP server.                                                                                           |
| <b>Options</b>                  | <b>logical</b> —Use the name that is configured for the logical interface.<br><b>device</b> —Use the name that is configured for the device interface.                                                                                                                              |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Including a Textual Description in DHCP Options on page 1468</a></li><li>• <a href="#">Using DHCP Relay Agent Option 82 Information on page 1464</a></li><li>• <a href="#">Configuring DHCPv6 Relay Agent Options</a></li></ul> |

## use-string

|                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                             | <code>use-string <i>string</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>For Platforms with Enhanced Layer 2 Software (ELS)</b> | <code>[edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> option-82 <i>remote-id</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>For Platforms Without ELS</b>                          | <code>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i>) dhcp-option82 <i>remote-id</i>],</code><br><code>[edit forwarding-options helpers bootp dhcp-option82 <i>remote-id</i>],</code><br><code>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82 <i>remote-id</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>For MX Series Platforms</b>                            | <code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82 <i>circuit-id</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                                | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Hierarchy level <code>[edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b>]</code> introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)<br>Hierarchy level <code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options <b>dhcp-security</b>]</code> introduced in Junos OS Release 14.1 for the MX Series.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                                        | Use a string rather than the MAC address of the host system (the default) in the remote ID value in the DHCP option 82 information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                            | <b><i>string</i></b> —Character string used as the remote ID value.<br><br><b>Range:</b> 1–255 characters                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b>                           | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                              | <ul style="list-style-type: none"> <li>• <a href="#">Setting Up DHCP Option 82 on an MX Series Router (CLI Procedure)</a></li> <li>• <a href="#">Understanding DHCP Option 82 for Port Security on Switching Devices on page 4503</a></li> <li>• <a href="#">Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</a></li> <li>• <a href="#">Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</a></li> <li>• <a href="#">Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</a></li> <li>• <a href="#">Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)</a></li> <li>• <a href="#">Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4577</a></li> </ul> |

- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>

---

## use-vlan-id

---

|                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                             | use-vlan-id;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>For Platforms with Enhanced Layer 2 Software (ELS)</b> | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> option-82 <b>circuit-id</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>For Platforms Without ELS</b>                          | [edit forwarding-options helpers bootp dhcp-option82-circuit-id],<br>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82-circuit-id]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>For MX Series Platforms</b>                            | [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82 <b>circuit-id</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>                                | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b>] introduced in Junos OS Release 13.2X50-D10. (See “<a href="#">Getting Started with Enhanced Layer 2 Software</a>” on page 3 for information about ELS.)</p> <p>Hierarchy level [edit bridge-domains <i>bridge-domain-name</i> forwarding-options <b>dhcp-security</b>] introduced in Junos OS Release 14.1 for the MX Series.</p>                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                        | Use the VLAN ID rather than the VLAN name (the default) in the circuit ID value in the DHCP option 82 information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                           | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>                              | <ul style="list-style-type: none"><li>• <i>Setting Up DHCP Option 82 on an MX Series Router (CLI Procedure)</i></li><li>• <i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i></li><li>• <i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i></li><li>• <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li><li>• <i>Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li><li>• <a href="#">Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4577</a></li><li>• RFC 3046, <i>DHCP Relay Agent Information Option</i>, at <a href="http://tools.ietf.org/html/rfc3046">http://tools.ietf.org/html/rfc3046</a></li></ul> |

## vendor-id

|                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                             | <code>vendor-id &lt;string&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>For Platforms with Enhanced Layer 2 Software (ELS)</b> | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> option-82]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>For Platforms Without ELS</b>                          | [edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i> ) dhcp-option82],<br>[edit forwarding-options helpers bootp dhcp-option82],<br>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>For MX Series Platforms</b>                            | [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                                | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> ] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 3</a> for information about ELS.)<br>Hierarchy level [edit bridge-domains <i>bridge-domain-name</i> forwarding-options <b>dhcp-security</b> ] introduced in Junos OS Release 14.1 for the MX Series.                                                                                     |
| <b>Description</b>                                        | Insert a vendor ID in the DHCP option 82 information in a DHCP request packet header before forwarding or relaying the request to a DHCP server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Default</b>                                            | If <b>vendor-id</b> is not explicitly configured for DHCP option 82, then no vendor ID is set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                                            | <b>string</b> —(Optional) A single string that designates the vendor ID.<br><br><b>Range:</b> 1–255 characters<br><br><b>Default:</b> If you specify <b>vendor-id</b> with no <b>string</b> value, then the default vendor ID <b>Juniper Networks</b> is configured.                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b>                           | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                              | <ul style="list-style-type: none"> <li><i>Setting Up DHCP Option 82 on an MX Series Router (CLI Procedure)</i></li> <li><i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i></li> <li><i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i></li> <li><i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li> <li><a href="#">Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 4577</a></li> </ul> |

- *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*

## write-interval

|                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                             | <code>write-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>For Platforms with Enhanced Layer 2 Software (ELS)</b> | <p>(See <a href="#">“Getting Started with Enhanced Layer 2 Software”</a> on page 3 for information about ELS)</p> <p>[edit system processes <a href="#">dhcp-service dhcp-snooping-file</a>],<br/>[edit system processes <a href="#">dhcp-service dhcpv6-snooping-file</a>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>For Platforms Without ELS</b>                          | <p>[edit ethernet-switching-options secure-access-port dhcp-snooping-file];<br/>[edit ethernet-switching-options secure-access-port <a href="#">dhcpv6-snooping-file</a>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>For MX Series Platforms</b>                            | [edit system processes <a href="#">dhcp-service dhcp-snooping-file</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                                | <p>Statement introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Support at the [edit system processes <a href="#">dhcp-service dhcp-snooping-file</a>] hierarchy level introduced in Junos OS Release 13.2X50-D10.</p> <p>Support at the [edit system processes <a href="#">dhcp-service dhcpv6-snooping-file</a>] hierarchy level introduced in Junos OS Release 13.2X51-D20.</p> <p>Statement introduced in Junos OS Release 14.1 for the MX Series.</p> <p>Support at the [edit ethernet-switching-options <a href="#">secure-access-port dhcpv6-snooping-file</a>] hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches.</p>                                                                                                                                                                                                                                                     |
| <b>Description</b>                                        | <p>Specify how frequently the device writes the database entries from memory into the DHCP snooping database file.</p> <ul style="list-style-type: none"> <li>• If you are configuring <b>write-interval</b> at the [edit ethernet-switching-options <a href="#">secure-access-port dhcp-snooping-file</a>] or the [edit ethernet-switching-options <a href="#">secure-access-port dhcpv6-snooping-file</a>] hierarchy level, see <i>Making IP-MAC Bindings in the DHCP Snooping Database Persistent (CLI Procedure)</i>.</li> <li>• If you are configuring <b>write-interval</b> at the [edit system processes <a href="#">dhcp-service dhcp-snooping-file</a>] or the [edit system processes <a href="#">dhcp-service dhcpv6-snooping-file</a>] hierarchy level, see <a href="#">“Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Performance (CLI Procedure)”</a> on page 4575.</li> </ul> |
| <b>Options</b>                                            | <p><b>seconds</b>—Value in seconds.</p> <p><b>Range:</b> 60 through 86,400 seconds.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b>                           | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                              | <ul style="list-style-type: none"> <li>• <i>Understanding DHCP Snooping for Port Security</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



## CHAPTER 76

# Administration

- [Operational Commands on page 4657](#)

### Operational Commands

---

- [clear arp](#)
- [clear dhcp-security binding](#)
- [clear dhcp-security ipv6 binding](#)
- [clear dot1x](#)
- [clear ethernet-switching recovery-timeout](#)
- [clear security mka statistics](#)
- [show dhcp-security arp inspection statistics](#)
- [show dhcp-security binding](#)
- [show dhcp-security binding ip-source-guard](#)
- [show dhcp-security ipv6 binding](#)
- [show dhcp-security ipv6 statistics](#)
- [show dhcp-security nd-inspection statistics](#)
- [show security macsec connections](#)
- [show security macsec statistics](#)
- [show security mka sessions](#)
- [show security mka statistics](#)

## clear arp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear arp</code><br><code>&lt;hostname <i>hostname</i>&gt;</code><br><code>&lt;interface <i>interface-name</i>&gt;</code><br><code>&lt;logical-system <i>logical-system-name</i>&gt;</code><br><code>&lt;vpn <i>vpn</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 14.1 for the MX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Remove entries from the Address Resolution Protocol (ARP) table for the current CLI view. To clear entries for a specific logical system, you must first enter the <b>set cli logical-system <i>logical-system-name</i></b> command, and then issue the <b>clear arp</b> command.                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b>none</b> —Clear all entries from the ARP table.<br><br><b>hostname <i>hostname</i></b> —(Optional) Clear only the specified host entry from the ARP table.<br><br><b>interface <i>interface-name</i></b> —(Optional) Clear entries only for the specified interface from the ARP table.<br><br><b>logical-system <i>logical-system-name</i></b> —(Optional) Clear entries for only the specified logical system from the ARP table (only available in main router context).<br><br><b>vpn <i>vpn</i></b> —(Optional) Clear entries from the ARP table for the specified virtual private network (VPN). |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>set cli logical-system</i></li><li>• <i>show arp</i></li><li>• <a href="#">show dhcp-security arp inspection statistics on page 4666</a></li><li>• <a href="#">Understanding Port Security on page 4497</a></li></ul>                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <a href="#">clear arp on page 4658</a><br><a href="#">clear arp logical-system ls1 on page 4659</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Sample Output

### clear arp

```
user@host> clear arp
192.168.71.254    deleted
192.168.65.46    deleted
192.168.64.10    deleted
10.0.12.14       deleted
10.0.17.14       deleted
```

### clear arp logical-system ls1

```
user@host> clear arp logical-system ls1
192.168.71.254    deleted
192.168.65.46    deleted
192.168.64.10    deleted
10.0.12.14       deleted
10.0.17.14       deleted
```

## clear dhcp-security binding

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear dhcp-security binding</code><br><code>&lt;interface <i>interface-name</i>&gt;</code><br><code>&lt;ip-address <i>ip-address</i>&gt;</code><br><code>&lt;statistics&gt;</code><br><code>&lt;vlan <i>vlan-name</i>&gt;</code>                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Command introduced in Junos OS Release 14.1 for the MX Series.                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Clear the DHCP snooping database information.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>interface <i>interface-name</i></b>—(Optional) Clear DHCP snooping database information for the specified interface.</p> <p><b>ip-address <i>ip-address</i></b>—(Optional) Clear DHCP snooping database information for the specified IP address.</p> <p><b>statistics</b>—(Optional) Clear all DHCP snooping database statistics.</p> <p><b>vlan <i>vlan-name</i></b>—(Optional) Clear DHCP snooping database information for the specified VLAN.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dhcp-security binding on page 4668</a></li><li>• <a href="#">Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529</a></li><li>• <a href="#">Example: Configuring IP Source Guard and Dynamic ARP Inspection on MX Series Routers</a></li><li>• <a href="#">Understanding Port Security on page 4497</a></li></ul>   |

## clear dhcp-security ipv6 binding

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear dhcp-security ipv6 binding<br><all><br><interface <i>interface-name</i> ><br><ipv6-address <i>ipv6-address</i> ><br><vlan <i>vlan-name</i> >                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X51-D20 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Clear the DHCPv6 snooping database information.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>all</b>—(Optional) Clear all DHCPv6 snooping database statistics.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear DHCPv6 snooping database information for the specified interface.</p> <p><b>ipv6-address <i>ipv6-address</i></b>—(Optional) Clear DHCPv6 snooping database information for the specified IPv6 address.</p> <p><b>vlan <i>vlan-name</i></b>—(Optional) Clear DHCPv6 snooping database information for the specified VLAN.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dhcp-security ipv6 binding on page 4673</a></li> <li>• <a href="#">Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect the Switch from IPv6 Address Spoofing on page 4534</a></li> </ul>                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">clear dhcp-security ipv6 binding on page 4661</a>                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | This command produces no output.                                                                                                                                                                                                                                                                                                                                                                                                                                    |

### Sample Output

#### clear dhcp-security ipv6 binding

```
user@switch> clear dhcp-security ipv6 binding
```

## clear dot1x

---

**Syntax** `clear dot1x (firewall <counter-name> | interface <[interface-name]> | mac-address [mac-addresses] | statistics <interface interface-name>)`

**Release Information** Command introduced in Junos OS Release 9.0 for EX Series switches.  
**firewall** option added in Junos OS Release 9.5 for EX Series switches.

**Description** Reset the authentication state of an interface or delete 802.1X statistics from the switch. When you reset an interface using the **interface** or **mac-address** options, reauthentication on the interface is also triggered. The switch sends out a multicast message on the interface to restart the authentication of all connected supplicants. If a MAC address is reset, then the switch sends out a unicast message to that specific MAC address to restart authentication.

If a supplicant is sending traffic when the **clear dot1x interface** command is issued, the authenticator immediately initiates reauthentication. This process happens quickly, and it might seem that reauthentication did not occur. To verify that reauthentication has happened, issue the **show dot1x interface detail** command. The values for **Reauthentication due** and **Reauthentication interval** will be about the same.



**CAUTION:** When you clear the learned MAC addresses from an interface using the **clear dot1x interface** command, all MAC addresses are cleared, including those in static MAC bypass list.

If you have enabled Media Access Control Security (MACsec) using static secure association key (SAK) security mode on an EX Series switch, the SAKs are rotated when the **clear dot1x** command is entered. The **clear dot1x** command has no impact on MACsec when MACsec is enabled using static connectivity association keys (CAK) or any other security mode.

**Options** **firewall <counter-name>**—Clear 802.1X firewall counter statistics. If the *counter-name* option is specified, clear 802.1X firewall statistics for that counter.

**interface <[interface-name]>**—Reset the authentication state of all the supplicants (also, clears all the authentication bypassed clients) connected to the specified interface (when the interface is an authenticator) or reset the authentication state for the interface itself (when the interface is a supplicant).

**mac-address [mac-addresses]**—Reset the authentication state of the specified MAC addresses.

**statistics <interface interface-name>**—Clear 802.1X statistics on all 802.1X-enabled interfaces. If the **interface** option is specified, clear 802.1X firewall statistics for that interface or interfaces.

**Required Privilege Level** view

**Related Documentation**

- [show dot1x on page 2009](#)
- [Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch on page 1852](#)
- [Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 1910](#)

**List of Sample Output**

[clear dot1x firewall c1 on page 4663](#)  
[clear dot1x interface ge-1/0/0 ge-2/0/0 ge-2/0/0 ge5/0/0 on page 4663](#)  
[clear dot1x mac-address 00:04:ae:cd:23:5f on page 4663](#)  
[clear dot1x statistics interface ge-1/0/1 on page 4663](#)

## Sample Output

`clear dot1x firewall c1`

```
user@switch> clear dot1x firewall c1
```

`clear dot1x interface ge-1/0/0 ge-2/0/0 ge-2/0/0 ge5/0/0`

```
user@switch> clear dot1x interface ge-1/0/0 ge-2/0/0 ge-2/0/0 ge5/0/0
```

`clear dot1x mac-address 00:04:ae:cd:23:5f`

```
user@switch> clear dot1x mac-address 00:04:ae:cd:23:5f
```

`clear dot1x statistics interface ge-1/0/1`

```
user@switch> clear dot1x statistics interface ge-1/0/1
```

## clear ethernet-switching recovery-timeout

---

|                                 |                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear ethernet-switching recovery-timeout                                                                                                                                             |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.                                                                                                            |
| <b>Description</b>              | Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch, and restore the interfaces to service.                  |
| <b>Options</b>                  | <b>none</b> —Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch and restore these interfaces to service.    |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)</a> on page 2206</li></ul> |
| <b>Output Fields</b>            | This command produces no output.                                                                                                                                                      |



## clear security mka statistics

---

|                                 |                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear security mka statistics<br><interface <i>interface-name</i> >                                                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                                     |
| <b>Description</b>              | Clear—reset to zero (0)—all MACsec Key Agreement (MKA) protocol statistics.<br><br>You are clearing the statistics that are viewed using the <b>show security mka statistics</b> when you enter this command.                                                                 |
| <b>Options</b>                  | <b>none</b> —Clear all MKA counters for all interfaces on the switch.<br><br><b>interface <i>interface-name</i></b> —(Optional) Clear MKA traffic counters for the specified interface only.                                                                                  |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show security mka statistics on page 4688</a></li> <li>• <a href="#">show security mka sessions on page 4686</a></li> <li>• <a href="#">Understanding Media Access Control Security (MACsec) on page 4520</a></li> </ul> |

## Sample Output

### clear security mka statistics

```
user@switch> clear security mka statistics
```

## show dhcp-security arp inspection statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show dhcp-security arp inspection statistics</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Command introduced in Junos OS Release 14.1 for the MX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display Address Resolution Protocol (ARP) inspection statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dhcp-security binding on page 4668</a></li> <li>• <a href="#">clear dhcp-security binding on page 4660</a></li> <li>• <a href="#">clear interfaces statistics</a></li> <li>• <a href="#">Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529</a></li> <li>• <a href="#">Example: Configuring IP Source Guard and Dynamic ARP Inspection on MX Series Routers</a></li> <li>• <a href="#">Understanding Port Security on page 4497</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show dhcp-security arp inspection statistics on page 4666</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>            | <p><a href="#">Table 501 on page 4666</a> lists the output fields for the <b>show dhcp-security arp inspection statistics</b> command. Output fields are listed in the approximate order in which they appear.</p> <p>The IP source guard database table shows the untrusted access interfaces in VLANs that have been enabled for IP source guard. The entries include the VLAN 802.1Q tag IDs if there are any, and the IP addresses and MAC addresses that are bound to one another.</p>                                                                                                |

**Table 501: show dhcp-security arp inspection statistics Output Fields**

| Field Name                 | Field Description                                      | Level of Output |
|----------------------------|--------------------------------------------------------|-----------------|
| <b>Interface</b>           | Interface on which ARP inspection has been applied.    | All levels      |
| <b>Packets received</b>    | Total number of packets that underwent ARP inspection. | All levels      |
| <b>ARP inspection pass</b> | Total number of packets that passed ARP inspection.    | All levels      |
| <b>ARP inspection fail</b> | Total number of packets that failed ARP inspection.    | All levels      |

## Sample Output

### show dhcp-security arp inspection statistics

```
user@device> show dhcp-security arp inspection statistics
```

| Interface   | Packets received | ARP inspection pass | ARP inspection fail |
|-------------|------------------|---------------------|---------------------|
| ge-0/0/30.0 | 7                | 7                   | 0                   |
| ge-0/0/4.0  | 3                | 3                   | 0                   |
| ge-0/0/6.0  | 72               | 4                   | 68                  |

## show dhcp-security binding

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show dhcp-security binding &lt;interface <i>interface-name</i>&gt; &lt;ip-address <i>ip-address</i>&gt; &lt;ip-source-guard <i>ip-sg-name</i>&gt; &lt;statistics&gt; &lt;vlan <i>vlan-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Command introduced in Junos OS Release 14.1 for the MX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Display the DHCP snooping database information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>interface <i>interface-name</i></b>—(Optional) Display the DHCP snooping database information for an interface.</p> <p><b>ip-address <i>ip-address</i></b>—(Optional) Display the DHCP snooping database information for an IP address.</p> <p><b>vlan <i>vlan-name</i></b>—(Optional) Display the DHCP snooping database information for a VLAN.</p>                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dhcp-security binding ip-source-guard on page 4671</a></li> <li>• <a href="#">clear dhcp-security binding on page 4660</a></li> <li>• <a href="#">Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529</a></li> <li>• <a href="#">Example: Configuring IP Source Guard and Dynamic ARP Inspection on MX Series Routers</a></li> <li>• <a href="#">Understanding Port Security on page 4497</a></li> </ul> |
| <b>List of Sample Output</b>    | <p><a href="#">show dhcp-security binding on page 4669</a></p> <p><a href="#">show dhcp-security binding interface on page 4669</a></p> <p><a href="#">show dhcp-security binding ip-address on page 4669</a></p> <p><a href="#">show dhcp-security binding vlan on page 4670</a></p>                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | <p><a href="#">Table 502 on page 4668</a> lists the output fields for the <b>show dhcp-security binding</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                    |

**Table 502: show dhcp-security binding Output Fields**

| Field Name  | Field Description                                           | Level of Output |
|-------------|-------------------------------------------------------------|-----------------|
| IP Address  | IP address of the network device; bound to the MAC address. | All levels      |
| MAC address | MAC address of the network device; bound to the IP address. | All levels      |
| VLAN        | VLAN name of the network device whose MAC address is shown. | All levels      |

Table 502: show dhcp-security binding Output Fields (*continued*)

| Field Name       | Field Description                                                                                                                                                                                                         | Level of Output |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Expires</b>   | The time, in seconds, remaining before the lease of the IP address to the MAC address expires.                                                                                                                            | All levels      |
| <b>State</b>     | Specifies whether the IP address is: <ul style="list-style-type: none"> <li>• <b>BOUND</b>: Leased to the MAC address for a limited period of time.</li> <li>• <b>STATIC</b>: Attached to a fixed MAC address.</li> </ul> | All levels      |
| <b>Interface</b> | Interface address (port).                                                                                                                                                                                                 | All levels      |

## Sample Output

### show dhcp-security binding

```
user@device> show dhcp-security binding
```

| IP address | MAC address       | Vlan   | Expires | State  | Interface  |
|------------|-------------------|--------|---------|--------|------------|
| 30.1.1.10  | 00:10:00:20:00:01 | vlan20 | 0       | STATIC | ge-0/0/4.0 |
| 30.1.1.18  | 00:10:94:00:00:34 | vlan20 | 86287   | BOUND  | ge-0/0/6.0 |
| 30.1.1.15  | 00:10:94:00:00:55 | vlan20 | 86265   | BOUND  | ge-0/0/4.0 |
| 30.1.1.16  | 00:10:94:00:00:56 | vlan20 | 86265   | BOUND  | ge-0/0/4.0 |
| 30.1.1.19  | 00:10:94:00:00:5b | vlan20 | 86287   | BOUND  | ge-0/0/6.0 |
| 30.1.1.20  | 00:10:94:00:00:5c | vlan20 | 86287   | BOUND  | ge-0/0/6.0 |
| 30.1.1.21  | 00:10:94:00:00:5d | vlan20 | 86287   | BOUND  | ge-0/0/6.0 |
| 30.1.1.17  | 00:10:94:00:00:68 | vlan20 | 86265   | BOUND  | ge-0/0/4.0 |

### show dhcp-security binding interface

```
user@device> show dhcp-security binding interface ge-0/0/6
```

| IP address | MAC address       | Vlan   | Expires | State | Interface  |
|------------|-------------------|--------|---------|-------|------------|
| 30.1.1.18  | 00:10:94:00:00:34 | vlan20 | 86282   | BOUND | ge-0/0/6.0 |
| 30.1.1.19  | 00:10:94:00:00:5b | vlan20 | 86282   | BOUND | ge-0/0/6.0 |
| 30.1.1.20  | 00:10:94:00:00:5c | vlan20 | 86282   | BOUND | ge-0/0/6.0 |
| 30.1.1.21  | 00:10:94:00:00:5d | vlan20 | 86282   | BOUND | ge-0/0/6.0 |

### show dhcp-security binding ip-address

```
user@device> show dhcp-security binding ip-address
```

| IP address | MAC address       | Vlan   | Expires | State | Interface  |
|------------|-------------------|--------|---------|-------|------------|
| 30.1.1.18  | 00:10:94:00:00:34 | vlan20 | 86282   | BOUND | ge-0/0/6.0 |

### show dhcp-security binding vlan

```
user@device> show dhcp-security binding vlan vlan20
```

| IIP address | MAC address       | Vlan   | Expires | State | Interface  |
|-------------|-------------------|--------|---------|-------|------------|
| 30.1.1.18   | 00:10:94:00:00:34 | vlan20 | 86282   | BOUND | ge-0/0/6.0 |

## show dhcp-security binding ip-source-guard

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show dhcp-security binding ip-source-guard</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Command introduced in Junos OS Release 14.1 for the MX Series.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Display IP source guard database table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dhcp-security binding on page 4668</a></li> <li>• <a href="#">clear dhcp-security binding on page 4660</a></li> <li>• <a href="#">Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 4529</a></li> <li>• <a href="#">Example: Configuring IP Source Guard and Dynamic ARP Inspection on MX Series Routers</a></li> <li>• <a href="#">Understanding Port Security on page 4497</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show dhcp-security binding ip-source-guard on page 4672</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <p><a href="#">Table 503 on page 4671</a> lists the output fields for the <b>show dhcp-security binding ip-source-guard</b> command. Output fields are listed in the approximate order in which they appear.</p> <p>The IP source guard database table shows the untrusted access interfaces in VLANs that have been enabled for IP source guard. The entries include the IP addresses and MAC addresses that are bound to one another.</p>                                                                                         |

**Table 503: show dhcp-security binding ip-source-guard Output Fields**

| Field Name         | Field Description                                                                                                                                                                                                                     | Level of Output |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>IP Address</b>  | IP address of the network device; bound to the MAC address.                                                                                                                                                                           | All levels      |
| <b>MAC address</b> | MAC address of the network device; bound to the IP address.                                                                                                                                                                           | All levels      |
| <b>VLAN</b>        | VLAN name of the network device whose MAC address is shown.                                                                                                                                                                           | All levels      |
| <b>Expires</b>     | The time, in seconds, remaining before the lease of the IP address to the MAC address expires.                                                                                                                                        | All levels      |
| <b>State</b>       | Specifies whether the IP address is: <ul style="list-style-type: none"> <li>• <b>BOUND:</b> Temporarily leased to the MAC address for a limited period of time.</li> <li>• <b>STATIC:</b> Attached to a fixed MAC address.</li> </ul> | All levels      |
| <b>Interface</b>   | Interface address (port).                                                                                                                                                                                                             | All levels      |

## Sample Output

### show dhcp-security binding ip-source-guard

```
user@device> show dhcp-security binding ip-source-guard
```

| IP address | MAC address       | Vlan   | Expires | State  | Interface  |
|------------|-------------------|--------|---------|--------|------------|
| 30.1.1.10  | 00:10:00:20:00:01 | vlan20 | 0       | STATIC | ge-0/0/4.0 |
| 30.1.1.18  | 00:10:94:00:00:34 | vlan20 | 86276   | BOUND  | ge-0/0/6.0 |
| 30.1.1.15  | 00:10:94:00:00:55 | vlan20 | 86254   | BOUND  | ge-0/0/4.0 |
| 30.1.1.16  | 00:10:94:00:00:56 | vlan20 | 86254   | BOUND  | ge-0/0/4.0 |
| 30.1.1.19  | 00:10:94:00:00:5b | vlan20 | 86276   | BOUND  | ge-0/0/6.0 |
| 30.1.1.20  | 00:10:94:00:00:5c | vlan20 | 86276   | BOUND  | ge-0/0/6.0 |
| 30.1.1.21  | 00:10:94:00:00:5d | vlan20 | 86276   | BOUND  | ge-0/0/6.0 |
| 30.1.1.17  | 00:10:94:00:00:68 | vlan20 | 86254   | BOUND  | ge-0/0/4.0 |



## show dhcp-security ipv6 binding

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show dhcp-security ipv6 binding<br><interface <i>interface-name</i> ><br><ipv6-address <i>ipv6-address</i> ><br><vlan <i>vlan-name</i> >                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X51-D20 for EX Series switches.                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display IPv6 address to MAC address bindings, also known as the DHCPv6 binding table or DHCPv6 snooping table.                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>interface <i>interface-name</i></b>—(Optional) Display the DHCPv6 snooping table for the specified interface.</p> <p><b>ipv6-address <i>ipv6-address</i></b>—(Optional) Display the DHCPv6 snooping table for the specified IPv6 address.</p> <p><b>vlan <i>vlan-name</i></b>—(Optional) Display the DHCPv6 snooping table for a VLAN.</p>                                                                         |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dhcp-security ipv6 statistics on page 4675</a></li> <li>• <a href="#">clear dhcp-security ipv6 binding on page 4661</a></li> <li>• <a href="#">Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect the Switch from IPv6 Address Spoofing on page 4534</a></li> </ul>                                                           |
| <b>List of Sample Output</b>    | <p><a href="#">show dhcp-security ipv6 binding on page 4674</a></p> <p><a href="#">show dhcp-security ipv6 binding interface ge-0/0/4.0 on page 4674</a></p>                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <p><a href="#">Table 503 on page 4671</a> lists the output fields for the <b>show dhcp-security ipv6 binding</b> command. Output fields are listed in the approximate order in which they appear.</p> <p>The DHCPv6 binding table shows the untrusted access interfaces in VLANs that have been enabled for DHCPv6 snooping. The entries include the IPv6 addresses and MAC addresses that are bound to one another.</p> |

**Table 504: show dhcp-security ipv6 binding Output Fields**

| Field Name   | Field Description                                                                                                                                                                                                                                                                                                               | Level of Output |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| IPv6 address | IPv6 addresses of the network device; bound to the MAC address. There are two entries added for each client: one with the link-local IPv6 address, which is used by the client for DHCP transactions, and another with the IPv6 address assigned by the server. The link-local address always has the prefix <b>fe80::/10</b> . | All levels      |
| MAC address  | MAC address of the network device; bound to the IPv6 address.                                                                                                                                                                                                                                                                   | All levels      |
| VLAN         | VLAN name of the network device whose MAC address is shown.                                                                                                                                                                                                                                                                     | All levels      |

Table 504: show dhcp-security ipv6 binding Output Fields (*continued*)

| Field Name       | Field Description                                                                                                                                                                                                                       | Level of Output |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Expires</b>   | The time, in seconds, remaining before the lease of the IPv6 address to the MAC address expires.                                                                                                                                        | All levels      |
| <b>State</b>     | Specifies whether the IPv6 address is: <ul style="list-style-type: none"> <li>• <b>BOUND</b>: Temporarily leased to the MAC address for a limited period of time.</li> <li>• <b>STATIC</b>: Attached to a fixed MAC address.</li> </ul> | All levels      |
| <b>Interface</b> | Interface address (port).                                                                                                                                                                                                               | All levels      |

## Sample Output

### show dhcp-security ipv6 binding

```

user@switch> show dhcp-security ipv6 binding
IPv6 address      MAC address      Vlan    Expires  State  Interface
2001::10:10:0:1    00:10:94:00:55:0b v1an20   0        BOUND  ge-0/0/1.0

fe80::210:94ff:fe00:550b 00:10:94:00:55:0b v1an20   0        BOUND  ge-0/0/1.0

2000::10:10:0:2    00:10:94:00:00:34 v1an20   0        BOUND  ge-0/0/2.0

fe80::210:94ff:fe00:0034 00:10:94:00:00:34 v1an20   0        BOUND  ge-0/0/2.0

2000::10:10:0:3    00:10:94:00:00:55 v1an20   0        BOUND  ge-0/0/3.0

fe80::210:94ff:fe00:0055 00:10:94:00:00:55 v1an20   0        BOUND  ge-0/0/3.0

```

## Sample Output

### show dhcp-security ipv6 binding interface ge-0/0/4.0

```

user@switch> show dhcp-security ipv6 binding interface ge-0/0/4.0
IPv6 address      MAC address      Vlan    Expires  State  Interface
2001::1           00:10:00:20:00:01 v1an20   0        STATIC ge-0/0/4.0

fe80::210:00ff:fe20:0001 00:10:00:20:00:01 v1an20   0        STATIC ge-0/0/4.0

```

---

## show dhcp-security ipv6 statistics

---

|                                 |                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show dhcp-security ipv6 statistics</b>                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X51-D20 for EX Series switches.                                                                                                                        |
| <b>Description</b>              | Display DHCPv6 statistics.                                                                                                                                                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dhcp-security ipv6 binding on page 4673</a></li><li>• <a href="#">show dhcp-security nd-inspection statistics on page 4678</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">show dhcp-security ipv6 statistics on page 4677</a>                                                                                                                                   |
| <b>Output Fields</b>            | <a href="#">Table 155 on page 1749</a> lists the output fields for the <b>show dhcp-security ipv6 statistics</b> command. Output fields are listed in the approximate order in which they appear. |

Table 505: show dhcp-security ipv6 statistics Output Fields

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DHCPv6 messages | <p>Number of DHCPv6 messages exchanged.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of DHCPv6 messages exchanged.</li> <li>• <b>Solicit</b>—Number of DHCPv6 messages of type Solicit. A client sends a Solicit message to locate servers.</li> <li>• <b>Advertise</b>—Number of DHCPv6 messages of type Advertise. A server sends an Advertise message, in response to a Solicit message, to indicate that it is available for DHCPv6 service.</li> <li>• <b>Request</b>—Number of DHCPv6 messages of type Request. A client sends a Request message to request configuration parameters from a server.</li> <li>• <b>Reply</b>—Number of DHCPv6 messages of type Reply. A server sends a Reply message in response to a Solicit, Request, Renew, Rebind, Confirm, Information Request, Release, or Decline message.</li> <li>• <b>Confirm</b>—Number of DHCPv6 messages of type Confirm. A client sends a Confirm message to any available server to determine whether the addresses it was assigned are still appropriate for the link to which the client is connected.</li> <li>• <b>Decline</b>—Number of DHCPv6 messages of type Decline. A client sends a Decline message to a server to indicate that one or more of the addresses assigned by the server are already in use on the link to which the client is connected.</li> <li>• <b>Release</b>—Number of DHCPv6 messages of type Release. A client sends a Release message to the server to indicate that the client will no longer use one or more of the assigned addresses.</li> <li>• <b>Renew</b>—Number of DHCPv6 messages of type Renew. A client sends a Renew message to the server to extend the lifetimes on the addresses assigned to the client by that server and to update other configuration parameters received by that server.</li> <li>• <b>Rebind</b>—Number of DHCPv6 messages of type Rebind. A client sends a Rebind message to any available server after receiving no reply to a Renew message.</li> <li>• <b>Relay-forward</b>—Number of DHCPv6 messages of type Relay-forward. A relay agent sends a Relay-forward message to relay messages to servers, either directly or through another relay agent. The received message is encapsulated in an option in the Relay-forward message.</li> <li>• <b>Relay-reply</b>—Number of DHCPv6 messages of type Relay-reply. A server sends a Relay-reply message to a relay agent containing a message that the relay agent delivers to a client.</li> <li>• <b>Information-request</b>—Number of DHCPv6 messages of type Information-request. A client sends an Information-request message to a server to request configuration parameters without the assignment of any IP addresses to the client.</li> <li>• <b>Reconfigure</b>—Number of DHCPv6 messages of type Reconfigure. A server sends a Reconfigure message to a client to inform the client that the server has new or updated configuration parameters, and that the client needs to initiate a Renew/Reply or Information-request/Reply transaction with the server in order to receive the updated information.</li> </ul> |
| Packets dropped | <p>Number of packets not considered for DHCPv6 snooping because of errors.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by DHCPv6 snooping.</li> <li>• <b>No configuration</b>—Number of packets discarded because they did not have a valid configuration.</li> <li>• <b>No VLAN</b>—Number of packets discarded because they did not belong to a valid VLAN.</li> <li>• <b>No interface</b>—Number of packets discarded because they did not belong to a valid interface.</li> <li>• <b>Request on trusted port</b>—Number of packets discarded because a Request message was received on a trusted port.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## Sample Output

### show dhcp-security ipv6 statistics

```
user@host> show dhcp-security ipv6 statistics
DHCPv6 messages:
  Total                32
  Solicit               1
  Advertise             1
  Request              3
  Reply                5
  Confirm              1
  Decline              2
  Release              9
  Renew               4
  Rebind               2
  Relay forward        1
  Relay reply          1
  Information request  1
  Reconfigure          2

Packets dropped:
  Total                0
  No configuration     0
  No VLAN              0
  No interface         0
  Request on trusted port 0
```

## show dhcp-security nd-inspection statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show dhcp-security nd-inspection statistics<br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X51-D20 for EX Series switches.                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Display IPv6 Neighbor Discovery inspection statistics to determine whether there is IPv6 address spoofing on the network.                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>interface <i>interface-name</i></b> —(Optional) Display Neighbor Discovery inspection statistics for the specified interface.                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dhcp-security ipv6 binding on page 4673</a></li> <li>• <a href="#">Enabling IPv6 Neighbor Discovery Inspection on page 4547</a></li> <li>• <a href="#">Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect the Switch from IPv6 Address Spoofing on page 4534</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show dhcp-security nd inspection statistics on page 4678</a><br><a href="#">show dhcp-security nd inspection statistics interface ge-0/0/1.0 on page 4679</a>                                                                                                                                                                                              |
| <b>Output Fields</b>            | Table 501 on page 4666 lists the output fields for the <b>show dhcp-security nd inspection statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                             |

Table 506: show dhcp-security nd inspection statistics Output Fields

| Field Name                | Field Description                                                     | Level of Output |
|---------------------------|-----------------------------------------------------------------------|-----------------|
| <b>Interface</b>          | Interface on which Neighbor Discovery inspection has been applied.    | All levels      |
| <b>Packets received</b>   | Total number of packets that underwent Neighbor Discovery inspection. | All levels      |
| <b>ND inspection pass</b> | Total number of packets that passed Neighbor Discovery inspection.    | All levels      |
| <b>ND inspection fail</b> | Total number of packets that failed Neighbor Discovery inspection.    | All levels      |

## Sample Output

### show dhcp-security nd inspection statistics

```

user@switch> show dhcp-security nd inspection statistics
Interface      ND Packets received  ND inspection pass  ND inspection failed
ge-0/0/1.0      7                    5                    2
ge-0/0/2.0      10                   10                   0
ge-0/0/3.0      12                   12                   0

```

## Sample Output

`show dhcp-security nd inspection statistics interface ge-0/0/1.0`

```
user@switch> show dhcp-security nd inspection statistics interface ge-0/0/1.0
Interface      ND Packets received  ND inspection pass  ND inspection failed
ge-0/0/1.0          7                   5                   2
```

## show security macsec connections

|                                 |                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show security macsec connections<br><interface <i>interface-name</i> >                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                              |
| <b>Description</b>              | Display the status of the active MACsec connections on the switch.                                                                                                                                                     |
| <b>Options</b>                  | <b>none</b> —Display MACsec connection information for all interfaces on the switch.<br><br><b>interface <i>interface-name</i></b> —(Optional) Display MACsec connection information for the specified interface only. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show security macsec statistics on page 4682</a></li> </ul>                                                                                                       |
| <b>List of Sample Output</b>    | <a href="#">show security macsec connections on page 4681</a>                                                                                                                                                          |
| <b>Output Fields</b>            | <a href="#">Table 507 on page 4680</a> lists the output fields for the <b>show security macsec connections</b> command. Output fields are listed in the approximate order in which they appear.                        |

**Table 507: show security macsec connections Output Fields**

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Fields for Interface</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Interface name</b>       | Name of the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>CA name</b>              | Name of the connectivity association.<br><br>A connectivity association is named using the <b>connectivity-association</b> statement when you are enabling MACsec.                                                                                                                                                                                                                                                                                     |
| <b>Cipher suite</b>         | Name of the cipher suite used for encryption.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Encryption</b>           | Encryption setting. Encryption is enabled when this output is <b>on</b> and disabled when this output is <b>off</b> .<br><br>The encryption setting is set using the <b>no-encryption</b> statement in the connectivity association when using static connectivity association key (CAK) security mode and is set using the <b>encryption</b> statement in the secure channel when using static secure association key (SAK) or dynamic security mode. |
| <b>Key server offset</b>    | Offset setting.<br><br>The offset is set using the <b>offset</b> statement when configuring the connectivity association when using static connectivity association key (CAK) or dynamic security mode or the secure channel when using static secure association key (SAK) security mode.                                                                                                                                                             |



Table 507: show security macsec connections Output Fields (*continued*)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Include SCI</b>    | <p>SCI tagging. The SCI tag is included on packets in a secure channel when this output is <b>yes</b>, and not included on packets in a secure channel when this output is <b>no</b>.</p> <p>SCI tagging is automatically enabled on EX4300 switch interfaces that have enabled MACsec using static connectivity association key (CAK) or dynamic security mode. You can enable SCI tagging using the <b>include-sci</b> statement in the connectivity association.</p> |
| <b>Replay protect</b> | <p>Replay protection setting. Replay protection is enabled when this output is <b>on</b> and disabled when this output is <b>off</b>.</p> <p>You can enable replay protection using the <b>replay-protect</b> statement in the connectivity association.</p>                                                                                                                                                                                                            |
| <b>Replay window</b>  | <p>Replay protection window setting. This output is set to <b>0</b> when replay protection is disabled, and is the size of the replay window, in number of packets, when replay protection is enabled.</p> <p>The size of the replay window is configured using the <b>replay-window-size</b> statement in the connectivity association.</p>                                                                                                                            |

## Sample Output

### show security macsec connections

```

user@host> show security macsec connections
Interface name: xe-0/1/0
  CA name: CA1
  Cipher suite: GCM-AES-128   Encryption: on
  Key server offset: 0        Include SCI: no
  Replay protect: off         Replay window: 0

```

## show security macsec statistics

**Syntax** show security macsec statistics  
<brief | detail>  
<interface *interface-name*>

**Release Information** Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches.  
Command introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.

**Description** Display Media Access Control Security (MACsec) statistics.

**Options** **none**—Display MACsec statistics in brief form for all interfaces on the switch.

**brief | detail**—(Optional) Display the specified level of output. Using the **brief** option is equivalent to entering the command with no options (the default). The **detail** option displays additional fields that are not visible in the **brief** output.



**NOTE:** The field names that only appear in this command output when you enter the **detail** option are mostly useful for debugging purposes by Juniper Networks support personnel.

**interface *interface-name***—(Optional) Display MACsec statistics for the specified interface only.

**Required Privilege Level** view

**Related Documentation** • [show security macsec connections on page 4680](#)

**List of Sample Output** [show security macsec statistics interface xe-0/1/0 detail on page 4684](#)

**Output Fields** [Table 508 on page 4682](#) lists the output fields for the **show security macsec statistics** command. Output fields are listed in the approximate order in which they appear.

The field names that appear in this command output only when you enter the **detail** option are mostly useful for debugging purposes by Juniper Networks support personnel. Those field names are, therefore, not included in this table.

**Table 508: show security macsec statistics Output Fields**

| Field Name                                   | Field Description      | Level of Output |
|----------------------------------------------|------------------------|-----------------|
| <b>Interface name</b>                        | Name of the interface. | All levels      |
| <b>Fields for Secure Channel transmitted</b> |                        |                 |

Table 508: show security macsec statistics Output Fields (*continued*)

| Field Name                                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output |
|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Encrypted packets</b>                         | <p>Total number of packets transmitted out of the interface in the secure channel that were secured and encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p>                                                                                                                                                              | All levels      |
| <b>Encrypted bytes</b>                           | <p>Total number of bytes transmitted out of the interface in the secure channel that were secured and encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p>                                                                                                                                                                | All levels      |
| <b>Protected packets</b>                         | <p>Total number of packets transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p>                                                                                                                                                          | All levels      |
| <b>Protected bytes</b>                           | <p>Total number of bytes transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p>                                                                                                                                                            | All levels      |
| <b>Fields for Secure Association transmitted</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                   |                 |
| <b>Encrypted packets</b>                         | <p>Total number of packets transmitted out of the interface in the connectivity association that were secured and encrypted using MACsec.</p> <p>The total includes the data packets transmitted in the secure channel and secured using a SAK and the control packets secured using a connectivity association key (CAK).</p>                                                                                                    | All levels      |
| <b>Protected packets</b>                         | <p>Total number of packets transmitted out of the interface in the connectivity association that were secured but not encrypted using MACsec.</p> <p>The total includes the data packets transmitted in the secure channel and secured using a SAK and the control packets secured using a connectivity association key (CAK).</p>                                                                                                | All levels      |
| <b>Fields for Secure Channel received</b>        |                                                                                                                                                                                                                                                                                                                                                                                                                                   |                 |
| <b>Accepted packets</b>                          | <p>The number of received packets that have been accepted by the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link.</p> <p>A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check.</p> <p>This counter increments for traffic that is and is not encrypted using MACsec.</p> | All levels      |

Table 508: show security macsec statistics Output Fields (*continued*)

| Field Name                                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Validated bytes</b>                        | <p>The number of bytes that have been validated by the MACsec integrity check and received on the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link.</p> <p>This counter does not increment when MACsec encryption is disabled.</p>                                                                                                 | All levels      |
| <b>Decrypted bytes</b>                        | <p>The number of bytes received in the secure channel on the interface that have been decrypted. The secure channel is used to send all data plane traffic on a MACsec-enabled link.</p> <p>An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.</p>          | All levels      |
| <b>Fields for Secure Association received</b> |                                                                                                                                                                                                                                                                                                                                                                                                       |                 |
| <b>Accepted packets</b>                       | <p>The number of received packets that have been accepted in the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface.</p> <p>A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check.</p>                                                  | All levels      |
| <b>Validated bytes</b>                        | <p>The number of bytes that have been validated by the MACsec integrity check and received on the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface.</p> <p>This counter does not increment when MACsec encryption is disabled.</p>                                                                                        | All levels      |
| <b>Decrypted bytes</b>                        | <p>The number of bytes received in the connectivity association on the interface that have been decrypted. The counter includes all control and data plane traffic accepted on the interface.</p> <p>An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.</p> | All levels      |

## Sample Output

### show security macsec statistics interface xe-0/1/0 detail

```
user@host> show security macsec statistics interface xe-0/1/0 detail
```

```
Interface name: xe-0/1/0
Secure Channel transmitted
  Encrypted packets: 123858
  Encrypted bytes:   32190903
  Protected packets: 0
  Protected bytes:   0
Secure Association transmitted
```

```
Encrypted packets: 123858
Protected packets: 0
Secure Channel received
  Accepted packets: 123877
  Validated bytes: 0
  Decrypted bytes: 32196238
Secure Association received
  Accepted packets: 123877
  Validated bytes: 0
  Decrypted bytes: 32196238
Error and debug
Secure Channel transmitted packets
  Untagged: 0, Too long: 0
Secure Channel received packets
  Control: 0, Tagged miss: 3202804
  Untagged hit: 0, Untagged: 0
  No tag: 0, Bad tag: 0
  Unknown SCI: 0, No SCI: 0
  Control pass: 0, Control drop: 0
  Uncontrol pass: 123877, Uncontrol drop: 0
  Hit dropped: 0, Invalid accept: 0
  Late drop: 0, Delayed accept: 0
  Unchecked: 0, Not valid drop: 0
  Not using SA drop: 0, Unused SA accept: 0
```

## show security mka sessions

|                                 |                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show security mka sessions<br><interface <i>interface-name</i> >                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                                |
| <b>Description</b>              | Display MACsec Key Agreement (MKA) session information.                                                                                                                                                                                                  |
| <b>Options</b>                  | <ul style="list-style-type: none"> <li><b>interface <i>interface-name</i></b>—(Optional) Display the MKA session information for the specified interface only.</li> </ul>                                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">show security mka statistics on page 4688</a></li> <li><a href="#">show security macsec connections on page 4680</a></li> <li><a href="#">show security macsec statistics on page 4682</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show security mka sessions on page 4687</a>                                                                                                                                                                                                  |
| <b>Output Fields</b>            | Table 509 on page 4686 lists the output fields for the <b>show security mka sessions</b> command. Output fields are listed in the approximate order in which they appear.                                                                                |

Table 509: show security mka sessions Output Fields

| Field Name        | Field Description                                                                                                                                    |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface name    | Name of the interface.                                                                                                                               |
| Member identifier | Name of the member identifier.                                                                                                                       |
| CAK name          | Name of the Connectivity Association Key (CAK).<br>The CAK is configured using the <b>cak</b> keyword when configuring the pre-shared key.           |
| Transmit interval | The transmit interval.                                                                                                                               |
| Outbound SCI      | Name of the outbound secure channel identifier.                                                                                                      |
| Message number    | Number of the last data message.                                                                                                                     |
| Key number        | Key number.                                                                                                                                          |
| Key server        | Key server status.<br>The switch is the key server when this output is <b>yes</b> . The switch is not the key server when this output is <b>no</b> . |

Table 509: show security mka sessions Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                  |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------|
| Key server priority         | The key server priority.<br><br>The key server priority can be set using the <b>key-server-priority</b> statement. |
| Latest SAK AN               | Name of the latest secure association key (SAK) association number.                                                |
| Latest SAK KI               | Name of the latest secure association key (SAK) key identifier.                                                    |
| <b>Fields for Peer list</b> |                                                                                                                    |
| Member identifier           | Name of the member identifier.                                                                                     |
| Hold time                   | Hold time, in seconds.                                                                                             |
| Message number              | Number of the last data message                                                                                    |
| SCI                         | Name of the secure channel identifier.                                                                             |
| Lowest acceptable PN        | Number of the lowest acceptable packet number (PN).                                                                |

## Sample Output

### show security mka sessions

```
user@host> show security mka sessions
```

```
Interface name: xe-0/1/0
Member identifier: 0CCBEE42F8778300F8D0C1DC
CAK name: 1234567890
Transmit interval: 2000(ms)
Outbound SCI: 2C:6B:F5:9D:4B:1B/1
Message number: 1526465    Key number: 0
Key server: no             Key server priority: 15
Latest SAK AN: 0           Latest SAK KI: 4F18CE25228178FD15976E4C/1
Previous SAK AN: 0         Previous SAK KI: 000000000000000000000000/0
Peer list
1. Member identifier: 4F18CE25228178FD15976E4C (live)
   Message number: 1526484 Hold time: 14500 (ms)
   SCI: 2C:6B:F5:9D:3A:1B/1
   Lowest acceptable PN: 121198
```

## show security mka statistics

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show security mka statistics<br><interface <i>interface-name</i> >                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.                                                                                              |
| <b>Description</b>              | Display MACsec Key Agreement (MKA) protocol statistics.<br><br>The output for this command does not include statistics for MACsec data traffic. For MACsec data traffic statistics, see <a href="#">show security macsec statistics</a> .              |
| <b>Options</b>                  | <ul style="list-style-type: none"> <li><b>interface <i>interface-name</i></b>—(Optional) Display the MKA information for the specified interface only.</li> </ul>                                                                                      |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">show security mka sessions on page 4686</a></li> <li><a href="#">show security macsec statistics on page 4682</a></li> <li><a href="#">show security macsec connections on page 4680</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show security mka statistics on page 4689</a>                                                                                                                                                                                              |
| <b>Output Fields</b>            | <a href="#">Table 510 on page 4688</a> lists the output fields for the <b>show security mka statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                            |

**Table 510: show security mka statistics Output Fields**

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Received packets</b>         | <p>Number of received MKA control packets.</p> <p>This counter increments for received MKA control packets only. This counter does not increment when data packets are received.</p>                                                                                                                        |
| <b>Transmitted packets</b>      | <p>Number of transmitted MKA packets</p> <p>This counter increments for transmitted MKA control packets only. This counter does not increment when data packets are transmitted.</p>                                                                                                                        |
| <b>Version mismatch packets</b> | Number of version mismatch packets.                                                                                                                                                                                                                                                                         |
| <b>CAK mismatch packets</b>     | <p>Number of Connectivity Association Key (CAK) mismatch packets.</p> <p>This counter increments when the connectivity association key (CAK) and connectivity association key name (CKN), which are user-configured values that have to match to enable MACsec, do not match for an MKA control packet.</p> |



Table 510: show security mka statistics Output Fields *(continued)*

| Field Name                           | Field Description                                                                                                                                                                 |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ICV mismatch packets                 | Number of ICV mismatched packets.<br><br>This counter increments when the connectivity association key (CAK) value does not match on both ends of a MACsec-secured Ethernet link. |
| Duplicate message identifier packets | Number of duplicate message identifier packets.                                                                                                                                   |
| Duplicate message number packets     | Number of duplicate message number packets.                                                                                                                                       |
| Duplicate address packets            | Number of duplicate source MAC address packets.                                                                                                                                   |
| Invalid destination address packets  | Number of invalid destination MAC address packets.                                                                                                                                |
| Formatting error packets             | Number of formatting error packets.                                                                                                                                               |
| Old Replayed message number packets  | Number of old replayed message number packets.                                                                                                                                    |

## Sample Output

### show security mka statistics

```
user@host> show security mka statistics
```

```

Received packets:          1525844
Transmitted packets:       1525841
Version mismatch packets:  0
CAK mismatch packets:      0
ICV mismatch packets:      0
Duplicate message identifier packets: 0
Duplicate message number packets: 0
Duplicate address packets:  0
Invalid destination address packets: 0
Formatting error packets:   0
Old Replayed message number packets: 0

```



## PART 24

# Routing Policy and Packet Filtering

- [Overview on page 4693](#)
- [Configuration on page 4773](#)
- [Administration on page 4873](#)
- [Troubleshooting Procedures on page 4897](#)



## CHAPTER 77

# Overview

- [Security Features Overview on page 4693](#)
- [Firewall Filters Overview on page 4695](#)

## Security Features Overview

---

- [Security Features for EX Series Switches Overview on page 4693](#)

## Security Features for EX Series Switches Overview

Juniper Networks Junos operating system (Junos OS) is a network operating system that has been hardened through the separation of control forwarding and services planes, with each function running in protected memory. The control-plane CPU is protected by rate limiting, routing policy, and firewall filters to ensure switch uptime even under severe attack. Access port security features such as dynamic Address Resolution Protocol (ARP) inspection, DHCP snooping, and MAC limiting are controlled through a single Junos OS CLI command.

Juniper Networks EX Series Ethernet Switches provide the following hardware and software security features:

**Console Port**—Allows use of the console port to connect to the Routing Engine through an RJ-45 cable. You then use the command-line interface (CLI) to configure the switch.

**Out-of-Band Management**—A dedicated management Ethernet port on the rear panel allows out-of-band management.

**Software Images**—All Junos OS images are signed by Juniper Networks certificate authority (CA) with public key infrastructure (PKI).

**User Authentication, Authorization, and Accounting (AAA)**—Features include:

- User and group accounts with password encryption and authentication.
- Access privilege levels configurable for login classes and user templates.
- RADIUS authentication, TACACS+ authentication, or both, for authenticating users who attempt to access the switch.
- Auditing of configuration changes through system logging or RADIUS/TACACS+.

**802.1X Authentication**—Provides network access control. Supplicants (hosts) are authenticated when they initially connect to a LAN. Authenticating supplicants before they receive an IP address from a DHCP server prevents unauthorized supplicants from gaining access to the LAN. EX Series switches support Extensible Authentication Protocol (EAP) methods, including EAP-MD5, EAP-TLS, EAP-TTLS, and EAP-PEAP.

**Port Security**—Access port security features include:

- **DHCP snooping**—Filters and blocks ingress DHCP server messages on untrusted ports; builds and maintains an IP-address/MAC-address binding database (called the DHCP snooping database).
- **Dynamic ARP inspection (DAI)**—Prevents ARP spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons.
- **MAC limiting**—Protects against flooding of the Ethernet switching table.
- **MAC move limiting**—Detects MAC movement and MAC spoofing on access ports.
- **Trusted DHCP server**—With a DHCP server on a trusted port, protects against rogue DHCP servers sending leases.
- **IP source guard**—Mitigates the effects of IP address spoofing attacks on the Ethernet LAN. The source IP address in the packet sent from an untrusted access interface is validated against the source MAC address in the DHCP snooping database. The packet is allowed for further processing if the source IP address to source MAC address binding is valid; if the binding is not valid, the packet is discarded.
- **DHCP option 82**—Also known as the DHCP relay agent information option. Helps protect the EX Series switch against attacks such as spoofing (forging) of IP addresses and MAC addresses and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.
- **Unrestricted proxy ARP**—The switch responds to all ARP messages with its own MAC address. Hosts that are connected to the switch's interfaces cannot communicate directly with other hosts. Instead, all communications between hosts go through the switch.
- **Restricted proxy ARP**—The switch does not respond to an ARP request if the physical networks of the source and target of the ARP request are the same. It does not matter whether the destination host has the same IP address as the incoming interface or a different (remote) IP address. An ARP request for a broadcast address elicits no reply.

**Device Security**—Storm control permits the switch to monitor unknown unicast and broadcast traffic and drop packets, or shut down, or temporarily disable the interface when a specified traffic level is exceeded, thus preventing packets from proliferating and degrading the LAN. You can enable storm control on access interfaces or trunk interfaces.

**Firewall Filters**—Allow auditing of various types of security violations, including attempts to access the switch from unauthorized locations. Firewall filters can detect such attempts and create audit log entries when they occur. The filters can also restrict access by limiting

traffic to source and destination MAC addresses, specific protocols, or, in combination with policers, to specified data rates to prevent denial of service (DoS) attacks.

**Policers**—Provide rate-limiting capability to control the amount of traffic that enters an interface, which acts to counter DoS attacks.

**Encryption Standards**—Supported standards include:

- 128-, 192-, and 256-bit Advanced Encryption Standard (AES)
- 56-bit Data Encryption Standard (DES) and 168-bit 3DES

#### Related Documentation

- [802.1X for EX Series Switches Overview on page 1821](#)
- [Firewall Filters for EX Series Switches Overview on page 4696](#)
- [Understanding Port Security on page 4497](#)
- [Understanding Proxy ARP on EX Series Switches on page 2278](#)
- [Understanding Storm Control on EX Series Switches](#)
- [Understanding the Use of Policers in Firewall Filters on page 4767](#)
- [Understanding Centralized Network Access Control and EX Series Switches](#)

## Firewall Filters Overview

- [Firewall Filters for EX Series Switches Overview on page 4696](#)
- [Understanding Planning of Firewall Filters on page 4699](#)
- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on EX Series Switches on page 4702](#)
- [Understanding How Firewall Filters Control Packet Flows on page 4703](#)
- [Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704](#)
- [Support for Match Conditions and Actions for Loopback Firewall Filters on Switches on page 4715](#)
- [Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches on page 4718](#)
- [Understanding How Firewall Filters Are Evaluated on page 4760](#)
- [Understanding Firewall Filter Match Conditions on page 4762](#)
- [Understanding How Firewall Filters Test a Packet's Protocol on page 4766](#)
- [Understanding the Use of Policers in Firewall Filters on page 4767](#)
- [Understanding Filter-Based Forwarding for EX Series Switches on page 4770](#)
- [Understanding Tricolor Marking Architecture on page 4770](#)

## Firewall Filters for EX Series Switches Overview

Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on a Juniper Networks EX Series Ethernet Switch from a source address to a destination address. You configure firewall filters to determine whether to permit, deny, or forward traffic before it enters or exits a port, VLAN, or Layer 3 (routed) interface to which the firewall filter is applied. To apply a firewall filter, you must first configure the filter and then apply it to an port, VLAN, or Layer 3 interface.

You can apply firewall filters to network interfaces, aggregated Ethernet interfaces (also known as link aggregation groups (LAGs)), loopback interfaces, management interfaces, virtual management Ethernet interfaces (VMEs), and routed VLAN interfaces (RVIs). For information on EX Series switches that support a firewall filter on these interfaces, see *EX Series Switch Software Features Overview*.

An *ingress* firewall filter is a filter that is applied to packets that are entering a network. An *egress* firewall filter is a filter that is applied to packets that are exiting a network. You can configure firewall filters to subject packets to filtering, class-of-service (CoS) marking (grouping similar types of traffic together, and treating each type of traffic as a class with its own level of service priority), and traffic policing (controlling the maximum rate of traffic sent or received on an interface).

This topic describes:

- [Firewall Filter Types on page 4696](#)
- [Firewall Filter Components on page 4697](#)
- [Firewall Filter Processing on page 4698](#)

---

### Firewall Filter Types

The following firewall filter types are supported for EX Series switches:

- **Port (Layer 2) firewall filter**—Port firewall filters apply to Layer 2 switch ports. You can apply port firewall filters in both ingress and egress directions on a physical port.
- **VLAN firewall filter**—VLAN firewall filters provide access control for packets that enter a VLAN, are bridged within a VLAN, or leave a VLAN. You can apply VLAN firewall filters in both ingress and egress directions on a VLAN. VLAN firewall filters are applied to all packets that are forwarded to or forwarded from the VLAN.
- **Router (Layer 3) firewall filter**—You can apply a router firewall filter in both ingress and egress directions on Layer 3 (routed) interfaces and routed VLAN interfaces (RVIs). You can apply a router firewall filter in the ingress direction on the loopback interface (**lo0**) also. Firewall filters configured on loopback interfaces are applied only to packets that are sent to the Routing Engine CPU for further processing.

You can apply port, VLAN, or router firewall filters to *both* IPv4 and IPv6 traffic on these switches:

- EX2200 switch
- EX3300 switch



- EX3200 switch
- EX4200 switch
- EX4300 switch
- EX4500 switch
- EX6200 switch
- EX8200 switch

For information on firewall filters supported on different switches, see [“Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches”](#) on page 4718.

### Firewall Filter Components

In a firewall filter, you first define the family address type (**ethernet-switching**, **inet**, or **inet6**), and then you define one or more terms that specify the filtering criteria (specified as terms with match conditions) and the action (specified as actions or action modifiers) to take if a match occurs.

The maximum number of terms allowed per firewall filter for EX Series switches is:

- 512 for EX2200 switches
- 1436 for EX3300 switches



**NOTE:** On EX3300 switches, if you add and delete filters with a large number of terms (on the order of 1000 or more) in the same commit operation, not all the filters are installed. You must add filters in one commit operation, and delete filters in a separate commit operation.

- 7,042 for EX3200 and EX4200 switches—as allocated by the dynamic allocation of ternary content addressable memory (TCAM) for firewall filters.
- On EX4300 switches, the following maximum number of terms are supported for ingress and egress traffic, for firewall filters configured on a port, VLAN and Layer 3 interface:
  - For ingress traffic:
    - 3500 terms for firewall filters configured on a port
    - 3500 terms for firewall filters configured on a VLAN
    - 7000 terms for firewall filters configured on Layer 3 interfaces for IPv4 traffic
    - 3500 terms for firewall filters configured on Layer 3 interfaces for IPv6 traffic
  - For egress traffic:
    - 512 terms for firewall filters configured on a port
    - 256 terms for firewall filters configured on a VLAN

- 512 terms for firewall filters configured on Layer 3 interfaces for IPv4 traffic
- 512 terms for firewall filters configured on Layer 3 interfaces for IPv6 traffic



**NOTE:** You can configure the maximum number of terms only when you configure one type of firewall filter (port, VLAN, or router (Layer 3) firewall filter) on the switch, and when storm control is not enabled on any interface in the switch.

- 1200 for EX4500 and EX4550 switches
- 1400 for EX6200 switches
- 32,768 for EX8200 switches



**NOTE:** The on-demand dynamic allocation of the shared space TCAM in EX8200 switches is achieved by assigning free space blocks to firewall filters. Firewall filters are categorized into two different pools. Port and VLAN filters are pooled together (the memory threshold for this pool is 22K) while router firewall filters are pooled separately (the threshold for this pool is 32K). The assignment happens based on the filter pool type. Free space blocks can be shared only among the firewall filters belonging to the same filter pool type. An error message is generated when you try to configure a firewall filter beyond the TCAM threshold.

Each term consists of the following components:

- Match conditions—Specify the values or fields that the packet must contain. You can define various match conditions, including the IP source address field, IP destination address field, Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port field, IP protocol field, Internet Control Message Protocol (ICMP) packet type, TCP flags, and interfaces.
- Action—Specifies what to do if a packet matches the match conditions. Possible actions are to accept or discard the packet or to send the packet to a specific virtual routing interface. In addition, packets can be counted to collect statistical information. If no action is specified for a term, the default action is to accept the packet.
- Action modifier—Specifies one or more actions for the switch if a packet matches the match conditions. You can specify action modifiers such as count, mirror, rate limit, and classify packets.

---

### Firewall Filter Processing

The order of the terms within a firewall filter configuration is important. Packets are tested against each term in the order in which the terms are listed in the firewall filter configuration. For information on how firewall filters process packets, see [“Understanding How Firewall Filters Are Evaluated” on page 4760](#).

- Related Documentation**
- [Understanding Planning of Firewall Filters on page 4699](#)
  - [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on EX Series Switches on page 4702](#)
  - [Understanding How Firewall Filters Are Evaluated on page 4760](#)
  - [Understanding Firewall Filter Match Conditions on page 4762](#)
  - [Understanding the Use of Policers in Firewall Filters on page 4767](#)
  - [Understanding Filter-Based Forwarding for EX Series Switches on page 4770](#)
  - [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
  - [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 4795](#)

## Understanding Planning of Firewall Filters

Before you create a firewall filter and apply it to an interface, determine what you want the firewall filter to accomplish and how to use its match conditions and actions to achieve your goals. You must understand how packets are matched to match conditions, the default and configured actions of the firewall filter, and proper placement of the firewall filter.

You can configure and apply no more than one firewall filter per port, VLAN, or router interface, per direction. The following limits apply for the number of firewall filter terms allowed per filter on various switch models:

- On EX2200 switches, the number of terms per filter cannot exceed 512.
- On EX3300 switches, the number of terms per filter cannot exceed 1436.
- On EX3200 and EX4200 switches, the number of terms per filter cannot exceed 7042.
- On EX4300 switches, the following maximum number of terms are supported for ingress and egress traffic, for firewall filters configured on a port, VLAN and Layer 3 interface:
  - For ingress traffic:
    - 3500 terms for firewall filters configured on a port
    - 3500 terms for firewall filters configured on a VLAN
    - 7000 terms for firewall filters configured on Layer 3 interfaces for IPv4 traffic
    - 3500 terms for firewall filters configured on Layer 3 interfaces for IPv6 traffic
  - For egress traffic:
    - 512 terms for firewall filters configured on a port
    - 256 terms for firewall filters configured on a VLAN

- 512 terms for firewall filters configured on Layer 3 interfaces for IPv4 traffic
- 512 terms for firewall filters configured on Layer 3 interfaces for IPv6 traffic



**NOTE:** You can configure the maximum number of terms only when you configure one type of firewall filter (port, VLAN, or router (Layer 3) firewall filter) on the switch, and when storm control is not enabled on any interface in the switch.

- On EX4500 and EX4550 switches, the number of terms per filter cannot exceed 1200.
- On EX6200 switches, the number of terms per filter cannot exceed 1400.
- On EX8200 switches, the number of terms per filter cannot exceed 32,768.

In addition, try to be conservative in the number of terms (rules) that you include in each firewall filter because a large number of terms requires longer processing time during a commit and also can make firewall filter testing and troubleshooting more difficult. Similarly, applying firewall filters across many switch and router interfaces can make testing and troubleshooting the rules of those filters difficult.

Before you configure and apply firewall filters, answer the following questions for each of those firewall filters:

1. What is the purpose of the firewall filter?

For example, you can use a firewall filter to limit traffic to source and destination MAC addresses, specific protocols, or certain data rates or to prevent denial of service (DoS) attacks.

2. What are the appropriate match conditions?

- a. Determine the packet header fields that the packet must contain for a match.  
Possible fields include:

- Layer 2 header fields—Source and destination MAC addresses, dot1q tag, Ethernet type, and VLAN
- Layer 3 header fields—Source and destination IP addresses, protocols, and IP options (IP precedence, IP fragmentation flags, TTL type)
- TCP header fields—Source and destination ports and flags
- ICMP header fields—Packet type and code

- b. Determine the port, VLAN, or router interface on which the packet was received.

3. What are the appropriate actions to take if a match occurs?

Possible actions to take if a match occurs are accept, discard, and forward to a routing instance.

4. What additional action modifiers might be required?

Determine whether additional actions are required if a packet matches a match condition; for example, you can specify an action modifier to count, analyze, or police packets.

5. On what interface should the firewall filter be applied?

Start with the following basic guidelines:

- If all the packets entering a port need to be exposed to filtering, then use port firewall filters.
- If all the packets that are bridged need filtering, then use VLAN firewall filters.
- If all the packets that are routed need filtering, then use router firewall filters.

Before you choose the interface on which to apply a firewall filter, understand how that placement can impact traffic flow to other interfaces. In general, apply a firewall filter that filters on source and destination IP addresses, IP protocols, or protocol information—such as ICMP message types, and TCP and UDP port numbers—nearest to the source devices. However, typically apply a firewall filter that filters only on a source IP address nearest to the destination devices. When applied too close to the source device, a firewall filter that filters only on a source IP address could potentially prevent that source device from accessing other services that are available on the network.



**NOTE:** Egress firewall filters do not affect the flow of locally generated control packets from the Routing Engine.

6. In which direction should the firewall filter be applied?

You can apply firewall filters to ports on the switch to filter packets that are entering a port. You can apply firewall filters to VLANs, and Layer 3 (routed) interfaces to filter packets that are entering or exiting a VLAN or routed interface. Typically, you configure different sets of actions for traffic entering an interface than you configure for traffic exiting an interface.

**Related Documentation**

- [Firewall Filters for EX Series Switches Overview on page 4696](#)
- [Understanding the Use of Policers in Firewall Filters on page 4767](#)
- [Understanding How Firewall Filters Are Evaluated on page 4760](#)
- [Understanding Filter-Based Forwarding for EX Series Switches on page 4770](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 4795](#)

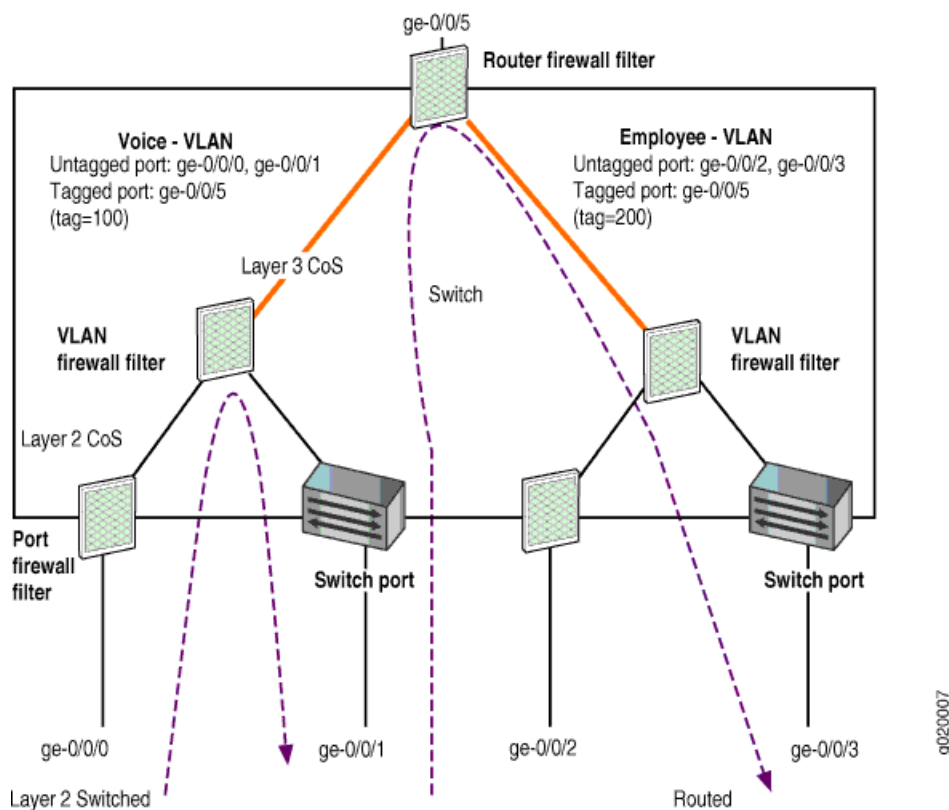
## Understanding Firewall Filter Processing Points for Bridged and Routed Packets on EX Series Switches

Juniper Networks EX Series Ethernet Switches are multilayered switches that provide Layer 2 switching and Layer 3 routing. You apply firewall filters at multiple processing points in the packet forwarding path on EX Series switches. At each processing point, the action to be taken on a packet is determined based on the results of the lookup in the switch's forwarding table. A table lookup determines which exit port on the switch to use to forward the packet.

For both bridged unicast packets and routed unicast packets, firewall filters are evaluated and applied hierarchically. First, a packet is checked against the port firewall filter, if present. If the packet is permitted, it is then checked against the VLAN firewall filter, if present. If the packet is permitted, it is then checked against the router firewall filter, if present. The packet must be permitted by the router firewall filter before it is processed.

Figure 66 on page 4702 shows the various firewall filter processing points in the packet forwarding path in a multilayered switching platform.

Figure 66: Firewall Filter Processing Points in the Packet Forwarding Path



For a multicast packet that results in replications, an egress firewall filter is applied to each copy of the packet based on its corresponding egress VLAN.

For Layer 2 (bridged) unicast packets, the following firewall filter processing points apply:

- Ingress port firewall filter
- Ingress VLAN firewall filter
- Egress port firewall filter
- Egress VLAN firewall filter

For Layer 3 (routed and multilayer-switched) unicast packets, the following firewall filter processing points apply:

- Ingress port firewall filter
- Ingress VLAN firewall filter (Layer 2 CoS)
- Ingress router firewall filter (Layer 3 CoS)
- Egress router firewall filter
- Egress VLAN firewall filter

#### Related Documentation

- [Firewall Filters for EX Series Switches Overview on page 4696](#)
- [Understanding How Firewall Filters Control Packet Flows on page 4703](#)
- [Understanding Bridging and VLANs on EX Series Switches on page 2245](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)

## Understanding How Firewall Filters Control Packet Flows

Juniper Networks EX Series Ethernet Switches support firewall filters that allow you to control flows of data packets and local packets. *Data packets* are chunks of data that transit the switch as they are forwarded from a source to a destination. *Local packets* are chunks of data that are destined for or sent by the switch. Local packets usually contain routing protocol data, data for IP services such as Telnet or SSH, and data for administrative protocols such as the Internet Control Message Protocol (ICMP).

You create firewall filters to protect your switch from excessive traffic transiting the switch to a network destination or destined for the Routing Engine on the switch. Firewall filters that control local packets can also protect your switch from external incidents such as denial-of-service (DoS) attacks.

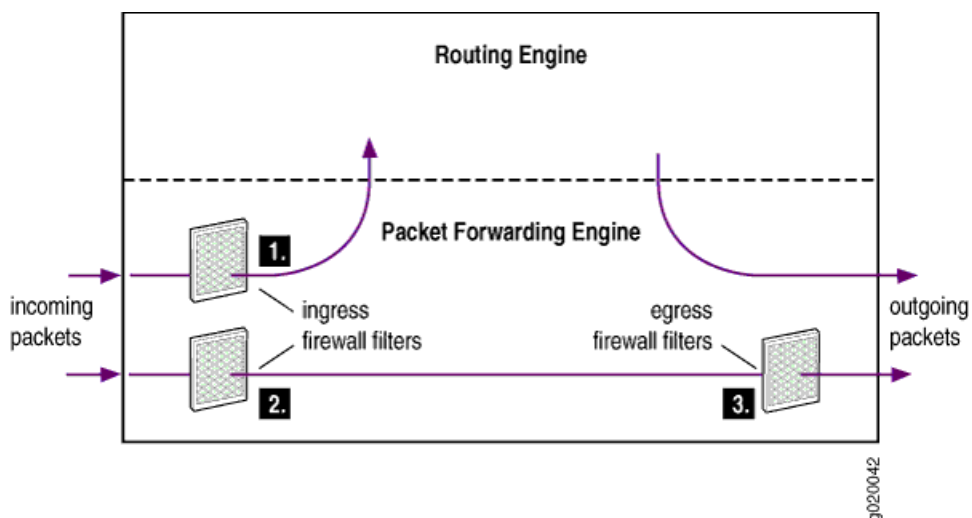
Firewall filters affect packet flows entering in to or exiting from the switch's interfaces:

- Ingress firewall filters affect the flow of data packets that are received by the switch's interfaces. The Packet Forwarding Engine handles this flow. When a switch receives a data packet on an interface, the switch determines where to forward the packet by looking in the forwarding table for the best route (Layer 2 switching, Layer 3 routing) to a destination. Data packets are forwarded to their destination through an outgoing interface. Locally destined packets are forwarded to the Routing Engine.

- Egress firewall filters affect the flow of data packets that are transmitted from the switch's interfaces but do not affect the flow of locally generated control packets from the Routing Engine. The Packet Forwarding Engine handles the flow of data packets that are transmitted from the switch, and egress firewall filters are applied here. The Packet Forwarding Engine also handles the flow of control packets from the Routing Engine.

Figure 67 on page 4704 illustrates the application of ingress and egress firewall filters to control the flow of packets through the switch.

Figure 67: Application of Firewall Filters to Control Packet Flow



- Ingress firewall filter applied to control locally destined packets that are received on the switch's interfaces and are destined for the Routing Engine.
- Ingress firewall filter applied to control incoming packets on the switch's interfaces.
- Egress firewall filter applied to control packets that are transiting the switch's interfaces.

#### Related Documentation

- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on EX Series Switches on page 4702](#)
- [Understanding How Firewall Filters Are Evaluated on page 4760](#)

## Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches

When you define a firewall filter for an EX Series switch, you define filtering criteria (*terms*, with *match conditions*) for the packets and an *action* (and, optionally, an *action modifier*) for the switch to take if the packets match the filtering criteria. You can define a firewall filter to monitor IPv4, IPv6, or non-IP traffic.

This topic describes in detail the various match conditions, actions, and action modifiers that you can define in a firewall filter. For information about support for match conditions



on various EX Series switches, see [“Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches”](#) on page 4718.

This topic describes:

- [Firewall Filter Elements on page 4705](#)
- [Match Conditions Supported on Switches on page 4705](#)
- [Actions for Firewall Filters on page 4712](#)
- [Action Modifiers for Firewall Filters on page 4713](#)

### Firewall Filter Elements

A firewall filter configuration contains a term, a match condition, an action, and, optionally, an action modifier. [Table 511 on page 4705](#) describes each element in a firewall filter configuration.

**Table 511: Elements of a Firewall Filter Configuration**

| Element Name    | Description                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Term            | Defines the filtering criteria for the packets. Each term in the firewall filter consists of match conditions and an action. You can define a single term or multiple terms in the firewall filter. If you define multiple terms, each term must have a unique name.                                                                                                                    |
| Match condition | Consists of a string (called a <i>match statement</i> ) that defines the match condition. Match conditions are the values or fields that a packet must contain. You can define a single match condition or multiple match conditions for a term. You can also opt not to define a match condition. If no match conditions are specified for a term, all packets are matched by default. |
| Action          | Specifies the action that the switch takes if a packet matches all the criteria specified in the match conditions.                                                                                                                                                                                                                                                                      |
| Action modifier | Specifies one or more actions that the switch takes if a packet matches the match conditions for the specific term.                                                                                                                                                                                                                                                                     |

### Match Conditions Supported on Switches

Based on the type of traffic that you want to monitor, you can configure a firewall filter to monitor IPv4, IPv6, or non-IP traffic. When you configure a firewall filter to monitor a particular type of traffic, ensure that you specify match conditions that are supported for that type of traffic. For information about match conditions supported for a specific type of traffic and switches on which they are supported, see [“Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches”](#) on page 4718.

[Table 512 on page 4705](#) describes all the match conditions that are supported for firewall filters on EX Series Switches.

**Table 512: Firewall Filter Match Conditions Supported on EX Series Switches**

| Match Condition                             | Description                                                                       |
|---------------------------------------------|-----------------------------------------------------------------------------------|
| <code>destination-address ip-address</code> | IP destination address field, which is the address of the final destination node. |

Table 512: Firewall Filter Match Conditions Supported on EX Series Switches (*continued*)

| Match Condition                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ip-destination-address</b> <i>ip-address</i>   | IP destination address field, which is the address of the final destination node.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>ip6-destination-address</b> <i>ip-address</i>  | IP destination address field, which is the address of the final destination node.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>destination-mac-address</b> <i>mac-address</i> | <p>Destination media access control (MAC) address of the packet.</p> <p>You can define a destination MAC address with a prefix, such as <b>destination-mac-address 00:01:02:03:04:05/24</b>. If no prefix is specified, the default value 48 is used.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>destination-port</b> <i>number</i>             | <p>TCP or UDP destination port field. Typically, you specify this match condition in conjunction with the <b>protocol</b> or <b>ip-protocol</b> match condition to determine which protocol is used on the port. For <i>number</i>, you can specify one of the following text synonyms (the port numbers are also listed):</p> <p><b>afs</b> (1483), <b>bgp</b> (179), <b>biff</b> (512), <b>bootpc</b> (68), <b>bootps</b> (67), <b>cmd</b> (514), <b>cvspserver</b> (2401), <b>dhcp</b> (67), <b>domain</b> (53), <b>eklogin</b> (2105), <b>ekshell</b> (2106), <b>exec</b> (512), <b>finger</b> (79), <b>ftp</b> (21), <b>ftp-data</b> (20), <b>http</b> (80), <b>https</b> (443), <b>ident</b> (113), <b>imap</b> (143), <b>kerberos-sec</b> (88), <b>klogin</b> (543), <b>kpasswd</b> (761), <b>krb-prop</b> (754), <b>krbupdate</b> (760), <b>kshell</b> (544), <b>ldap</b> (389), <b>login</b> (513), <b>mobileip-agent</b> (434), <b>mobilip-mn</b> (435), <b>msdp</b> (639), <b>netbios-dgm</b> (138), <b>netbios-ns</b> (137), <b>netbios-ssn</b> (139), <b>nfsd</b> (2049), <b>nntp</b> (119), <b>ntalk</b> (518), <b>ntp</b> (123), <b>pop3</b> (110), <b>pptp</b> (1723), <b>printer</b> (515), <b>radacct</b> (1813), <b>radius</b> (1812), <b>rip</b> (520), <b>rkinit</b> (2108), <b>smtp</b> (25), <b>snmp</b> (161), <b>snmptrap</b> (162), <b>snpp</b> (444), <b>socks</b> (1080), <b>ssh</b> (22), <b>sunrpc</b> (111), <b>syslog</b> (514), <b>tacacs-ds</b> (65), <b>talk</b> (517), <b>telnet</b> (23), <b>tftp</b> (69), <b>timed</b> (525), <b>who</b> (513), <b>xdmcp</b> (177), <b>zephyr-clt</b> (2103), <b>zephyr-hm</b> (2104)</p> |
| <b>destination-prefix-list</b> <i>prefix-list</i> | <p>IP destination prefix list field.</p> <p>You can define a list of IP address prefixes under a prefix-list alias for frequent use. You define this match condition at the <b>[edit policy-options]</b> hierarchy level.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>dot1q-tag</b> <i>number</i>                    | <p>The tag field in the Ethernet header. The tag values range from 1 through 4095. The <b>dot1q-tag</b> match condition and the <b>vlan</b> match condition are mutually exclusive.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>user-vlan-id</b> <i>number</i>                 | <p>The tag field in the Ethernet header. The tag values range from 1 through 4095. The <b>user-vlan-id</b> match condition and the <b>learn-vlan-id</b> match condition are mutually exclusive.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 512: Firewall Filter Match Conditions Supported on EX Series Switches (*continued*)

| Match Condition                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>dot1q-user-priority <i>number</i></b>   | <p>User-priority field of the tagged Ethernet packet. User-priority values can range from 0 through 7.</p> <p>For <i>number</i>, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>background (1)</b>—Background</li> <li>• <b>best-effort (0)</b>—Best effort</li> <li>• <b>controlled-load (4)</b>—Controlled load</li> <li>• <b>excellent-load (3)</b>—Excellent load</li> <li>• <b>network-control (7)</b>—Network control reserved traffic</li> <li>• <b>standard (2)</b>—Standard or spare</li> <li>• <b>video (5)</b>—Video</li> <li>• <b>voice (6)</b>—Voice</li> </ul>                                                                                                                                                                                                   |
| <b>user-vlan-1p-priority <i>number</i></b> | <p>User-priority field of the tagged Ethernet packet. User-priority values can range from 0 through 7.</p> <p>For <i>number</i>, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>background (1)</b>—Background</li> <li>• <b>best-effort (0)</b>—Best effort</li> <li>• <b>controlled-load (4)</b>—Controlled load</li> <li>• <b>excellent-load (3)</b>—Excellent load</li> <li>• <b>network-control (7)</b>—Network control reserved traffic</li> <li>• <b>standard (2)</b>—Standard or spare</li> <li>• <b>video (5)</b>—Video</li> <li>• <b>voice (6)</b>—Voice</li> </ul>                                                                                                                                                                                                   |
| <b>dscp <i>number</i></b>                  | <p>Specifies the Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p> <p>For <i>number</i>, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>ef (46)</b>—as defined in <a href="#">RFC 2598</a>, <i>An Expedited Forwarding PHB</i>.</li> <li>• <b>af11 (10), af12 (12), af13 (14), af21 (18), af22 (20), af23 (22), af31 (26), af32 (28), af33 (30), af41 (34), af42 (36), af43 (38)</b></li> </ul> <p>These four classes, with three drop precedences in each class, are defined for 12 code points in <a href="#">RFC 2597</a>, <i>Assured Forwarding PHB Group</i>.</p> |

Table 512: Firewall Filter Match Conditions Supported on EX Series Switches (*continued*)

| Match Condition                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ether-type</b> <i>value</i>              | <p>Ethernet type field of a packet. The <i>value</i> specifies what protocol is being transported in the Ethernet frame. For <i>value</i>, you can specify one of the following text synonyms:</p> <ul style="list-style-type: none"> <li><b>aarp</b>—EtherType value AARP (0x80F3)</li> <li><b>appletalk</b>—EtherType value AppleTalk (0x809B)</li> <li><b>arp</b>—EtherType value ARP (0x0806)</li> <li><b>ipv4</b>—EtherType value IPv4 (0x0800)</li> <li><b>ipv6</b>—EtherType value IPv6 (0x08DD)</li> <li><b>mpls multicast</b>—EtherType value MPLS multicast (0x8848)</li> <li><b>mpls unicast</b>—EtherType value MPLS unicast (0x8847)</li> <li><b>oam</b>—EtherType value OAM (0x88A8)</li> <li><b>ppp</b>—EtherType value PPP (0x880B)</li> <li><b>pppoe-discovery</b>—EtherType value PPPoE Discovery Stage (0x8863)</li> <li><b>pppoe-session</b>—EtherType value PPPoE Session Stage (0x8864)</li> <li><b>sna</b>—EtherType value SNA (0x80D5)</li> </ul> <p><b>NOTE:</b> The following match conditions are not supported when <b>ether-type</b> is set to <b>ipv6</b>:</p> <ul style="list-style-type: none"> <li><b>dscp</b></li> <li><b>fragment-flags</b></li> <li><b>is-fragment</b></li> <li><b>precedence</b> or <b>ip-precedence</b></li> <li><b>protocol</b> or <b>ip-protocol</b></li> </ul> |
| <b>fragment-flags</b> <i>fragment-flags</i> | <p>IP fragmentation flags, specified in symbolic or hexadecimal formats. You can specify one of the following options:</p> <ul style="list-style-type: none"> <li><b>dont-fragment</b> (0x4000)</li> <li><b>more-fragments</b> (0x2000)</li> <li><b>reserved</b> (0x8000)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Table 512: Firewall Filter Match Conditions Supported on EX Series Switches (*continued*)

| Match Condition                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>icmp-code</b> <i>number</i>         | <p>ICMP code field. This value or option provides more specific information than <b>icmp-type</b>. Because the value's meaning depends upon the associated <b>icmp-type</b>, you must specify <b>icmp-type</b> along with <b>icmp-code</b>. For <i>number</i>, you can specify one of the following text synonyms (the field values are also listed). The options are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> <li>• <b>parameter-problem</b>—<b>ip-header-bad</b> (0), <b>required-option-missing</b> (1)</li> <li>• <b>redirect</b>—<b>redirect-for-host</b> (1), <b>redirect-for-network</b> (0), <b>redirect-for-tos-and-host</b> (3), <b>redirect-for-tos-and-net</b> (2)</li> <li>• <b>time-exceeded</b>—<b>ttl-eq-zero-during-reassembly</b> (1), <b>ttl-eq-zero-during-transit</b> (0)</li> <li>• <b>unreachable</b>—<b>communication-prohibited-by-filtering</b> (13), <b>destination-host-prohibited</b> (10), <b>destination-host-unknown</b> (7), <b>destination-network-prohibited</b> (9), <b>destination-network-unknown</b> (6), <b>fragmentation-needed</b> (4), <b>host-precedence-violation</b> (14), <b>host-unreachable</b> (1), <b>host-unreachable-for-TOS</b> (12), <b>network-unreachable</b> (0), <b>network-unreachable-for-TOS</b> (11), <b>port-unreachable</b> (3), <b>precedence-cutoff-in-effect</b> (15), <b>protocol-unreachable</b> (2), <b>source-host-isolated</b> (8), <b>source-route-failed</b> (5)</li> </ul> |
| <b>icmp-type</b> <i>number</i>         | <p>ICMP packet type field. Typically, you specify this match condition in conjunction with the <b>protocol</b> or <b>ip-protocol</b> match condition to determine which protocol is being used on the port. For <i>number</i>, you can specify one of the following text synonyms (the field values are also listed):</p> <p><b>echo-reply</b> (0), <b>echo-request</b> (8), <b>info-reply</b> (16), <b>info-request</b> (15), <b>mask-request</b> (17), <b>mask-reply</b> (18), <b>parameter-problem</b> (12), <b>redirect</b> (5), <b>router-advertisement</b> (9), <b>router-solicit</b> (10), <b>source-quench</b> (4), <b>time-exceeded</b> (11), <b>timestamp</b> (13), <b>timestamp-reply</b> (14), <b>unreachable</b> (3)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>interface</b> <i>interface-name</i> | <p>Interface on which the packet is received. You can specify the wildcard character (*) as part of an interface name.</p> <p><b>NOTE:</b> The <b>interface</b> match condition is not supported for egress traffic on an EX8200 Virtual Chassis.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>ip-options</b>                      | Presence of the options field in the IP header.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 512: Firewall Filter Match Conditions Supported on EX Series Switches (*continued*)

| Match Condition                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ip-version</b> <i>version match_condition(s)</i> | <p>Version of the IP protocol for port and VLAN firewall filters. The value for <i>version</i> can be <b>ipv4</b> or <b>ipv6</b>.</p> <p>For <i>match_condition(s)</i>, you can specify one or more of the following match conditions:</p> <ul style="list-style-type: none"> <li>• <b>destination-address</b>, <b>ip-destination-address</b>, or <b>ip6-destination-address</b></li> <li>• <b>destination-port</b></li> <li>• <b>destination-prefix-list</b></li> <li>• <b>dscp</b></li> <li>• <b>fragment-flags</b></li> <li>• <b>icmp-code</b></li> <li>• <b>icmp-type</b></li> <li>• <b>is-fragment</b></li> <li>• <b>precedence</b> or <b>ip-precedence</b></li> <li>• <b>protocol</b> or <b>ip-protocol</b></li> <li>• <b>source-address</b> or <b>ip-source-address</b></li> <li>• <b>source-port</b></li> <li>• <b>source-prefix-list</b></li> <li>• <b>tcp-established</b></li> <li>• <b>tcp-flags</b></li> <li>• <b>tcp-initial</b></li> </ul> |
| <b>is-fragment</b>                                  | <p>If the packet is a trailing fragment, this match condition does not match the first fragment of a fragmented packet. Use two terms to match both first and trailing fragments.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>l2-encap-type llc-non-snap</b>                   | <p>Match on logical link control (LLC) layer packets for non-Subnet Access Protocol (SNAP) Ethernet Encapsulation type.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>next-header bytes</b>                            | <p>8-bit protocol field that identifies the type of header immediately following the IPv6 header. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <p><b>ah (51), dstops (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmp6 (1), igmp (2), ipip (4), ipv6 (41), no-next-header (59), ospf (89), pim (103), routing (43), rsvp (46), sctp (132), tcp (6), udp (17), vrrp (112)</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>packet-length bytes</b>                          | <p>Length of the received packet, in bytes.</p> <p>The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>precedence precedence</b>                        | <p>IP precedence. For <i>precedence</i>, you can specify one of the following text synonyms (the field values are also listed):</p> <p><b>critical-ecp (5), flash (3), flash-override (4), immediate (2), internet-control (6), net-control (7), priority (1), routine (0)</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 512: Firewall Filter Match Conditions Supported on EX Series Switches (*continued*)

| Match Condition                                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ip-precedence</b> <i>precedence</i>                              | <p>IP precedence. For <i>precedence</i>, you can specify one of the following text synonyms (the field values are also listed):</p> <p><b>critical-ecp (5), flash (3), flash-override (4), immediate (2), internet-control (6), net-control (7), priority (1), routine (0)</b></p>                                                                                                                                                           |
| <b>protocol</b> <i>list of protocol</i>                             | <p>IPv4 protocol value. For <i>protocols</i>, you can specify one of the following text synonyms:</p> <p><b>egp (8), esp (50), gre (47), icmp (1), igmp (2), ipip (4), ospf (89), pim (103), rsvp (46), tcp (6), udp (17)</b></p>                                                                                                                                                                                                            |
| <b>ip-protocol</b> <i>list of protocol</i>                          | <p>IPv4 protocol value. For <i>protocols</i>, you can specify one of the following text synonyms:</p> <p><b>egp (8), esp (50), gre (47), icmp (1), igmp (2), ipip (4), ospf (89), pim (103), rsvp (46), tcp (6), udp (17)</b></p>                                                                                                                                                                                                            |
| <b>source-address</b> <i>ip-address</i>                             | <p>IP source address field, which is the address of the source node sending the packet. For IPv6, the source-address field is 128 bits in length. The filter description syntax supports the text representations for IPv6 addresses that are described in <a href="#">RFC 2373</a>, <i>IP Version 6 Addressing Architecture</i>.</p>                                                                                                        |
| <b>ip-source-address</b> ( <i>ip-address</i>   <i>ip6-address</i> ) | <p>IP source address field, which is the address of the source node sending the packet. You can specify either an IPv4 address (<b>ip-address</b>) or an IPv6 address (<b>ip6-address</b>). For IPv6, the ip-source-address field is 128 bits in length. The filter description syntax supports the text representations for IPv6 addresses that are described in <a href="#">RFC 2373</a>, <i>IP Version 6 Addressing Architecture</i>.</p> |
| <b>source-mac-address</b> <i>mac-address</i>                        | <p>Source MAC address.</p> <p>You can define a source MAC address with a prefix, such as <b>source-mac-address 00:01:02:03:04:05/24</b>. If no prefix is specified, the default value 48 is used.</p>                                                                                                                                                                                                                                        |
| <b>source-port</b> <i>number</i>                                    | <p>TCP or UDP <b>source-port</b> field. Typically, you specify this match in conjunction with the <b>protocol</b> or <b>ip-protocol</b> match condition to determine which protocol is being used on the port. For <i>number</i>, you can specify one of the text synonyms listed under <b>destination-port</b>.</p>                                                                                                                         |
| <b>source-prefix-list</b> <i>prefix-list</i>                        | <p>IP source prefix list field.</p> <p>You can define a list of IP address prefixes under a prefix-list alias for frequent use. You define this match condition at the <b>[edit policy-options]</b> hierarchy level.</p>                                                                                                                                                                                                                     |
| <b>tcp-established</b>                                              | <p>TCP packets of an established TCP connection. This condition matches packets other than the first packet of a connection. <b>tcp-established</b> is a synonym for the bit names "(ack   rst)".</p> <p><b>tcp-established</b> does not implicitly check whether the protocol is TCP. To do so, specify the <b>next-header tcp</b> match condition.</p>                                                                                     |

Table 512: Firewall Filter Match Conditions Supported on EX Series Switches (*continued*)

| Match Condition                                            | Description                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>tcp-flags</b> ( <i>flags tcp-initial</i> )              | <p>One or more TCP flags:</p> <ul style="list-style-type: none"> <li>bit-name—<b>fin</b>, <b>syn</b>, <b>rst</b>, <b>push</b>, <b>ack</b>, <b>urgent</b></li> <li>logical operators—<b>&amp;</b> (logical AND), <b> </b> (logical OR), <b>!</b> (negation)</li> <li>numerical value—0x01 through 0x20</li> <li>text synonym—<b>tcp-initial</b></li> </ul> <p>To specify multiple flags, use logical operators.</p> |
| <b>tcp-initial</b>                                         | <p>Matches the first TCP packet of a connection. <b>tcp-initial</b> is a synonym for the bit names "<b>syn&amp;!ack</b>".</p> <p><b>tcp-initial</b> does not implicitly check whether the protocol is TCP. To do so, specify the <b>protocol tcp</b> or <b>ip-protocol tcp</b> match condition.</p>                                                                                                                |
| <b>traffic-class</b> <i>number</i>                         | Specifies the DSCP code point for a packet.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>ttl</b> <i>value</i>                                    | TTL type to match. The value ranges from 1 through 255.                                                                                                                                                                                                                                                                                                                                                            |
| <b>vlan</b> ( <i>vlan-name</i>   <i>vlan-id</i> )          | The VLAN that is associated with the packet. For <i>vlan-id</i> , you can specify either the VLAN ID or a VLAN range. The <b>vlan</b> match condition and the <b>dot1q-tag</b> match condition are mutually exclusive.                                                                                                                                                                                             |
| <b>learn-vlan-id</b> ( <i>vlan-name</i>   <i>vlan-id</i> ) | The VLAN that is associated with the packet. For <i>vlan-id</i> , you can specify either the VLAN ID or a VLAN range. The <b>vlan</b> match condition and the <b>user-vlan-id</b> match condition are mutually exclusive.                                                                                                                                                                                          |

### Actions for Firewall Filters

You can define an action for the switch to take if a packet matches the filtering criteria defined in a match condition. [Table 513 on page 4712](#) describes the actions supported in a firewall filter configuration.

Table 513: Actions for Firewall Filters

| Action         | Description                                                                                    |
|----------------|------------------------------------------------------------------------------------------------|
| <b>accept</b>  | Accept a packet.                                                                               |
| <b>discard</b> | Discard a packet silently without sending an Internet Control Message Protocol (ICMP) message. |



Table 513: Actions for Firewall Filters (*continued*)

| Action                                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>reject message-type</code>                    | <p>Discard a packet, and send the ICMPv4 message (type 3) <b>destination unreachable</b>. You can log the rejected packets if you configure the <b>syslog</b> action modifier.</p> <p>You can specify one of the following message codes: <b>administratively-prohibited (default)</b>, <b>bad-host-tos</b>, <b>bad-network-tos</b>, <b>host-prohibited</b>, <b>host-unknown</b>, <b>host-unreachable</b>, <b>network-prohibited</b>, <b>network-unknown</b>, <b>network-unreachable</b>, <b>port-unreachable</b>, <b>precedence-cutoff</b>, <b>precedence-violation</b>, <b>protocol-unreachable</b>, <b>source-host-isolated</b>, <b>source-route-failed</b>, <b>tcp-reset</b>.</p> <p>If you specify <b>tcp-reset</b>, a TCP reset is returned if the packet is a TCP packet. Otherwise nothing is returned.</p> <p>If you do not specify a message type, the ICMP notification <b>destination unreachable</b> is sent with the default message <b>communication administratively filtered</b>.</p> |
| <code>routing-instance routing-instance-name</code> | <p>Forward matched packets to a virtual routing instance.</p> <p><b>NOTE:</b> EX4200 switches do not support firewall-filter-based redirection to the default routing instance.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <code>vlan vlan-name</code>                         | <p>Forward matched packets to a specific VLAN. Ensure that you specify the VLAN name or VLAN ID and not a VLAN range, because the <b>vlan</b> action does not support the <b>vlan-range</b> option.</p> <p><b>NOTE:</b> If you have defined a VLAN that is enabled for dot1q tunneling, then that particular VLAN is not supported as an action (using the <b>vlan vlan-name</b> action) for an ingress VLAN firewall filter.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

### Action Modifiers for Firewall Filters

In addition to the actions described in [Table 513 on page 4712](#), you can define action modifiers in a firewall filter configuration for a switch if packets match the filtering criteria defined in the match condition. [Table 514 on page 4713](#) describes the action modifiers supported in a firewall filter configuration.

Table 514: Action Modifiers for Firewall Filters

| Action Modifier                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>analyzer analyzer-name</code> | <p>Mirror port traffic to a specified destination port or VLAN that is connected to a protocol analyzer application. Mirroring copies all packets seen on one switch port to a network monitoring connection on another switch port. The analyzer name must be configured under <b>[edit ethernet-switching-options analyzer]</b>.</p> <p><b>NOTE:</b> <b>analyzer</b> is not a supported action modifier for a management interface.</p> <p><b>NOTE:</b> On EX4500 switches, you can configure only one analyzer and include it in a firewall filter. If you configure multiple analyzers, you cannot include any one of those analyzers in a firewall filter.</p> |

Table 514: Action Modifiers for Firewall Filters (*continued*)

| Action Modifier                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>dscp</b> <i>number</i>              | <p>Change the DSCP value for matched packets to the DSCP value specified with this action modifier. <i>number</i> specifies the Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p> <p>For <i>number</i>, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>ef</b> (46)—as defined in <a href="#">RFC 2598</a>, <i>An Expedited Forwarding PHB</i>.</li> <li>• <b>af11</b> (10), <b>af12</b> (12), <b>af13</b> (14), <b>af21</b> (18), <b>af22</b> (20), <b>af23</b> (22), <b>af31</b> (26), <b>af32</b> (28), <b>af33</b> (30), <b>af41</b> (34), <b>af42</b> (36), <b>af43</b> (38)</li> </ul> <p>These four classes, with three drop precedences in each class, are defined for 12 code points in <a href="#">RFC 2597</a>, <i>Assured Forwarding PHB Group</i>.</p> |
| <b>count</b> <i>counter-name</i>       | <p>Count the number of packets that pass this filter, term, or policer. A policer enables you to specify rate limits on traffic that enters an interface on a switch.</p> <p><b>NOTE:</b> On EX4300 switches, you can configure the same number of counters and policers as the number of terms in the ternary content addressable memory (TCAM).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>forwarding-class</b> <i>class</i>   | <p>Classify the packet in one of the following forwarding classes:</p> <ul style="list-style-type: none"> <li>• <b>assured-forwarding</b></li> <li>• <b>best-effort</b></li> <li>• <b>expedited-forwarding</b></li> <li>• <b>network-control</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>interface</b> <i>interface-name</i> | Forward the traffic to the specified interface bypassing the switching lookup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>log</b>                             | <p>Log the packet's header information in the Routing Engine. To view this information, issue the <b>show firewall log</b> command in the CLI.</p> <p><b>NOTE:</b> If the <b>log</b> or the <b>syslog</b> action modifier is configured along with a <b>vlan</b> action or an <b>interface</b> action modifier, the events might not be logged. However, the redirect interface functionality works as expected.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>loss-priority</b> (high   low)      | Set the packet loss priority (PLP).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>policer</b> <i>policer-name</i>     | <p>Apply rate limits to the traffic.</p> <p>You can specify a policer in a firewall filter only for ingress traffic on a port, VLAN, and router.</p> <p><b>NOTE:</b> A counter for a policer is not supported on EX8200 switches.</p> <p><b>NOTE:</b> On EX4300 switches, you can configure the same number of counters and policers as the number of terms in the TCAM.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>port-mirror</b>                     | Mirror packets to the interface defined in the <b>[edit forwarding-options analyzer]</b> hierarchy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 514: Action Modifiers for Firewall Filters (*continued*)

| Action Modifier                                        | Description                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>port-mirror-instance</code> <i>instance-name</i> | Mirror packets to the instance defined in the <code>[edit forwarding-options analyzer]</code> hierarchy.                                                                                                                                                                                                                                                                                           |
| <code>syslog</code>                                    | Log an alert for this packet. You can specify that the log be sent to a server for storage and analysis.<br><br><b>NOTE:</b> If the <code>log</code> or the <code>syslog</code> action modifier is configured along with a <code>vlan</code> action or an <code>interface</code> action modifier, the events might not be logged. However, the redirect interface functionality works as expected. |
| <code>three-color-policer</code>                       | Apply a three-color policer.                                                                                                                                                                                                                                                                                                                                                                       |

**Related Documentation**

- [Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches on page 4718](#)
- [Understanding Firewall Filter Match Conditions on page 4762](#)
- [Firewall Filter Configuration Statements Supported by Junos OS for EX Series Switches on page 4832](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 4795](#)

## Support for Match Conditions and Actions for Loopback Firewall Filters on Switches

On EX Series Ethernet switches, a loopback interface is a gateway for all the control traffic that enters the Routing Engine of the switch. If you want to monitor this control traffic, you must configure a firewall filter on the loopback interface (lo0). Loopback firewall filters are applied only to packets that are sent to the Routing Engine CPU for further processing. Therefore, you can apply a firewall filter only in the ingress direction on the loopback interface.

Each term in a firewall filter consists of *match conditions* and an *action*. Match conditions are the values or fields that a packet must contain. You can define multiple, single, or no match conditions. If no match conditions are specified for the term, all packets are matched by default. The string that defines a match condition is called a *match statement*. The action is the action that the switch takes if a packet matches the match conditions for the specific term. Action modifiers are optional and specify one or more actions that the switch takes if a packet matches the match conditions for the specific term.

The following tables list match conditions, actions, and action modifiers that are supported for a firewall filter configured on a loopback interface on a switch:

- [Table 515 on page 4716](#)
- [Table 516 on page 4717](#)

- [Table 517 on page 4718](#)

For information on match conditions, actions, and action modifiers supported for a firewall filter configured on a network interface, see [“Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches” on page 4718](#).

**Table 515: Match Conditions for Firewall Filters on Loopback Interfaces for IPv4 and IPv6 Traffic—Support per Switch**

| Match Condition                    | EX2200 | EX3200,<br>EX4200 | EX3300 | EX4500 | EX6200 | EX8200 |
|------------------------------------|--------|-------------------|--------|--------|--------|--------|
| Match conditions for IPv4 traffic: |        |                   |        |        |        |        |
| destination-address                | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| destination-port                   | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| destination-prefix-list            | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| dscp                               | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| icmp-code                          | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| icmp-type                          | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| interface                          | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| is-fragment                        | ✓      | ✓                 | ✓      | ✓      | –      | –      |
| packet-length                      | –      | –                 | –      | –      | –      | ✓      |
| precedence                         | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| protocol                           | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| source-address                     | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| source-port                        | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| source-prefix-list                 | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| Match conditions for IPv6 traffic: |        |                   |        |        |        |        |
| destination-address                | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| destination-port                   | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| destination-prefix-list            | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| icmp-code                          | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |

**Table 515: Match Conditions for Firewall Filters on Loopback Interfaces for IPv4 and IPv6 Traffic—Support per Switch** (*continued*)

| Match Condition    | EX2200 | EX3200,<br>EX4200 | EX3300 | EX4500 | EX6200 | EX8200 |
|--------------------|--------|-------------------|--------|--------|--------|--------|
| icmp-type          | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| interface          | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| next-header        | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| packet-length      | –      | –                 | –      | –      | –      | ✓      |
| source-address     | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| source-port        | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| source-prefix-list | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| tcp-established    | ✓      | ✓                 | ✓      | ✓      | ✓      | –      |
| tcp-flags          | ✓      | ✓                 | ✓      | ✓      | ✓      | –      |
| tcp-initial        | ✓      | ✓                 | ✓      | ✓      | ✓      | –      |
| traffic-class      | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |

**Table 516: Actions for Firewall Filters on Loopback Interfaces for IPv4 and IPv6 Traffic—Support per Switch**

| Action                    | EX2200 | EX3200,<br>EX4200 | EX3300 | EX4500 | EX6200 | EX8200 |
|---------------------------|--------|-------------------|--------|--------|--------|--------|
| Actions for IPv4 traffic: |        |                   |        |        |        |        |
| accept                    | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| discard                   | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| Actions for IPv6 traffic: |        |                   |        |        |        |        |
| accept                    | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |
| discard                   | ✓      | ✓                 | ✓      | ✓      | ✓      | ✓      |

**Table 517: Action Modifiers for Firewall Filters on Loopback Interfaces for IPv4 and IPv6 Traffic—Support per Switch**

| Action                             | EX2200 | EX3200,<br>EX4200 | EX3300 | EX4500 | EX6200 | EX8200 |
|------------------------------------|--------|-------------------|--------|--------|--------|--------|
| Action modifiers for IPv4 traffic: |        |                   |        |        |        |        |
| count                              | –      | ✓                 | –      | ✓      | ✓      | –      |
| forwarding-class                   | ✓      | ✓                 | ✓      | ✓      | –      | ✓      |
| loss-priority                      | ✓      | ✓                 | ✓      | ✓      | –      | ✓      |
| Action modifiers for IPv6 traffic: |        |                   |        |        |        |        |
| count                              | –      | ✓                 | –      | ✓      | –      | –      |
| forwarding-class                   | ✓      | ✓                 | ✓      | ✓      | –      | ✓      |
| loss-priority                      | ✓      | ✓                 | ✓      | ✓      | –      | ✓      |



**NOTE:** On EX8200 switches, if an implicit or explicit discard action is configured on a loopback interface for IPv4 traffic, next hop resolve packets are accepted and allowed to pass through the switch. However, for IPv6 traffic, you must explicitly configure a rule to allow the neighbor discovery IPv6 resolve packets to pass through the switch.

#### Related Documentation

- [Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704](#)
- [Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches on page 4718](#)
- [Understanding Firewall Filter Match Conditions on page 4762](#)
- [Understanding How Firewall Filters Are Evaluated on page 4760](#)
- [Understanding How Firewall Filters Test a Packet's Protocol on page 4766](#)
- [Understanding the Use of Policers in Firewall Filters on page 4767](#)

## Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches

After you define a firewall filter on an EX Series switch, you must associate the filter to a bind point so that the filter can filter the packets that enter or exit the bind point. Port firewall filters, VLAN firewall filters, and Layer 3 (or router) firewall filters are the different types of firewall filters you can apply on a switch, depending on the bind points the filters are associated with. While a port firewall filter applies to Layer 2 interfaces, a VLAN

firewall filter applies to packets that enter or leave a VLAN and also to packets that are bridged within a VLAN. A Layer 3 firewall filter applies to Layer 3 (routed) interfaces and routed VLAN interfaces (RVIs).



**NOTE:** If you want to control the traffic that enters the Routing Engine of the switch, you must configure a firewall filter on the loopback interface (lo0) of the switch. For information about match conditions, actions, and action modifiers supported on the loopback interface of a switch, see [“Support for Match Conditions and Actions for Loopback Firewall Filters on Switches”](#) on page 4715.

This topic describes the supported switches and bind points for match conditions, actions, and action modifiers for firewall filters supported on EX Series switches. For descriptions of the match conditions, actions, and action modifiers, see [“Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches”](#) on page 4704.

This topic describes:

- [Firewall Filter Types and Their Bind Points](#) on page 4719
- [Support for IPv4 and IPv6 Firewall Filters on Switches](#) on page 4720
- [Platform Support for Match Conditions for IPv4 Traffic](#) on page 4720
- [Platform Support for Match Conditions for IPv6 Traffic](#) on page 4733
- [Platform Support for Match Conditions for Non-IP Traffic](#) on page 4742
- [Platform Support for Actions for IPv4 Traffic](#) on page 4742
- [Platform Support for Actions for IPv6 Traffic](#) on page 4745
- [Platform Support for Action Modifiers for IPv4 Traffic](#) on page 4748
- [Platform Support for Action Modifiers for IPv6 Traffic](#) on page 4754

### Firewall Filter Types and Their Bind Points

You can apply a firewall filter at specific bind points to filter IPv4, IPv6, or non-IP traffic. See the remaining sections in this topic for information about support on individual switches for different traffic types.

[Table 518 on page 4719](#) lists the firewall filter types and their associated bind points that are supported on the switches.

**Table 518: Bind Points Associated with Firewall Filter Types**

| Bind Points                                                                       | Firewall Filter Type   |
|-----------------------------------------------------------------------------------|------------------------|
| Ports (Layer 2 interfaces)                                                        | Port firewall filter   |
| VLANs                                                                             | VLAN firewall filter   |
| Layer 3 interfaces (Layer 3 (routed) interfaces or routed VLAN interfaces (RVIs)) | Router firewall filter |

### Support for IPv4 and IPv6 Firewall Filters on Switches

On EX2200, EX3200/EX4200, EX3300, EX4500, and EX6200 switches port and VLAN filters on IPv6 traffic can match only layer 2 header fields. On an EX8200 switch, port and VLAN traffic can match on layer 3 and layer 4 header fields, in addition to layer 2 header fields, of IPv6 traffic. On an EX4300 switch, port and VLAN filters on IPv6 traffic can match layer 3 and layer 4 header fields

[Table 519 on page 4720](#) briefly summarizes the support for IPv4 and IPv6 firewall filters on different switches. The support for port, VLAN, and router firewall filters on different switches is further discussed in the subsequent sections in this topic.

**Table 519: Support for IPv4 and IPv6 Firewall Filters on Switches**

| Switch            | Support for IPv4 Firewall Filter | Support for IPv6 Firewall Filter |
|-------------------|----------------------------------|----------------------------------|
| EX2200            | Yes                              | Yes                              |
| EX3200 and EX4200 | Yes                              | Yes                              |
| EX3300            | Yes                              | Yes                              |
| EX4300            | Yes                              | Yes                              |
| EX4500            | Yes                              | Yes                              |
| EX6200            | Yes                              | Yes                              |
| EX8200            | Yes                              | Yes                              |

### Platform Support for Match Conditions for IPv4 Traffic

You can define port, VLAN, and router firewall filters for ingress and egress IPv4 traffic on all EX Series switches. [Table 520 on page 4721](#) summarizes the support for match conditions on different bind points for ingress and egress IPv4 traffic on different switches.



Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches

| Match Condition                                      | Switch            | Supported Bind Points                |                                      |
|------------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                      |                   | Ingress                              | Egress                               |
| <b>destination-address</b><br><i>ip-address</i>      | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                      | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                      | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                      | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                      | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                      | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                      | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>ip-destination-address</b>                        | EX4300            | Ports and VLANs                      | Not Supported                        |
| <b>destination-mac-address</b><br><i>mac-address</i> | EX2200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX3200 and EX4200 | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX3300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX4300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX4500            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX6200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX8200            | Ports and VLANs                      | Ports and VLANs                      |

Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition                                       | Switch            | Supported Bind Points                |                                      |
|-------------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                       |                   | Ingress                              | Egress                               |
| <b>destination-port <i>number</i></b>                 | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                       | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                       | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                       | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                                       | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                       | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                       | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>destination-prefix-list<br/><i>prefix-list</i></b> | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                       | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                       | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                       | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                                       | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                       | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                       | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |

Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition                            | Switch            | Supported Bind Points |                 |
|--------------------------------------------|-------------------|-----------------------|-----------------|
|                                            |                   | Ingress               | Egress          |
| <b>dot1q-tag <i>number</i></b>             | EX2200            | Ports and VLANs       | Ports and VLANs |
|                                            | EX3200 and EX4200 | Ports and VLANs       | Ports and VLANs |
|                                            | EX3300            | Ports and VLANs       | Ports and VLANs |
|                                            | EX4300            | Not Supported         | Not Supported   |
|                                            | EX4500            | Ports and VLANs       | Ports and VLANs |
|                                            | EX6200            | Ports and VLANs       | Ports and VLANs |
|                                            | EX8200            | Ports and VLANs       | Not supported   |
| <b>user-vlan-id <i>number</i></b>          | EX4300            | Ports and VLANs       | Ports and VLANs |
| <b>dot1q-user-priority <i>number</i></b>   | EX2200            | Ports and VLANs       | Ports and VLANs |
|                                            | EX3200 and EX4200 | Ports and VLANs       | Ports and VLANs |
|                                            | EX3300            | Ports and VLANs       | Ports and VLANs |
|                                            | EX4300            | Not Supported         | Not Supported   |
|                                            | EX4500            | Ports and VLANs       | Ports and VLANs |
|                                            | EX6200            | Ports and VLANs       | Ports and VLANs |
|                                            | EX8200            | Ports and VLANs       | Ports and VLANs |
| <b>user-vlan-1p-priority <i>number</i></b> | EX4300            | Ports and VLANs       | Ports and VLANs |

Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition         | Switch            | Supported Bind Points                |                                      |
|-------------------------|-------------------|--------------------------------------|--------------------------------------|
|                         |                   | Ingress                              | Egress                               |
| <b>dscp number</b>      | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                         | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                         | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                         | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                         | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                         | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                         | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>ether-type value</b> | EX2200            | Ports and VLANs                      | Ports and VLANs                      |
|                         | EX3200 and EX4200 | Ports and VLANs                      | Ports and VLANs                      |
|                         | EX3300            | Ports and VLANs                      | Ports and VLANs                      |
|                         | EX4300            | Ports and VLANs                      | Ports and VLANs                      |
|                         | EX4500            | Ports and VLANs                      | Ports and VLANs                      |
|                         | EX6200            | Ports and VLANs                      | Ports and VLANs                      |
|                         | EX8200            | Ports and VLANs                      | Not supported                        |

Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition                                | Switch            | Supported Bind Points                |                                      |
|------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                |                   | Ingress                              | Egress                               |
| <b>fragment-flags</b><br><i>fragment-flags</i> | EX2200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                                | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                                | EX3300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                                | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                                | EX4500            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                                | EX6200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                                | EX8200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
| <b>icmp-code <i>number</i></b>                 | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports and Layer 3 interfaces         |
|                                                | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                                | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |

Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition                                                                                                                                                                                                                                             | Switch            | Supported Bind Points                |                                       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|--------------------------------------|---------------------------------------|
|                                                                                                                                                                                                                                                             |                   | Ingress                              | Egress                                |
| <b>icmp-type <i>number</i></b>                                                                                                                                                                                                                              | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces  |
|                                                                                                                                                                                                                                                             | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                    |
|                                                                                                                                                                                                                                                             | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces  |
|                                                                                                                                                                                                                                                             | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                    |
|                                                                                                                                                                                                                                                             | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces  |
|                                                                                                                                                                                                                                                             | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces  |
|                                                                                                                                                                                                                                                             | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces  |
| <b>interface <i>interface-name</i></b><br><br><b>NOTE:</b> This match condition is not supported by firewall filters configured on ingress L3 interfaces and ingress VLAN interfaces when the interface to be matched is aggregate Ethernet (ae) interface. | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces  |
|                                                                                                                                                                                                                                                             | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces  |
|                                                                                                                                                                                                                                                             | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces  |
|                                                                                                                                                                                                                                                             | EX4300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces. |
|                                                                                                                                                                                                                                                             | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces  |
|                                                                                                                                                                                                                                                             | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces  |
|                                                                                                                                                                                                                                                             | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces  |

Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition                                  | Switch            | Supported Bind Points                |                                      |
|--------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                  |                   | Ingress                              | Egress                               |
| <b>ip-options</b>                                | EX2200            | Layer 3 interfaces                   | Not supported                        |
|                                                  | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports and VLANs                      |
|                                                  | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports and VLANs                      |
|                                                  | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                  | EX4500            | Layer 3 interfaces                   | Not supported                        |
|                                                  | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                  | EX8200            | Layer 3 interfaces                   | Not supported                        |
| <b>ip-version version<br/>match_condition(s)</b> | EX2200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                  | EX3200 and EX4200 | Ports and VLANs                      | Ports and VLANs                      |
|                                                  | EX3300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                  | EX4300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                  | EX4500            | Ports and VLANs                      | Ports and VLANs                      |
|                                                  | EX6200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                  | EX8200            | Ports and VLANs                      | Ports and VLANs                      |

Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition                 | Switch            | Supported Bind Points                |                                      |
|---------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                 |                   | Ingress                              | Egress                               |
| <b>is-fragment</b>              | EX2200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                 | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                 | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                 | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                 | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                 | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                 | EX8200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
| <b>precedence precedence</b>    | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                 | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                 | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                 | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                 | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                 | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                 | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>ip-precedence precedence</b> | EX4300            | Ports and VLANs                      | Not supported                        |



Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition                               | Switch            | Supported Bind Points                |                                      |
|-----------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                               |                   | Ingress                              | Egress                               |
| <b>protocol</b> <i>list of protocols</i>      | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                               | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>ip-protocol</b> <i>list of protocols</i>   | EX4300            | Ports and VLANs                      | Ports and VLANs                      |
| <b>source-address</b><br><b>ip-address</b>    | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                               | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>ip-source-address</b><br><b>ip-address</b> | EX4300            | Ports and VLANs                      | Ports and VLANs                      |

Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition                                 | Switch            | Supported Bind Points                |                                      |
|-------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                 |                   | Ingress                              | Egress                               |
| <b>source-mac-address</b><br><i>mac-address</i> | EX2200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX3200 and EX4200 | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX3300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX4300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX4500            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX6200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX8200            | Ports and VLANs                      | Ports and VLANs                      |
| <b>source-port number</b>                       | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                 | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                 | EX3300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                                 | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                                 | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                 | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                 | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |

Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition                              | Switch            | Supported Bind Points                |                                      |
|----------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                              |                   | Ingress                              | Egress                               |
| <b>source-prefix-list</b> <i>prefix-list</i> | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                              | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                              | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                              | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                              | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                              | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                              | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>tcp-established</b>                       | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                              | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                              | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                              | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                              | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                              | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                              | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |

Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition                               | Switch            | Supported Bind Points                |                                      |
|-----------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                               |                   | Ingress                              | Egress                               |
| <b>tcp-flags</b> ( <i>flags tcp-initial</i> ) | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                               | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>tcp-initial</b>                            | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                               | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                               | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |

Table 520: Firewall Filter Match Conditions Supported for IPv4 Traffic on Switches (*continued*)

| Match Condition                   | Switch            | Supported Bind Points |                 |
|-----------------------------------|-------------------|-----------------------|-----------------|
|                                   |                   | Ingress               | Egress          |
| <b>ttl value</b>                  | EX2200            | Layer 3 interfaces    | Not supported   |
|                                   | EX3200 and EX4200 | Layer 3 interfaces    | Not supported   |
|                                   | EX3300            | Layer 3 interfaces    | Not supported   |
|                                   | EX4300            | Layer 3 interfaces    | Not supported   |
|                                   | EX4500            | Layer 3 interfaces    | Not supported   |
|                                   | EX6200            | Layer 3 interfaces    | Not supported   |
|                                   | EX8200            | Layer 3 interfaces    | Not supported   |
| <b>vlan (vlan-name   vlan-id)</b> | EX2200            | Ports and VLANs       | Ports and VLANs |
|                                   | EX3200 and EX4200 | Ports and VLANs       | Ports and VLANs |
|                                   | EX3300            | Ports and VLANs       | Ports and VLANs |
|                                   | EX4500            | Ports and VLANs       | Ports           |
|                                   | EX6200            | Ports and VLANs       | Ports and VLANs |
|                                   | EX8200            | Ports and VLANs       | Ports and VLANs |
| <b>learn-vlan-id vlan-id</b>      | EX4300            | Ports and VLANs       | Ports and VLANs |

#### Platform Support for Match Conditions for IPv6 Traffic

Table 521 on page 4734 summarizes support for match conditions on different bind points for ingress and egress IPv6 traffic on different switches.

Table 521: Firewall Filter Match Conditions Supported for IPv6 Traffic on Switches

| Match Condition                                                                                          | Switch            | Supported Bind Points                |                                      |
|----------------------------------------------------------------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                                                                          |                   | Ingress                              | Egress                               |
| <b>destination-address</b><br><b>ip-addressor</b><br><b>ip6-destination-address</b><br><b>ip-address</b> | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                          | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 (routed) interfaces only     |
|                                                                                                          | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                          | EX4300            | Layer 3 interfaces                   | Layer 3 (routed) interfaces only     |
|                                                                                                          | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                          | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                          | EX8200            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
| <b>destination-mac-address</b><br><b>mac-address</b>                                                     | EX2200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                                                                          | EX3200 and EX4200 | Ports and VLANs                      | Ports and VLANs                      |
|                                                                                                          | EX3300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                                                                          | EX4300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                                                                          | EX4500            | Ports and VLANs                      | Ports and VLANs                      |
|                                                                                                          | EX6200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                                                                          | EX8200            | Ports and VLANs                      | Ports and VLANs                      |
| <b>destination-port</b><br><b>number</b>                                                                 | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                          | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                          | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                          | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                          | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                          | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                          | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |

Table 521: Firewall Filter Match Conditions Supported for IPv6 Traffic on Switches (*continued*)

| Match Condition                                      | Switch            | Supported Bind Points                |                                      |
|------------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                      |                   | Ingress                              | Egress                               |
| <b>destination-prefix-list</b><br><i>prefix-list</i> | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                      | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 (routed) interfaces only     |
|                                                      | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                      | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                      | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                      | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                      | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>dot1q-tag number</b>                              | EX2200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX3200 and EX4200 | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX3300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX4300            | Not Supported                        | Not Supported                        |
|                                                      | EX4500            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX6200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX8200            | Ports and VLANs                      | Not supported                        |
| <b>user-vlan-id number</b>                           | EX4300            | Ports and VLANs                      | Ports and VLANs                      |
| <b>dot1q-user-priority</b><br><i>number</i>          | EX2200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX3200 and EX4200 | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX3300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX4300            | Not Supported                        | Not Supported                        |
|                                                      | EX4500            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX6200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                      | EX8200            | Ports and VLANs                      | Ports and VLANs                      |

Table 521: Firewall Filter Match Conditions Supported for IPv6 Traffic on Switches (*continued*)

| Match Condition                     | Switch            | Supported Bind Points                |                                      |
|-------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                     |                   | Ingress                              | Egress                               |
| <b>user-vlan-1p-priority number</b> | EX4300            | Ports and VLANs                      | Ports and VLANs                      |
| <b>ether-type (ipv6) value</b>      | EX2200            | Ports and VLANs                      | Ports and VLANs                      |
|                                     | EX3200 and EX4200 | Ports and VLANs                      | Ports and VLANs                      |
|                                     | EX3300            | Ports and VLANs                      | Ports and VLANs                      |
|                                     | EX4300            | Ports and VLANs                      | Ports and VLANs                      |
|                                     | EX4500            | Ports and VLANs                      | Ports and VLANs                      |
|                                     | EX6200            | Ports and VLANs                      | Ports and VLANs                      |
|                                     | EX8200            | Ports and VLANs                      | Ports and VLANs.                     |
| <b>icmp-code number</b>             | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                     | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                     | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                     | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                     | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                     | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                     | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |



Table 521: Firewall Filter Match Conditions Supported for IPv6 Traffic on Switches (*continued*)

| Match Condition                                                                                                                                                                                                                                      | Switch            | Supported Bind Points                |                                      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                                                                                                                                                                                                                      |                   | Ingress                              | Egress                               |
| <b>icmp-type number</b>                                                                                                                                                                                                                              | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                                                                                                                                                                      | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                                                                                                                                                                      | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                                                                                                                                                                      | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                                                                                                                                                                      | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                                                                                                                                                                      | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                                                                                                                                                                      | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>interface interface-name</b><br><br><b>NOTE:</b> This match condition is not supported by firewall filters configured on ingress L3 interfaces and ingress VLAN interfaces when the interface to be matched is aggregate Ethernet (ae) interface. | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                                                                                                                                                                                                                      | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                                                                                                                                                                                                                      | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                                                                                                                                                                                                                      | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not Supported                        |
|                                                                                                                                                                                                                                                      | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                                                                                                                                                                                                                                      | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                                                                                                                                                                                                                      | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |

Table 521: Firewall Filter Match Conditions Supported for IPv6 Traffic on Switches (*continued*)

| Match Condition                                     | Switch            | Supported Bind Points                |                                      |
|-----------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                     |                   | Ingress                              | Egress                               |
| <b>ip-version</b> <i>version match_condition(s)</i> | EX2200            | Not supported                        | Not supported                        |
|                                                     | EX3200 and EX4200 | Not supported                        | Not supported                        |
|                                                     | EX3300            | Not supported                        | Not supported                        |
|                                                     | EX4300            | Not supported                        | Not supported                        |
|                                                     | EX4500            | Not supported                        | Not supported                        |
|                                                     | EX6200            | Not supported                        | Not supported                        |
|                                                     | EX8200            | Ports and VLANs                      | Ports and VLANs                      |
| <b>next-header</b> <i>bytes</i>                     | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                     | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                     | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                     | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                     | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                     | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                     | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>packet-length</b> <i>bytes</i>                   | EX2200            | Not supported                        | Not supported                        |
|                                                     | EX3200 and EX4200 | Not supported                        | Not supported                        |
|                                                     | EX3300            | Not supported                        | Not supported                        |
|                                                     | EX4300            | Not supported                        | Not supported                        |
|                                                     | EX4500            | Not supported                        | Not supported                        |
|                                                     | EX6200            | Not supported                        | Not supported                        |
|                                                     | EX8200            | Layer 3 interfaces                   | Not supported                        |

Table 521: Firewall Filter Match Conditions Supported for IPv6 Traffic on Switches (*continued*)

| Match Condition                                 | Switch            | Supported Bind Points                |                                      |
|-------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                 |                   | Ingress                              | Egress                               |
| <b>source-address</b><br><i>ip-address</i>      | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX8200            | Ports, VLANs, Layer 3 interfaces     | Ports, VLANs, and Layer 3 interfaces |
| <b>ip6-source-address</b><br><i>ip-address</i>  | EX4300            | Layer 3 interfaces                   | Not supported                        |
| <b>source-mac-address</b><br><i>mac-address</i> | EX2200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX3200 and EX4200 | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX3300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX4300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX4500            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX6200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX8200            | Ports and VLANs                      | Ports and VLANs                      |
| <b>source-port number</b>                       | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |

Table 521: Firewall Filter Match Conditions Supported for IPv6 Traffic on Switches (*continued*)

| Match Condition                                         | Switch            | Supported Bind Points                |                                      |
|---------------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                         |                   | Ingress                              | Egress                               |
| <b>source-prefix-list</b><br><i>prefix-list</i>         | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>tcp-established</b>                                  | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                                         | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>tcp-flags</b> ( <i>flags</i><br><i>tcp-initial</i> ) | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                                         | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                         | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |

Table 521: Firewall Filter Match Conditions Supported for IPv6 Traffic on Switches (*continued*)

| Match Condition                                 | Switch            | Supported Bind Points                |                                      |
|-------------------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                                 |                   | Ingress                              | Egress                               |
| <b>tcp-initial</b>                              | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX4300            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                                 | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX6200            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces                   |
|                                                 | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>traffic-class <i>number</i></b>              | EX2200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX3200 and EX4200 | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX3300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX4300            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX6200            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                                 | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>vlan (<i>vlan-id</i>   <i>vlan-name</i>)</b> | EX2200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX3200 and EX4200 | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX3300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX4300            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX4500            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX6200            | Ports and VLANs                      | Ports and VLANs                      |
|                                                 | EX8200            | Ports and VLANs                      | Not supported                        |

### Platform Support for Match Conditions for Non-IP Traffic

You can define port, VLAN, and router firewall filters for ingress and egress non-IP traffic on all EX Series switches. [Table 522 on page 4742](#) summarizes support for match conditions on different bind points for ingress and egress non-IP traffic on different switches.

**Table 522: Firewall Filter Match Condition Supported for Non-IP Traffic on Switches**

| Match Condition            | Switch            | Supported Bind Points |                 |
|----------------------------|-------------------|-----------------------|-----------------|
|                            |                   | Ingress               | Egress          |
| l2-encap-type llc-non-snap | EX2200            | Ports and VLANs       | Ports and VLANs |
|                            | EX3200 and EX4200 | Ports and VLANs       | Ports and VLANs |
|                            | EX3300            | Ports and VLANs       | Ports and VLANs |
|                            | EX4300            | Ports and VLANs       | Ports and VLANs |
|                            | EX4500            | Ports and VLANs       | Ports and VLANs |
|                            | EX6200            | Ports and VLANs       | Ports and VLANs |
|                            | EX8200            | Ports and VLANs       | Ports and VLANs |

### Platform Support for Actions for IPv4 Traffic

[Table 523 on page 4743](#) summarizes the support for actions on different bind points for ingress and egress IPv4 traffic on different switches.

Table 523: Firewall Filter Actions Supported for IPv4 Traffic on Switches

| Action         | Switch            | Supported Bind Points                |                                      |
|----------------|-------------------|--------------------------------------|--------------------------------------|
|                |                   | Ingress                              | Egress                               |
| <b>accept</b>  | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX4300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
| <b>discard</b> | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX4300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |

Table 523: Firewall Filter Actions Supported for IPv4 Traffic on Switches (*continued*)

| Action                                                  | Switch            | Supported Bind Points |               |
|---------------------------------------------------------|-------------------|-----------------------|---------------|
|                                                         |                   | Ingress               | Egress        |
| <b>reject</b> <i>message-type</i>                       | EX2200            | Layer 3 interfaces    | Not supported |
|                                                         | EX3200 and EX4200 | Layer 3 interfaces    | Not supported |
|                                                         | EX3300            | Layer 3 interfaces    | Not supported |
|                                                         | EX4300            | Layer 3 interfaces    | Not supported |
|                                                         | EX4500            | Layer 3 interfaces    | Not supported |
|                                                         | EX6200            | Layer 3 interfaces    | Not supported |
|                                                         | EX8200            | Layer 3 interfaces    | Not supported |
| <b>routing-instance</b><br><i>routing-instance-name</i> | EX2200            | Layer 3 interfaces    | Not supported |
|                                                         | EX3200 and EX4200 | Layer 3 interfaces    | Not supported |
|                                                         | EX3300            | Layer 3 interfaces    | Not supported |
|                                                         | EX4300            | Layer 3 interfaces    | Not supported |
|                                                         | EX4500            | Layer 3 interfaces    | Not supported |
|                                                         | EX6200            | Layer 3 interfaces    | Not supported |
|                                                         | EX8200            | Layer 3 interfaces    | Not supported |



Table 523: Firewall Filter Actions Supported for IPv4 Traffic on Switches (*continued*)

| Action                                                                                                                                                                                | Switch            | Supported Bind Points |                 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|-----------------------|-----------------|
|                                                                                                                                                                                       |                   | Ingress               | Egress          |
| <b>vlan</b> <i>vlan-name</i>                                                                                                                                                          | EX2200            | Ports and VLANs       | Not supported   |
|                                                                                                                                                                                       | EX3200 and EX4200 | Ports and VLANs       | Not supported   |
|                                                                                                                                                                                       | EX3300            | Ports and VLANs       | Ports and VLANs |
|                                                                                                                                                                                       | EX4300            | Ports and VLANs       | Not supported   |
|                                                                                                                                                                                       | EX4500            | Ports and VLANs       | Ports           |
|                                                                                                                                                                                       | EX6200            | Ports and VLANs       | Ports and VLANs |
|                                                                                                                                                                                       | EX8200            | Ports and VLANs       | Not supported   |
| <p><b>NOTE:</b> Supported only when used in conjunction with the <b>interface</b> action modifier. On EX8200 Virtual Chassis, the <b>vlan</b> action is supported only for VLANs.</p> |                   |                       |                 |

#### Platform Support for Actions for IPv6 Traffic

Table 524 on page 4746 summarizes the support for actions on different bind points for ingress and egress IPv6 traffic.

Table 524: Firewall Filter Actions Supported for IPv6 Traffic on Switches

| Action         | Switch            | Supported Bind Points                   |                                         |
|----------------|-------------------|-----------------------------------------|-----------------------------------------|
|                |                   | Ingress                                 | Egress                                  |
| <b>accept</b>  | EX2200            | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
|                | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
|                | EX3300            | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
|                | EX4300            | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
|                | EX4500            | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
|                | EX6200            | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
|                | EX8200            | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
| <b>discard</b> | EX2200            | Ports and VLANs, and Layer 3 interfaces | Ports and VLANs, and Layer 3 interfaces |
|                | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
|                | EX3300            | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
|                | EX4300            | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
|                | EX4500            | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
|                | EX6200            | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |
|                | EX8200            | Ports, VLANs, and Layer 3 interfaces    | Ports, VLANs, and Layer 3 interfaces    |

Table 524: Firewall Filter Actions Supported for IPv6 Traffic on Switches (*continued*)

| Action                                                  | Switch            | Supported Bind Points |               |
|---------------------------------------------------------|-------------------|-----------------------|---------------|
|                                                         |                   | Ingress               | Egress        |
| <b>reject</b> <i>message-type</i>                       | EX2200            | Layer 3 interfaces    | Not supported |
|                                                         | EX3200 and EX4200 | Layer 3 interfaces    | Not supported |
|                                                         | EX3300            | Layer 3 interfaces    | Not supported |
|                                                         | EX4300            | Layer 3 interfaces    | Not supported |
|                                                         | EX4500            | Layer 3 interfaces    | Not supported |
|                                                         | EX6200            | Layer 3 interfaces    | Not supported |
|                                                         | EX8200            | Layer 3 interfaces    | Not supported |
| <b>routing-instance</b><br><i>routing-instance-name</i> | EX2200            | Layer 3 interfaces    | Not supported |
|                                                         | EX3200 and EX4200 | Layer 3 interfaces    | Not supported |
|                                                         | EX3300            | Layer 3 interfaces    | Not supported |
|                                                         | EX4300            | Not supported         | Not supported |
|                                                         | EX4500            | Layer 3 interfaces    | Not supported |
|                                                         | EX6200            | Layer 3 interfaces    | Not supported |
|                                                         | EX8200            | Layer 3 interfaces    | Not supported |

Table 524: Firewall Filter Actions Supported for IPv6 Traffic on Switches (*continued*)

| Action                                                                                                                                                                                | Switch            | Supported Bind Points |               |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|-----------------------|---------------|
|                                                                                                                                                                                       |                   | Ingress               | Egress        |
| <b>vlan</b> <i>vlan-name</i>                                                                                                                                                          | EX2200            | Ports and VLANs       | Not supported |
|                                                                                                                                                                                       | EX3200 and EX4200 | Ports and VLANs       | Not supported |
|                                                                                                                                                                                       | EX3300            | Ports and VLANs       | Not supported |
|                                                                                                                                                                                       | EX4300            | Ports and VLANs       | Not supported |
|                                                                                                                                                                                       | EX4500            | Ports and VLANs       | Not supported |
|                                                                                                                                                                                       | EX6200            | Ports and VLANs       | Not supported |
|                                                                                                                                                                                       | EX8200            | Ports and VLANs       | Not supported |
| <p><b>NOTE:</b> Supported only when used in conjunction with the <b>interface</b> action modifier. On EX8200 Virtual Chassis, the <b>vlan</b> action is supported only for VLANs.</p> |                   |                       |               |

#### Platform Support for Action Modifiers for IPv4 Traffic

Table 525 on page 4749 summarizes support for action modifiers on different bind points for ingress and egress IPv4 traffic on different switches.

Table 525: Firewall Filter Action Modifiers Supported for IPv4 Traffic on Switches

| Action Modifier | Switch            | Supported Bind Points                |                    |
|-----------------|-------------------|--------------------------------------|--------------------|
|                 |                   | Ingress                              | Egress             |
| <b>analyzer</b> | EX2200            | Ports, VLANs, and Layer 3 interfaces | Not supported      |
|                 | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Not supported      |
|                 | EX3300            | Ports, VLANs, and Layer 3 interfaces | Not supported      |
|                 | EX4300            | Not supported                        | Not supported      |
|                 | EX4500            | Ports, VLANs, and Layer 3 interfaces | Layer 3 interfaces |
|                 | EX6200            | Ports, VLANs, and Layer 3 interfaces | Not supported      |
|                 | EX8200            | Ports, VLANs, and Layer 3 interfaces | Not supported      |
| <b>dscp</b>     | EX2200            | Not supported                        | Not supported      |
|                 | EX3200 and EX4200 | Not supported                        | Not supported      |
|                 | EX3300            | Not supported                        | Not supported      |
|                 | EX4300            | Not supported                        | Not supported      |
|                 | EX4500            | Not supported                        | Not supported      |
|                 | EX6200            | Not supported                        | Not supported      |
|                 | EX8200            | Layer 3 interfaces                   | Not supported      |

Table 525: Firewall Filter Action Modifiers Supported for IPv4 Traffic on Switches (*continued*)

| Action Modifier               | Switch            | Supported Bind Points                                                      |                                                                  |
|-------------------------------|-------------------|----------------------------------------------------------------------------|------------------------------------------------------------------|
|                               |                   | Ingress                                                                    | Egress                                                           |
| <b>count</b>                  | EX2200            | VLANs and Layer 3 interfaces ( <b>me0</b> interfaces only)                 | Layer 3 interfaces ( <b>me0</b> interfaces only)                 |
|                               | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces                                       | Ports, VLANs, and Layer 3 interfaces                             |
|                               | EX3300            | VLANs and Layer 3 interfaces ( <b>me0</b> and <b>vme0</b> interfaces only) | Layer 3 interfaces ( <b>me0</b> and <b>vme0</b> interfaces only) |
|                               | EX4300            | Ports, VLANs, and Layer 3 interfaces                                       | Ports, VLANs, and Layer 3 interfaces                             |
|                               | EX4500            | Ports, VLANs, and Layer 3 interfaces                                       | Ports, VLANs, and Layer 3 interfaces                             |
|                               | EX6200            | Ports, VLANs, and Layer 3 interfaces                                       | Ports, VLANs, and Layer 3 interfaces                             |
|                               | EX8200            | Ports, VLANs, and Layer 3 interfaces                                       | Not supported                                                    |
| <b>forwarding-class class</b> | EX2200            | Ports, VLANs, and Layer 3 interfaces                                       | Ports and Layer 3 interfaces                                     |
|                               | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces                                       | Ports and Layer 3 interfaces                                     |
|                               | EX3300            | Ports, VLANs, and Layer 3 interfaces                                       | Ports and Layer 3 interfaces                                     |
|                               | EX4300            | Ports, VLANs, and Layer 3 interfaces                                       | Not supported                                                    |
|                               | EX4500            | Ports, VLANs, and Layer 3 interfaces                                       | Ports and Layer 3 interfaces                                     |
|                               | EX6200            | Ports, VLANs, and Layer 3 interfaces                                       | Ports, VLANs, and Layer 3 interfaces                             |
|                               | EX8200            | Ports, VLANs, and Layer 3 interfaces                                       | Ports and Layer 3 interfaces                                     |

Table 525: Firewall Filter Action Modifiers Supported for IPv4 Traffic on Switches (*continued*)

| Action Modifier                           | Switch            | Supported Bind Points                                                                              |               |
|-------------------------------------------|-------------------|----------------------------------------------------------------------------------------------------|---------------|
|                                           |                   | Ingress                                                                                            | Egress        |
| <b>interface</b><br><i>interface-name</i> | EX2200            | Ports and VLANs                                                                                    | Not supported |
|                                           | EX3200 and EX4200 | Ports and VLANs                                                                                    | Not supported |
|                                           | EX3300            | Ports and VLANs                                                                                    | Not supported |
|                                           | EX4300            | Ports and VLANs                                                                                    | Not supported |
|                                           | EX4500            | Ports and VLANs                                                                                    | Not supported |
|                                           | EX6200            | Ports and VLANs                                                                                    | Not supported |
|                                           | EX8200            | Ports and VLANs                                                                                    | Not supported |
|                                           |                   | NOTE: On EX8200 Virtual Chassis, the <b>interface</b> action modifier is supported only for VLANs. |               |
| <b>log</b>                                | EX2200            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX3300            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX4300            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX4500            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX6200            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX8200            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |

Table 525: Firewall Filter Action Modifiers Supported for IPv4 Traffic on Switches (*continued*)

| Action Modifier                    | Switch            | Supported Bind Points                |                                      |
|------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                    |                   | Ingress                              | Egress                               |
| <b>loss-priority (high   low)</b>  | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports and Layer 3 interfaces         |
|                                    | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports and Layer 3 interfaces         |
|                                    | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports and Layer 3 interfaces         |
|                                    | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX4500            | Ports, VLANs, and Layer 3 interfaces | Ports and Layer 3 interfaces         |
|                                    | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                    | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports and Layer 3 interfaces         |
| <b>policer <i>policer-name</i></b> | EX2200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX3300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX4500            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX6200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX8200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |



Table 525: Firewall Filter Action Modifiers Supported for IPv4 Traffic on Switches (*continued*)

| Action Modifier                               | Switch            | Supported Bind Points                |               |
|-----------------------------------------------|-------------------|--------------------------------------|---------------|
|                                               |                   | Ingress                              | Egress        |
| <b>port-mirror</b>                            | EX2200            | Not supported                        | Not supported |
|                                               | EX3200 and EX4200 | Not supported                        | Not supported |
|                                               | EX3300            | Not supported                        | Not supported |
|                                               | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                                               | EX4500            | Not supported                        | Not supported |
|                                               | EX6200            | Not supported                        | Not supported |
|                                               | EX8200            | Not supported                        | Not supported |
| <b>port-mirror-instance<br/>instance-name</b> | EX2200            | Not supported                        | Not supported |
|                                               | EX3200 and EX4200 | Not supported                        | Not supported |
|                                               | EX3300            | Not supported                        | Not supported |
|                                               | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                                               | EX4500            | Not supported                        | Not supported |
|                                               | EX6200            | Not supported                        | Not supported |
|                                               | EX8200            | Not supported                        | Not supported |

Table 525: Firewall Filter Action Modifiers Supported for IPv4 Traffic on Switches (*continued*)

| Action Modifier            | Switch            | Supported Bind Points                |                                      |
|----------------------------|-------------------|--------------------------------------|--------------------------------------|
|                            |                   | Ingress                              | Egress                               |
| <b>syslog</b>              | EX2200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                            | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                            | EX3300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                            | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                            | EX4500            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                            | EX6200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                            | EX8200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
| <b>three-color-policer</b> | EX2200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                            | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                            | EX3300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                            | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                            | EX4500            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                            | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                            | EX8200            | Not supported                        | Not supported                        |

#### Platform Support for Action Modifiers for IPv6 Traffic

Table 526 on page 4755 summarizes support for action modifiers on different bind points for ingress and egress IPv6 traffic.

Table 526: Firewall Filter Action Modifiers Supported for IPv6 Traffic on Switches

| Action Modifier | Switch            | Supported Bind Points                |               |
|-----------------|-------------------|--------------------------------------|---------------|
|                 |                   | Ingress                              | Egress        |
| <b>analyzer</b> | EX2200            | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                 | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                 | EX3300            | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                 | EX4300            | Ports, VLANs, and layer 3 interfaces | Not supported |
|                 | EX4500            | Layer 3 interfaces                   | Not supported |
|                 | EX6200            | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                 | EX8200            | Ports, VLANs, and Layer 3 interfaces | Not supported |
| <b>dscp</b>     | EX2200            | Not supported                        | Not supported |
|                 | EX3200 and EX4200 | Not supported                        | Not supported |
|                 | EX3300            | Not supported                        | Not supported |
|                 | EX4300            | Not supported                        | Not supported |
|                 | EX4500            | Not supported                        | Not supported |
|                 | EX6200            | Not supported                        | Not supported |
|                 | EX8200            | Layer 3 interfaces                   | Not supported |

Table 526: Firewall Filter Action Modifiers Supported for IPv6 Traffic on Switches (*continued*)

| Action Modifier               | Switch            | Supported Bind Points                                            |                                                                  |
|-------------------------------|-------------------|------------------------------------------------------------------|------------------------------------------------------------------|
|                               |                   | Ingress                                                          | Egress                                                           |
| <b>count</b>                  | EX2200            | VLANs and Layer 3 interfaces (me0 and vme0 interfaces only)      | Layer 3 interfaces (me0 and vme0 interfaces only)                |
|                               | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces                             | Ports, VLANs, and Layer 3 interfaces                             |
|                               | EX3300            | Layer 3 interfaces ( <b>me0</b> and <b>vme0</b> interfaces only) | Layer 3 interfaces ( <b>me0</b> and <b>vme0</b> interfaces only) |
|                               | EX4300            | Ports, VLANs, and Layer 3 interfaces                             | Ports, VLANs, and Layer 3 interfaces                             |
|                               | EX4500            | Ports, VLANs, and Layer 3 interfaces                             | Ports, VLANs, and Layer 3 interfaces                             |
|                               | EX6200            | Ports, VLANs, and Layer 3 interfaces                             | Ports, VLANs, and Layer 3 interfaces                             |
|                               | EX8200            | Ports, VLANs, and Layer 3 interfaces                             | Not supported                                                    |
| <b>forwarding-class class</b> | EX2200            | Ports, VLANs, and Layer 3 interfaces                             | Ports and Layer 3 interfaces                                     |
|                               | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces                             | Ports and Layer 3 interfaces                                     |
|                               | EX3300            | Ports, VLANs, and Layer 3 interfaces                             | Ports and Layer 3 interfaces                                     |
|                               | EX4300            | Ports, VLANs, and Layer 3 interfaces                             | Not supported                                                    |
|                               | EX4500            | Ports, VLANs, and Layer 3 interfaces                             | Ports, VLANs, and Layer 3 interfaces                             |
|                               | EX6200            | Ports, VLANs, and Layer 3 interfaces                             | Ports and Layer 3 interfaces                                     |
|                               | EX8200            | Ports, VLANs, and Layer 3 interfaces                             | Ports and Layer 3 interfaces                                     |

Table 526: Firewall Filter Action Modifiers Supported for IPv6 Traffic on Switches (*continued*)

| Action Modifier                           | Switch            | Supported Bind Points                                                                              |               |
|-------------------------------------------|-------------------|----------------------------------------------------------------------------------------------------|---------------|
|                                           |                   | Ingress                                                                                            | Egress        |
| <b>interface</b><br><i>interface-name</i> | EX2200            | Ports and VLANs                                                                                    | Not supported |
|                                           | EX3200 and EX4200 | Ports and VLANs                                                                                    | Not supported |
|                                           | EX3300            | Ports and VLANs                                                                                    | Not supported |
|                                           | EX4300            | Ports and VLANs                                                                                    | Not supported |
|                                           | EX4500            | Ports and VLANs                                                                                    | Not supported |
|                                           | EX6200            | Ports and VLANs                                                                                    | Not supported |
|                                           | EX8200            | Ports and VLANs                                                                                    | Not supported |
|                                           |                   | NOTE: On EX8200 Virtual Chassis, the <b>interface</b> action modifier is supported only for VLANs. |               |
| <b>log</b>                                | EX2200            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX3300            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX4300            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX4500            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX6200            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |
|                                           | EX8200            | Ports, VLANs, and Layer 3 interfaces                                                               | Not supported |

Table 526: Firewall Filter Action Modifiers Supported for IPv6 Traffic on Switches (*continued*)

| Action Modifier                    | Switch            | Supported Bind Points                |                                      |
|------------------------------------|-------------------|--------------------------------------|--------------------------------------|
|                                    |                   | Ingress                              | Egress                               |
| <b>loss-priority (high   low)</b>  | EX2200            | Ports, VLANs, and Layer 3 interfaces | Ports and Layer 3 interfaces         |
|                                    | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Ports and Layer 3 interfaces         |
|                                    | EX3300            | Ports, VLANs, and Layer 3 interfaces | Ports and Layer 3 interfaces         |
|                                    | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                    | EX6200            | Ports, VLANs, and Layer 3 interfaces | Ports and Layer 3 interfaces         |
|                                    | EX8200            | Ports, VLANs, and Layer 3 interfaces | Ports and Layer 3 interfaces         |
| <b>policer <i>policer-name</i></b> | EX2200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX3300            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX4300            | Ports, VLANs, and Layer 3 interfaces | Ports, VLANs, and Layer 3 interfaces |
|                                    | EX4500            | Layer 3 interfaces                   | Layer 3 interfaces                   |
|                                    | EX6200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |
|                                    | EX8200            | Ports, VLANs, and Layer 3 interfaces | Not supported                        |

Table 526: Firewall Filter Action Modifiers Supported for IPv6 Traffic on Switches (*continued*)

| Action Modifier                               | Switch            | Supported Bind Points                |               |
|-----------------------------------------------|-------------------|--------------------------------------|---------------|
|                                               |                   | Ingress                              | Egress        |
| <b>port-mirror</b>                            | EX2200            | Not supported                        | Not supported |
|                                               | EX3200 and EX4200 | Not supported                        | Not supported |
|                                               | EX3300            | Not supported                        | Not supported |
|                                               | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                                               | EX6200            | Not supported                        | Not supported |
|                                               | EX8200            | Not supported                        | Not supported |
| <b>port-mirror-instance<br/>instance-name</b> | EX2200            | Not supported                        | Not supported |
|                                               | EX3200 and EX4200 | Not supported                        | Not supported |
|                                               | EX3300            | Not supported                        | Not supported |
|                                               | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                                               | EX6200            | Not supported                        | Not supported |
|                                               | EX8200            | Not supported                        | Not supported |
| <b>syslog</b>                                 | EX2200            | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                                               | EX3200 and EX4200 | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                                               | EX3300            | Ports, VLAN, and Layer 3 interfaces  | Not supported |
|                                               | EX4300            | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                                               | EX4500            | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                                               | EX6200            | Ports, VLANs, and Layer 3 interfaces | Not supported |
|                                               | EX8200            | Ports, VLANs, and Layer 3 interfaces | Not supported |

Table 526: Firewall Filter Action Modifiers Supported for IPv6 Traffic on Switches (*continued*)

| Action Modifier            | Switch            | Supported Bind Points |               |
|----------------------------|-------------------|-----------------------|---------------|
|                            |                   | Ingress               | Egress        |
| <b>three-color-policer</b> | EX2200            | Not supported         | Not supported |
|                            | EX3200 and EX4200 | Not Supported         | Not Supported |
|                            | EX3300            | Not supported         | Not supported |
|                            | EX4300            | Not Supported         | Not Supported |
|                            | EX4500            | Not supported         | Not supported |
|                            | EX6200            | Not supported         | Not supported |
|                            | EX8200            | Not Supported         | Not Supported |

**Related Documentation**

- [Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704](#)
- [Support for Match Conditions and Actions for Loopback Firewall Filters on Switches on page 4715](#)
- [Understanding Firewall Filter Match Conditions on page 4762](#)
- [Firewall Filter Configuration Statements Supported by Junos OS for EX Series Switches on page 4832](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 4795](#)

## Understanding How Firewall Filters Are Evaluated

A firewall filter consists of one or more terms, and the order of the terms within a firewall filter is important. Before you configure firewall filters, you should understand how Juniper Networks EX Series Ethernet Switches evaluate the terms within a firewall filter and how packets are evaluated against the terms.

When a firewall filter consists of a single term, the filter is evaluated as follows:

- If the packet matches all the conditions, the action in the **then** statement is taken.
- If the packet matches all the conditions, and no action is specified in the **then** statement, the default action **accept** is taken.



When a firewall filter consists of more than one term, the firewall filter is evaluated sequentially:

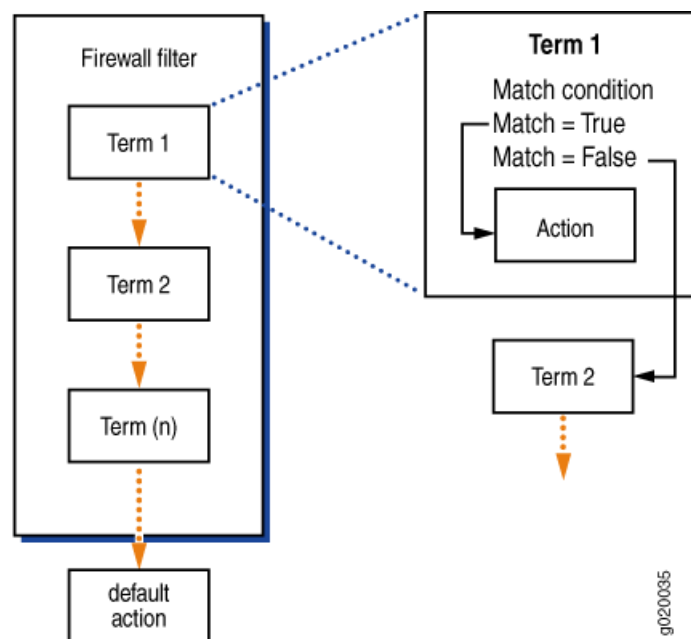
1. The packet is evaluated against the conditions in the **from** statement in the first term.
2. If the packet matches all the conditions in the term, the action in the **then** statement is taken and the evaluation ends. Subsequent terms in the filter are not evaluated.
3. If the packet does not match all the conditions in the term, the packet is evaluated against the conditions in the **from** statement in the second term.

This process continues until either the packet matches the conditions in the **from** statement in one of the subsequent terms or there are no more terms in the filter.

4. If a packet passes through all the terms in the filter without a match, the packet is discarded.

Figure 68 on page 4761 shows how an EX Series switch evaluates the terms within a firewall filter.

Figure 68: Evaluation of Terms Within a Firewall Filter



If a term does not contain a **from** statement, the packet is considered to match and the action in the **then** statement of the term is taken.

If a term does not contain a **then** statement, or if an action has not been configured in the **then** statement, and the packet matches the conditions in the **from** statement of the term, the packet is accepted.

Every firewall filter contains an implicit **deny** statement at the end of the filter, which is equivalent to the following explicit filter term:

```
term implicit-rule {
  then discard;
```

}

Consequently, if a packet passes through all the terms in a filter without matching any conditions, the packet is discarded. If you configure a firewall filter that has no terms, all packets that pass through the filter are discarded.

**Related  
Documentation**

- [Firewall Filters for EX Series Switches Overview on page 4696](#)
- [Understanding Firewall Filter Match Conditions on page 4762](#)
- [Understanding the Use of Policers in Firewall Filters on page 4767](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)

## Understanding Firewall Filter Match Conditions

Before you define terms for firewall filters, you must understand how the match conditions that you specify in a term are handled and how to specify various types of match conditions to achieve the desired filtering results. A match condition consists of a string (called a match statement) that defines the match condition. Match conditions are the values or fields that a packet must contain.

This topic describes:

- [Filter Match Conditions on page 4762](#)
- [Numeric Filter Match Conditions on page 4763](#)
- [Interface Filter Match Conditions on page 4763](#)
- [IP Address Filter Match Conditions on page 4764](#)
- [MAC Address Filter Match Conditions on page 4764](#)
- [Bit-Field Filter Match Conditions on page 4765](#)

### Filter Match Conditions

---

In the **from** statement of a firewall filter term, you specify the packet conditions that trigger the action in one of the **then** statements: **then** with various options, **then interface** or **then vlan**. All conditions in the **from** statement must match for the action to be taken. The order in which you specify match conditions is not important, because a packet must match all the conditions in a term for a match to occur.

If you specify no match conditions in a term, that term matches all packets.

An individual condition in a **from** statement cannot contain a list of values. For example, you cannot specify numeric ranges or multiple source or destination addresses.

Individual conditions in a **from** statement cannot be negated. A negated condition is an explicit mismatch.

## Numeric Filter Match Conditions

Numeric filter conditions match packet fields that are identified by a numeric value, such as port and protocol numbers. For numeric filter match conditions, you specify a keyword that identifies the condition and a single value that a field in a packet must match.

You can specify the numeric value in one of the following ways:

- Single number—A match occurs if the value of the field matches the number. For example:  

```
source-port 25;
```
- Text synonym for a single number— A match occurs if the value of the field matches the number that corresponds to the synonym. For example:  

```
source-port http;
```

To specify more than one value in a filter term, you enter each value in its own match statement, which is a string that defines a match condition. For example, a match occurs in the following term if the value of **vlan** is 10 or 30.

```
[edit firewall family family-name filter filter-name term term-name from]
vlan 10;
vlan 30;
```

The following restrictions apply to numeric filter match conditions:

- You cannot specify a range of values.
- You cannot specify a list of comma-separated values.
- You cannot exclude a specific value in a numeric filter match condition. For example, you cannot specify a condition that would match only if the match condition was not equal to a given value.

## Interface Filter Match Conditions

Interface filter match conditions can match interface name values in a packet. For interface filter match conditions, you specify the name of the interface, for example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set interface ge-0/0/1
```

Port and VLAN interfaces do not use logical unit numbers. However, a firewall filter that is applied to a router interface can specify the logical unit number in the interface filter match condition, for example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set interface ge-0/1/0.0
```

You can include the **\*** wildcard as part of the interface name, for example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set interface ge-0/*/1
user@switch# set interface ge-0/1/*
user@switch# set interface ge-*
```

### IP Address Filter Match Conditions

---

Address filter match conditions can match prefix values in a packet, such as IP source and destination prefixes. For address filter match conditions, you specify a keyword that identifies the field and one prefix of that type that a packet must match.

You specify the address as a single prefix. A match occurs if the value of the field matches the prefix. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-address 10.2.1.0/28;
```

Each prefix contains an implicit 0/0 except statement, which means that any prefix that does not match the prefix that is specified is explicitly considered not to match.

To specify the address prefix, use the notation prefix/prefix-length. If you omit prefix-length, it defaults to /32. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-address 10
[edit firewall family family-name filter filter-name term term-name from]
user@switch# show destination-address
10.0.0.0/32;
```

To specify more than one IP address in a filter term, you enter each address in its own match statement. For example, a match occurs in the following term if the value of the **source-address** field matches either of the following source-address prefixes:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-address 10.0.0.0/8
user@switch# set source-address 10.1.0.0/16
```

### MAC Address Filter Match Conditions

---

MAC address filter match conditions can match source and destination MAC address values in a packet. For MAC address filter match conditions, you specify a keyword that identifies the field and one value of that type that a packet must match.

You can specify the MAC address as six hexadecimal bytes in the following formats:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 0011.2233.4455
```

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 00:11:22:33:44:55
```

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 001122334455
```

To specify more than one MAC address in a filter term, you enter each MAC address in its own match statement. For example, a match occurs in the following term if the value of the **source-mac-address** field matches either of the following addresses.

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-mac-address 00:11:22:33:44:55
user@switch# set source-mac-address 00:11:22:33:20:15
```

### Bit-Field Filter Match Conditions

Bit-field filter conditions match packet fields if particular bits in those fields are or are not set. You can match the IP options, TCP flags, and IP fragmentation fields. For bit-field filter match conditions, you specify a keyword that identifies the field and tests to determine that the option is present in the field.

To specify the bit-field value to match, enclose the value in double quotation marks. For example, a match occurs if the **RST** bit in the TCP flags field is set:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags "rst"
```

Typically, you specify the bits to be tested by using keywords. Bit-field match keywords always map to a single bit value. You also can specify bit fields as hexadecimal or decimal numbers.

To match multiple bit-field values, use the logical operators, which are described in [Table 527 on page 4765](#). The operators are listed in order from highest precedence to lowest precedence. Operations are left-associative.

**Table 527: Logical Operators for Matching Multiple Bit-Field Operators**

| Logical Operators | Description  |
|-------------------|--------------|
| !                 | Negation.    |
| &                 | Logical AND. |
|                   | Logical OR.  |

To negate a match, precede the value with an exclamation point. For example, a match occurs only if the RST bit in the TCP flags field is not set:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags "!rst"
```

In the following example of a logical AND operation, a match occurs if the packet is the initial packet on a TCP session:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags "syn&!ack"
```

In the following example of a logical OR operation, a match occurs if the packet is not the initial packet on a TCP session:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags "syn|ack"
```

For a logical OR operation, you can specify a maximum of two match conditions in a single term. If you need to match more than two bit-field values in a logical OR operation, configure the same match condition in consecutive terms with additional bit-field values. In the following example, the two terms configured match the SYN, ACK, FIN, or RST bit in the TCP flags field:

```
[edit firewall family family-name filter filter-name term term-name1 from]
user@switch# set tcp-flags "syn|ack"
[edit firewall family family-name filter filter-name term term-name2 from]
```

```
user@switch# set tcp-flags "fin|rst"
```

You can use text synonyms to specify some common bit-field matches. You specify these matches as a single keyword. In the following example of a text synonym, a match occurs if the packet is the initial packet on a TCP session:

```
[edit firewall family family-name filter filter-name term term-name from]  
user@switch# set tcp-flags tcp-initial
```

#### Related Documentation

- [Firewall Filters for EX Series Switches Overview on page 4696](#)
- [Understanding How Firewall Filters Test a Packet's Protocol on page 4766](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 4795](#)
- [Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704](#)

## Understanding How Firewall Filters Test a Packet's Protocol

When examining match conditions, Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches tests only the field that is specified. The software does not implicitly test the IP header to determine whether a packet is an IP packet. Therefore, in some cases, you must specify **protocol** field match conditions in conjunction with other match conditions to ensure that the filters are performing the expected matches.

If you specify a protocol match condition or a match of the ICMP type or TCP flags field, there is no implied protocol match. For the following match conditions, you must explicitly specify the protocol match condition in the same term:

- **destination-port**—Specify the match **protocol tcp** or **protocol udp**.
- **source-port**—Specify the match **protocol tcp** or **protocol udp**.

If you do not specify the protocol when using the preceding fields, design your filters carefully to ensure that they perform the expected matches. For example, if you specify a match of **destination-port ssh**, the switch deterministically matches any packets that have a value of **22** in the two-byte field that is two bytes beyond the end of the IP header without ever checking the IP protocol field.

#### Related Documentation

- [Firewall Filters for EX Series Switches Overview on page 4696](#)
- [Understanding Firewall Filter Match Conditions on page 4762](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)

## Understanding the Use of Policers in Firewall Filters

Policing, or rate limiting, is an important component of firewall filters that lets you control the amount of traffic that enters an interface on Juniper Networks EX Series Ethernet Switches. You can achieve policing by including policers in firewall filter configurations.

This topic describes:

- [Policers Overview on page 4767](#)
- [Policer Types on page 4767](#)
- [Policer Actions on page 4768](#)
- [Policer Levels on page 4769](#)
- [Color Modes on page 4769](#)
- [Naming Conventions for Policers on page 4770](#)

### Policers Overview

---

You can use policers to specify rate limits on traffic. A firewall filter configured with a policer permits only traffic within a specified set of rate limits, thereby providing protection from denial-of-service (DoS) attacks. Traffic that exceeds the rate limits specified by the policer is either discarded immediately or is marked as lower priority than traffic that is within the rate limits. The switch discards the lower-priority traffic when there is traffic congestion.

A policer applies two types of rate limits on traffic:

- **Bandwidth**—The number of bits per second permitted, on average.
- **Maximum burst size**—The maximum size permitted for bursts of data that exceed the given bandwidth limit.

Policing uses an algorithm to enforce a limit on average bandwidth while allowing bursts up to a specified maximum value. You can define specific classes of traffic on an interface and apply a set of rate limits to each class. After you name and configure a policer, it is stored as a template. You can then use the policer in a firewall filter configuration.

On all EX Series switches except Juniper Networks EX8200 Ethernet Switches, each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limit specified for the policer. Each EX8200 switch contains three global management counters. You must assign ingress policers to these global management counters to obtain policer statistics. You can assign any number of ingress policers to each global management counter. The policer statistics for each global management counter are the aggregate of the policer statistics for all policers associated with that global management counter.

To get filter-specific packet counts, you must configure a different policer for each firewall filter. Policers give term-specific counts by default.

### Policer Types

---

Switches support three types of policers:

- **Single-rate two-color**—A two-color policer (sometimes called simply “policer”) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit or simply discard them. A two-color policer is most useful for metering traffic at the port (physical interface) level.
- **Single-rate three-color**—This type of policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and the excess burst size (EBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets are arriving at rates that are below the CBS (green), exceed the CBS but not the EBS (yellow), or exceed the EBS (red). A single-rate three-color policer is most useful when a service is structured according to packet size and not according to peak arrival rate.
- **Two-rate three-color**—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and the peak information rate (PIR), along with their associated burst sizes; the CBS, and the peak burst size (PBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on packets are arriving at rates that are below the CIR (green), exceed the CIR but not the PIR (yellow), or exceed the PIR (red). A two-rate three-color policer is most useful when a service is structured according to arrival rates and not to packet size.

### Policer Actions

Policer actions can be implicit or explicit and vary by policer type. The term implicit means that Junos OS assigns a loss-priority value automatically; explicit means that you configure the action. [Table 528 on page 4768](#) lists policer actions.

**Table 528: Policer Actions**

| Policer               | Marking             | Implicit Action          | Configurable Action                                                     |
|-----------------------|---------------------|--------------------------|-------------------------------------------------------------------------|
| Single-rate two-color | Green (Conforming)  | Assign low loss priority | None                                                                    |
|                       | Red (Nonconforming) | None                     | Assign low or high loss priority, assign a forwarding class, or discard |
|                       | Yellow              | Not supported            | Not supported                                                           |



Table 528: Policer Actions (*continued*)

| Policer                 | Marking                                  | Implicit Action                                                           | Configurable Action                                   |
|-------------------------|------------------------------------------|---------------------------------------------------------------------------|-------------------------------------------------------|
| Single-rate three-color | Green (Conforming)                       | Assign low loss priority                                                  | None                                                  |
|                         | Red (Above the EBS)                      | Assign low loss priority                                                  | Discard                                               |
|                         | Yellow (Exceeds the CBS but not the EBS) | Assign low loss priority<br><i>NOTE:</i> Not supported on EX8200 switches | None<br><i>NOTE:</i> Not supported on EX8200 switches |
| Two-rate three-color    | Green (Conforming)                       | Assign low loss priority                                                  | None                                                  |
|                         | Red (Above the PIR)                      | Assign low loss priority                                                  | Discard                                               |
|                         | Yellow (Exceeds the CIR but not the PIR) | Assign low loss priority<br><i>NOTE:</i> Not supported on EX8200 switches | None<br><i>NOTE:</i> Not supported on EX8200 switches |

### Policer Levels

You can configure policers at the queue level, logical interface level, or Layer 2 (MAC) level. Only a single policer is applied to a packet at the egress queue. The search for policers occurs in this order:

- Queue level
- Logical interface level
- Layer 2 (MAC) level

### Color Modes

Tricolor marking (TCM) policers are not bound by a green-yellow-red coloring convention. Packets are marked with low or high PLP bit configurations based on color. Therefore, both three-color policer types (single-rate and two-rate) extend the functionality of class-of-service (CoS) traffic policing by providing three levels of drop precedence (loss priority) instead of the two normally available in policers. Both single-rate and two-rate three-color policer types can operate in two modes:

- Color-blind—In color-blind mode, the three-color policer operates without reference to whether the examined packets have been previously marked or metered. In other words, the three-color policer is *blind* to any previous coloring a packet might have had.
- Color-aware—In color-aware mode, the three-color policer operates with reference to any previous marking or metering of the examined packets. In other words, the three-color policer is *aware* of the previous coloring a packet might have had. In color-aware mode, the three-color policer can increase the PLP of a packet but can never decrease it. For example, if a color-aware three-color policer meters a packet

with a low PLP marking, it can raise the PLP level to high. But it cannot reduce a high PLP level to low.

### Naming Conventions for Policers

---

We recommend you use the naming convention *rate-TCMnumber-color*type when configuring three-color policers. TCM stands for tricolor marking. Because policers can be numerous and must be applied correctly to work, observing a simple naming convention makes it easier to apply the policers properly.

For example, if you configure a single-rate, three-color, color-aware policer, name it srTCM1-ca. If you configure a two-rate, three-color, color-blind policer, name it trTCM2-cb.

#### Related Documentation

- [Firewall Filters for EX Series Switches Overview on page 4696](#)
- [Understanding Tricolor Marking Architecture on page 4770](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704](#)

## Understanding Filter-Based Forwarding for EX Series Switches

Administrators of Juniper Networks EX Series Ethernet Switches can use firewall filters in conjunction with virtual routing instances to specify different routes for packets to travel in their networks. To set up this feature, which is called filter-based forwarding, you specify a filter and match criteria and then specify the virtual routing instance to send packets to.

You might want to use filter-based forwarding to route specific types of traffic through a firewall or security device before the traffic continues on its path. You can also use filter-based forwarding to give certain types of traffic preferential treatment or to improve load balancing of switch traffic.

#### Related Documentation

- [Understanding Virtual Routing Instances on EX Series Switches on page 2263](#)
- [Firewall Filters for EX Series Switches Overview on page 4696](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 4795](#)

## Understanding Tricolor Marking Architecture

Tricolor marking (TCM) policers provide two functions: metering and marking. A policer meters each packet and passes the packet and the metering result to the marker.

The meter operates in two modes. In the color-blind mode, the meter treats the packet stream as uncolored. Any preset loss priorities are ignored. In the color-aware mode, the meter inspects the packet loss priority (PLP) field, which has been set by an upstream device as high or low; in other words, the PLP field has already been set by a behavior

aggregate (BA) or multifield (MF) classifier. The marker changes the PLP of each incoming IP packet according to the results of the meter.

Single-rate TCM is so called because traffic is policed according to one rate—the committed burst rate (CBR)—and two burst sizes: the committed burst size (CBS) and the excess burst size (EBS). The configured information rate (CIR) specifies the average rate at which bits are admitted to the network. The CBS specifies the usual burst size in bytes and the EBS specifies the maximum burst size in bytes for packets that are admitted to the network. The EBS is greater than or equal to the CBS, and neither can be 0. As each packet enters the network, its bytes are counted. Packets that do not exceed the CBS are marked low PLP. Packets that exceed the peak information rate (PIR) are marked high PLP.

Two-rate TCM is so called because traffic is policed according to two rates: the CIR and the PIR. The PIR is greater than or equal to the CIR. The CIR specifies the average rate at which bits are admitted to the network, and the PIR specifies the maximum rate at which bits are admitted to the network. As each packet enters the network, its bits are counted. Bits in packets that do not exceed the CIR have their packets marked low PLP. Bits in packets that exceed the PIR have their packets marked high PLP.

**Related  
Documentation**

- [Understanding the Use of Policers in Firewall Filters on page 4767](#)
- [Configuring Tricolor Marking Policers on page 4827](#)



## CHAPTER 78

# Configuration

- [Configuration Examples on page 4773](#)
- [Configuration Tasks on page 4803](#)
- [Configuration Statements on page 4829](#)

### Configuration Examples

---

- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch on page 4791](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 4795](#)
- [Example: Applying Firewall Filters to Multiple Supplicants on Interfaces Enabled for 802.1X or MAC RADIUS Authentication on page 4798](#)

### Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches

This example shows how to configure and apply firewall filters to control traffic that is entering or exiting a port on the switch, a VLAN on the network, and a Layer 3 interface on the switch. Firewall filters define the rules that determine whether to forward or deny packets at specific processing points in the packet flow.

- [Requirements on page 4774](#)
- [Overview on page 4774](#)
- [Configuring an Ingress Port Firewall Filter to Prioritize Voice Traffic and Rate-Limit TCP and ICMP Traffic on page 4777](#)
- [Configuring a VLAN Ingress Firewall Filter to Prevent Rogue Devices from Disrupting VoIP Traffic on page 4782](#)
- [Configuring a VLAN Firewall Filter to Count, Monitor, and Analyze Egress Traffic on the Employee VLAN on page 4784](#)
- [Configuring a VLAN Firewall Filter to Restrict Guest-to-Employee Traffic and Peer-to-Peer Applications on the Guest VLAN on page 4786](#)

- [Configuring a Router Firewall Filter to Give Priority to Egress Traffic Destined for the Corporate Subnet on page 4788](#)
- [Verification on page 4789](#)

## Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.0 or later for EX Series switches.
- Two Juniper Networks EX3200-48T switches: one to be used as an access switch, the other to be used as a distribution switch
- One Juniper Networks EX-UM-4SFP uplink module
- One Juniper Networks J-series router

Before you configure and apply the firewall filters in this example, be sure you have:

- An understanding of firewall filter concepts, policers, and CoS
- Installed the uplink module in the distribution switch. See *Installing an Uplink Module in an EX3200 Switch*.

## Overview

This configuration example show how to configure and apply firewall filters to provide rules to evaluate the contents of packets and determine when to discard, forward, classify, count, and analyze packets that are destined for or originating from the EX Series switches that handle all **voice-vlan**, **employee-vlan**, and **guest-vlan** traffic. [Table 529 on page 4774](#) shows the firewall filters that are configured for the EX Series switches in this example.

**Table 529: Configuration Components: Firewall Filters**

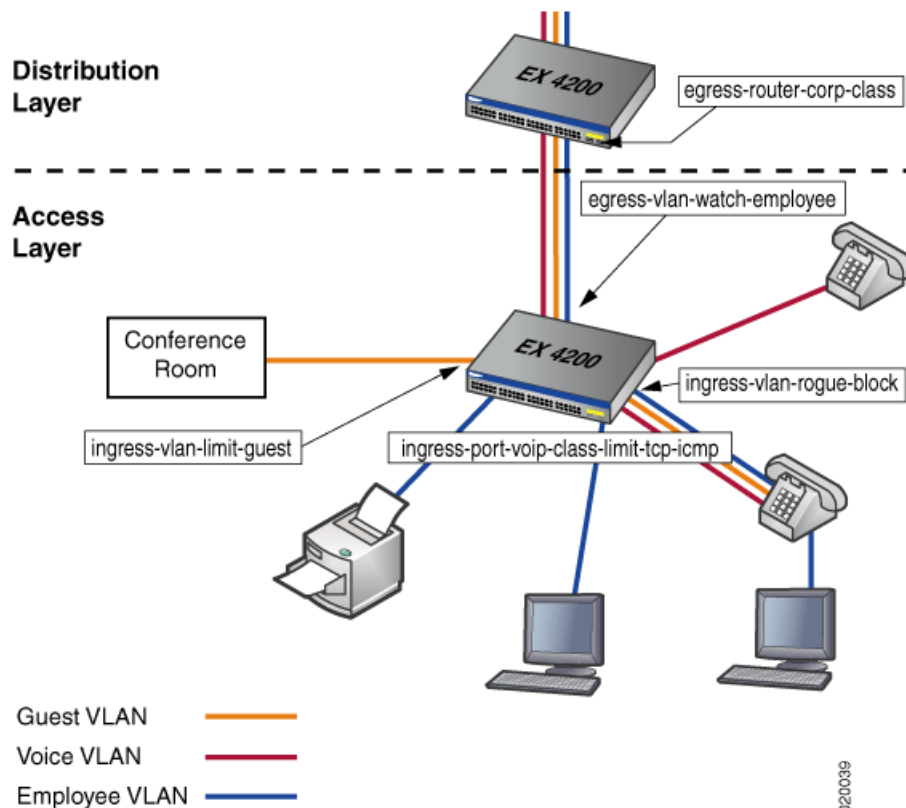
| Component                                                              | Purpose/Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port firewall filter,<br><b>ingress-port-voip-class-limit-tcp-icmp</b> | <p>This firewall filter performs two functions:</p> <ul style="list-style-type: none"> <li>• Assigns priority queueing to packets with a source MAC address that matches the phone MAC addresses. The forwarding class <b>expedited-forwarding</b> provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service for all <b>voice-vlan</b> traffic.</li> <li>• Performs rate limiting on packets that enter the ports for <b>employee-vlan</b>. The traffic rate for TCP and ICMP packets is limited to 1 Mbps with a burst size up to 30,000 bytes.</li> </ul> <p>This firewall filter is applied to port interfaces on the access switch.</p> |
| VLAN firewall filter,<br><b>ingress-vlan-rogue-block</b>               | <p>Prevents rogue devices from using HTTP sessions to mimic the gatekeeper device that manages call registration, admission, and call status for VoIP calls. Only TCP or UDP ports should be used; and only the gatekeeper uses HTTP. That is, all <b>voice-vlan</b> traffic on TCP ports should be destined for the gatekeeper device. This firewall filter applies to all phones on <b>voice-vlan</b>, including communication between any two phones on the VLAN and all communication between the gatekeeper device and VLAN phones.</p> <p>This firewall filter is applied to VLAN interfaces on the access switch.</p>                                          |

Table 529: Configuration Components: Firewall Filters (*continued*)

| Component                                                  | Purpose/Description                                                                                                                                                                                                                                                                                           |
|------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VLAN firewall filter,<br><b>egress-vlan-watch-employee</b> | Accepts <b>employee-vlan</b> traffic destined for the corporate subnet, but does not monitor this traffic. Employee traffic destined for the Web is counted and analyzed.<br><br>This firewall filter is applied to vlan interfaces on the access switch.                                                     |
| VLAN firewall filter,<br><b>ingress-vlan-limit-guest</b>   | Prevents guests (non-employees) from talking with employees or employee hosts on <b>employee-vlan</b> . Also prevents guests from using peer-to-peer applications on <b>guest-vlan</b> , but allows guests to access the Web.<br><br>This firewall filter is applied to VLAN interfaces on the access switch. |
| Router firewall filter,<br><b>egress-router-corp-class</b> | Prioritizes <b>employee-vlan</b> traffic, giving highest forwarding-class priority to employee traffic destined for the corporate subnet.<br><br>This firewall filter is applied to a routed port (Layer 3 uplink module) on the distribution switch.                                                         |

Figure 69 on page 4775 shows the application of port, VLAN, and Layer 3 routed firewall filters on the switch.

Figure 69: Application of Port, VLAN, and Layer 3 Routed Firewall Filters



**Network Topology**

The topology for this configuration example consists of one EX-3200-48T switch at the access layer, and one EX-3200-48T switch at the distribution layer. The distribution switch's uplink module is configured to support a Layer 3 connection to a J-series router.

The EX Series switches are configured to support VLAN membership.

[Table 530 on page 4776](#) shows the VLAN configuration components for the VLANs.

**Table 530: Configuration Components: VLANs**

| VLAN Name     | VLAN ID | VLAN Subnet and Available IP Addresses                                                    | VLAN Description                                                                                                                                                                                                                                                                                   |
|---------------|---------|-------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| voice-vlan    | 10      | 192.0.2.0/28 192.0.2.1 through 192.0.2.14<br><br>192.0.2.15 is subnet's broadcast address | Voice VLAN used for employee VoIP traffic                                                                                                                                                                                                                                                          |
| employee-vlan | 20      | 192.0.2.16/28 192.0.2.17 through 192.0.2.30 192.0.2.31 is subnet's broadcast address      | VLAN standalone PCs, PCs connected to the network through the hub in VoIP telephones, wireless access points, and printers. This VLAN completely includes the voice VLAN. Two VLANs ( <b>voice-vlan</b> and <b>employee-vlan</b> ) must be configured on the ports that connect to the telephones. |
| guest-vlan    | 30      | 192.0.2.32/28 192.0.2.33 through 192.0.2.46 192.0.2.47 is subnet's broadcast address      | VLAN for guests' data devices (PCs). The scenario assumes that the corporation has an area open to visitors, either in the lobby or in a conference room, that has a hub to which visitors can plug in their PCs to connect to the Web and to their company's VPN.                                 |
| camera-vlan   | 40      | 192.0.2.48/28 192.0.2.49 through 192.0.2.62 192.0.2.63 is subnet's broadcast address      | VLAN for the corporate security cameras.                                                                                                                                                                                                                                                           |

Ports on the EX Series switches support Power over Ethernet (PoE) to provide both network connectivity and power for VoIP telephones connecting to the ports.

[Table 531 on page 4777](#) shows the switch ports that are assigned to the VLANs and the IP and MAC addresses for devices connected to the switch ports:



Table 531: Configuration Components: Switch Ports on a 48-Port All-PoE Switch

| Switch and Port Number | VLAN Membership           | IP and MAC Addresses                                                                                       | Port Devices                                                                                                                         |
|------------------------|---------------------------|------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| ge-0/0/0, ge-0/0/1     | voice-vlan, employee-vlan | IP addresses: 192.0.2.1 through 192.0.2.2<br><br>MAC addresses:<br>00.05.85.00.00.01,<br>00.05.85.00–00.02 | Two VoIP telephones, each connected to one PC.                                                                                       |
| ge-0/0/2, ge-0/0/3     | employee-vlan             | 192.0.2.17 through 192.0.2.18                                                                              | Printer, wireless access points                                                                                                      |
| ge-0/0/4, ge-0/0/5     | guest-vlan                | 192.0.2.34 through 192.0.2.35                                                                              | Two hubs into which visitors can plug in their PCs. Hubs are located in an area open to visitors, such as a lobby or conference room |
| ge-0/0/6, ge-0/0/7     | camera-vlan               | 192.0.2.49 through 192.0.2.50                                                                              | Two security cameras                                                                                                                 |
| ge-0/0/9               | voice-vlan                | IP address: 192.0.2.14<br><br>MAC address: 00.05.85.00.00.0E                                               | Gatekeeper device. The gatekeeper manages call registration, admission, and call status for VoIP phones.                             |
| ge-0/1/0               |                           | IP address: 192.0.2.65                                                                                     | Layer 3 connection to a router; note that this is a port on the switch's uplink module                                               |

### Configuring an Ingress Port Firewall Filter to Prioritize Voice Traffic and Rate-Limit TCP and ICMP Traffic

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

#### CLI Quick Configuration

To quickly configure and apply a port firewall filter to prioritize voice traffic and rate-limit packets that are destined for the **employee-vlan** subnet, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall policer tcp-connection-policer if-exceeding burst-size-limit 30k bandwidth-limit 1m
set firewall policer tcp-connection-policer then discard
set firewall policer icmp-connection-policer if-exceeding burst-size-limit 30k bandwidth-limit 1m
set firewall policer icmp-connection-policer then discard
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
from source-mac-address 00.05.85.00.00.01
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
from source-mac-address 00.05.85.00.00.02
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
from protocol udp
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
then forwarding-class expedited-forwarding
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
then loss-priority low
```

```
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
network-control from precedence net-control
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
network-control then forwarding-class network-control
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
network-control then loss-priority low
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection from destination-address 192.0.2.16/28
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection from protocol tcp
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then policer tcp-connection-policer
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then count tcp-counter
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then forwarding-class best-effort
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then loss-priority high
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection from destination-address 192.0.2.16/28
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection from protocol icmp
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then policer icmp-connection-policer
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then count icmp-counter
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then forwarding-class best-effort
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then loss-priority high
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term best-effort
then forwarding-class best-effort
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term best-effort
then loss-priority high
set interfaces ge-0/0/0 description "voice priority and tcp and icmp traffic rate-limiting filter at
ingress port"
set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
set interfaces ge-0/0/1 description "voice priority and tcp and icmp traffic rate-limiting filter at
ingress port"
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
set class-of-service schedulers voice-high buffer-size percent 15
set class-of-service schedulers voice-high priority high
set class-of-service schedulers net-control buffer-size percent 10
set class-of-service schedulers net-control priority high
set class-of-service schedulers best-effort buffer-size percent 75
set class-of-service schedulers best-effort priority low
set class-of-service scheduler-maps ethernet-diffsrv-cos-map forwarding-class
expedited-forwarding scheduler voice-high
set class-of-service scheduler-maps ethernet-diffsrv-cos-map forwarding-class network-control
scheduler net-control
set class-of-service scheduler-maps ethernet-diffsrv-cos-map forwarding-class best-effort
scheduler best-effort
```

**Step-by-Step Procedure** To configure and apply a port firewall filter to prioritize voice traffic and rate-limit packets that are destined for the **employee-vlan** subnet:

1. Define the policers **tcp-connection-policer** and **icmp-connection-policer**:
 

```
[edit]
user@switch# set firewall policer tcp-connection-policer if-exceeding burst-size-limit 30k
bandwidth-limit 1m
user@switch# set firewall policer tcp-connection-policer then discard
user@switch# set firewall policer icmp-connection-policer if-exceeding burst-size-limit
30k bandwidth-limit 1m
user@switch# set firewall policer icmp-connection-policer then discard
```
2. Define the firewall filter **ingress-port-voip-class-limit-tcp-icmp**:
 

```
[edit firewall]
user@switch# set family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp
```
3. Define the term **voip-high**:
 

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp ]
user@switch# set term voip-high from source-mac-address 00.05.85.00.00.01
user@switch# set term voip-high from source-mac-address 00.05.85.00.00.02
user@switch# set term voip-high from protocol udp
user@switch# set term voip-high then forwarding-class expedited-forwarding
user@switch# set term voip-high then loss-priority low
```
4. Define the term **network-control**:
 

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp ]
user@switch# set term network-control from precedence net-control
user@switch# set term network-control then forwarding-class network-control
user@switch# set term network-control then loss-priority low
```
5. Define the term **tcp-connection** to configure rate limits for TCP traffic:
 

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp]
user@switch# set term tcp-connection from destination-address 192.0.2.16/28
user@switch# set term tcp-connection from protocol tcp
user@switch# set term tcp-connection then policer tcp-connection-policer
user@switch# set term tcp-connection then count tcp-counter
user@switch# set term tcp-connection then forwarding-class best-effort
user@switch# set term tcp-connection then loss-priority high
```
6. Define the term **icmp-connection** to configure rate limits for ICMP traffic:
 

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp]
user@switch# set term icmp-connection from destination-address 192.0.2.16/28
user@switch# set term icmp-connection from protocol icmp
user@switch# set term icmp-connection then policer icmp-policer
user@switch# set term icmp-connection then count icmp-counter
user@switch# set term icmp-connection then forwarding-class best-effort
user@switch# set term icmp-connection then loss-priority high
```
7. Define the term **best-effort** with no match conditions for an implicit match on all packets that did not match any other term in the firewall filter:
 

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp]
user@switch# set term best-effort then forwarding-class best-effort
user@switch# set term best-effort then loss-priority high
```

8. Apply the firewall filter **ingress-port-voip-class-limit-tcp-icmp** as an input filter to the port interfaces for **employee-vlan** :

```
[edit interfaces]
user@switch# set ge-0/0/0 description "voice priority and tcp and icmp traffic rate-limiting
filter at ingress port"
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
user@switch# set ge-0/0/1 description "voice priority and tcp and icmp traffic rate-limiting
filter at ingress port"
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
```

9. Configure the parameters that are desired for the different schedulers.



**NOTE:** When you configure parameters for the schedulers, define the numbers to match your network traffic patterns.

```
[edit class-of-service]
user@switch# set schedulers voice-high buffer-size percent 15
user@switch# set schedulers voice-high priority high
user@switch# set schedulers network-control buffer-size percent 10
user@switch# set schedulers network-control priority high
user@switch# set schedulers best-effort buffer-size percent 75
user@switch# set schedulers best-effort priority low
```

10. Assign the forwarding classes to schedulers with a scheduler map:

```
[edit class-of-service]
user@switch# set scheduler-maps ethernet-diffsrv-cos-map
user@switch# set scheduler-maps ethernet-diffsrv-cos-map forwarding-class
expedited-forwarding scheduler voice-high
user@switch# set scheduler-maps ethernet-diffsrv-cos-map forwarding-class
network-control scheduler net-control
user@switch# set scheduler-maps ethernet-diffsrv-cos-map forwarding-class best-effort
scheduler best-effort
```

11. Associate the scheduler map with the outgoing interface:

```
[edit class-of-service]
user@switch# set interfaces ge-0/1/0 scheduler-map ethernet-diffsrv-cos-map
```

**Results** Display the results of the configuration:

```
user@switch# show
firewall {
  policer tcp-connection-policer {
    if-exceeding {
      bandwidth-limit 1m;
      burst-size-limit 30k;
    }
    then {
      discard;
    }
  }
  policer icmp-connection-policer {
    if-exceeding {
      bandwidth-limit 1m;
    }
  }
}
```

```

        burst-size-limit 30k;
    }
    then {
        discard;
    }
}
family ethernet-switching {
    filter ingress-port-voip-class-limit-tcp-icmp {
        term voip-high {
            from {
                destination-mac-address 00.05.85.00.00.01;
                destination-mac-address 00.05.85.00.00.02;
                protocol udp;
            }
            then {
                forwarding-class expedited-forwarding;
                loss-priority low;
            }
        }
        term network-control {
            from {
                precedence net-control ;
            }
            then {
                forwarding-class network-control;
                loss-priority low;
            }
        }
        term tcp-connection {
            from {
                destination-address 192.0.2.16/28;
                protocol tcp;
            }
            then {
                policer tcp-connection-policer;
                count tcp-counter;
                forwarding-class best-effort;
                loss-priority high;
            }
        }
        term icmp-connection
            from {
                protocol icmp;
            }
            then {
                policer icmp-connection-policer;
                count icmp-counter;
                forwarding-class best-effort;
                loss-priority high;
            }
        }
        term best-effort {
            then {
                forwarding-class best-effort;
                loss-priority high;
            }
        }
    }
}

```

```
    }
  }
}
interfaces {
  ge-0/0/0 {
    description "voice priority and tcp and icmp traffic rate-limiting filter at ingress port";
    unit 0 {
      family ethernet-switching {
        filter {
          input ingress-port-voip-class-limit-tcp-icmp;
        }
      }
    }
  }
  ge-0/0/1 {
    description "voice priority and tcp and icmp traffic rate-limiting filter at ingress port";
    unit 0 {
      family ethernet-switching {
        filter {
          input ingress-port-voip-class-limit-tcp-icmp;
        }
      }
    }
  }
}
scheduler-maps {
  ethernet-diffsrv-cos-map {
    forwarding-class expedited-forwarding scheduler voice-high;
    forwarding-class network-control scheduler net-control;
    forwarding-class best-effort scheduler best-effort;
  }
}
interfaces {
  ge/0/1/0 {
    scheduler-map ethernet-diffsrv-cos-map;
  }
}
```

---

### Configuring a VLAN Ingress Firewall Filter to Prevent Rogue Devices from Disrupting VoIP Traffic

---

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

#### CLI Quick Configuration

To quickly configure a VLAN firewall filter on **voice-vlan** to prevent rogue devices from using HTTP sessions to mimic the gatekeeper device that manages VoIP traffic, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family ethernet-switching filter ingress-vlan-rogue-block term to-gatekeeper from
destination-address 192.0.2.14
set firewall family ethernet-switching filter ingress-vlan-rogue-block term to-gatekeeper from
destination-port 80
set firewall family ethernet-switching filter ingress-vlan-rogue-block term to-gatekeeper then
accept
```

```

set firewall family ethernet-switching filter ingress-vlan-rogue-block term from-gatekeeper from
source-address 192.0.2.14
set firewall family ethernet-switching filter ingress-vlan-rogue-block term from-gatekeeper from
source-port 80
set firewall family ethernet-switching filter ingress-vlan-rogue-block term from-gatekeeper then
accept
set firewall family ethernet-switching filter ingress-vlan-rogue-block term not-gatekeeper from
destination-port 80
set firewall family ethernet-switching filter ingress-vlan-rogue-block term not-gatekeeper then
count rogue-counter
set firewall family ethernet-switching filter ingress-vlan-rogue-block term not-gatekeeper then
discard
set vlans voice-vlan description "block rogue devices on voice-vlan"
set vlans voice-vlan filter input ingress-vlan-rogue-block

```

**Step-by-Step Procedure** To configure and apply a VLAN firewall filter on **voice-vlan** to prevent rogue devices from using HTTP to mimic the gatekeeper device that manages VoIP traffic:

1. Define the firewall filter **ingress-vlan-rogue-block** to specify filter matching on the traffic you want to permit and restrict:  

```

[edit firewall]
user@switch# set family ethernet-switching filter ingress-vlan-rogue-block

```
2. Define the term **to-gatekeeper** to accept packets that match the destination IP address of the gatekeeper:  

```

[edit firewall family ethernet-switching filter ingress-vlan-rogue-block]
user@switch# set term to-gatekeeper from destination-address 192.0.2.14
user@switch# set term to-gatekeeper from destination-port 80
user@switch# set term to-gatekeeper then accept

```
3. Define the term **from-gatekeeper** to accept packets that match the source IP address of the gatekeeper:  

```

[edit firewall family ethernet-switching filter ingress-vlan-rogue-block]
user@switch# set term from-gatekeeper from source-address 192.0.2.14
user@switch# set term from-gatekeeper from source-port 80
user@switch# set term from-gatekeeper then accept

```
4. Define the term **not-gatekeeper** to ensure all **voice-vlan** traffic on TCP ports is destined for the gatekeeper device:  

```

[edit firewall family ethernet-switching filter ingress-vlan-rogue-block]
user@switch# set term not-gatekeeper from destination-port 80
user@switch# set term not-gatekeeper then count rogue-counter
user@switch# set term not-gatekeeper then discard

```
5. Apply the firewall filter **ingress-vlan-rogue-block** as an input filter to the VLAN interface for the VoIP telephones:  

```

[edit]
user@switch# set vlans voice-vlan description "block rogue devices on voice-vlan"
user@switch# set vlans voice-vlan filter input ingress-vlan-rogue-block

```

**Results** Display the results of the configuration:

```

user@switch# show
firewall {
  family ethernet-switching {
    filter ingress-vlan-rogue-block {
      term to-gatekeeper {
        from {

```

```
        destination-address 192.0.2.14/32
        destination-port 80;
    }
    then {
        accept;
    }
}
term from-gatekeeper {
    from {
        source-address 192.0.2.14/32
        source-port 80;
    }
    then {
        accept;
    }
}
term not-gatekeeper {
    from {
        destination-port 80;
    }
    then {
        count rogue-counter;
        discard;
    }
}
}
vlands {
    voice-vlan {
        description "block rogue devices on voice-vlan";
        filter {
            input ingress-vlan-rogue-block;
        }
    }
}
```

### Configuring a VLAN Firewall Filter to Count, Monitor, and Analyze Egress Traffic on the Employee VLAN

---

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

#### CLI Quick Configuration

A firewall filter is configured and applied to VLAN interfaces to filter **employee-vlan** egress traffic. Employee traffic destined for the corporate subnet is accepted but not monitored. Employee traffic destined for the Web is counted and analyzed.

To quickly configure and apply a VLAN firewall filter, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-corp
from destination-address 192.0.2.16/28
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-corp
then accept
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-web
from destination-port 80
```



```

set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-web
then count employee-web-counter
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-web
then analyzer employee-monitor
set vlans employee-vlan description "filter at egress VLAN to count and analyze employee to
Web traffic"
set vlans employee-vlan filter output egress-vlan-watch-employee

```

**Step-by-Step Procedure** To configure and apply an egress port firewall filter to count and analyze **employee-vlan** traffic that is destined for the Web:

1. Define the firewall filter **egress-vlan-watch-employee**:  

```

[edit firewall]
user@switch# set family ethernet-switching filter egress-vlan-watch-employee

```
2. Define the term **employee-to-corp** to accept but not monitor all **employee-vlan** traffic destined for the corporate subnet:  

```

[edit firewall family ethernet-switching filter egress-vlan-watch-employee]
user@switch# set term employee-to-corp from destination-address 192.0.2.16/28
user@switch# set term employee-to-corp then accept

```
3. Define the term **employee-to-web** to count and monitor all **employee-vlan** traffic destined for the Web:  

```

[edit firewall family ethernet-switching filter egress-vlan-watch-employee]
user@switch# set term employee-to-web from destination-port 80
user@switch# set term employee-to-web then count employee-web-counter
user@switch# set term employee-to-web then analyzer employee-monitor

```



**NOTE:** See *Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches* for information about configuring the **employee-monitor** analyzer.

4. Apply the firewall filter **egress-vlan-watch-employee** as an output filter to the port interfaces for the VoIP telephones:  

```

[edit]
user@switch# set vlans employee-vlan description "filter at egress VLAN to count and
analyze employee to Web traffic"
user@switch# set vlans employee-vlan filter output egress-vlan-watch-employee

```

**Results** Display the results of the configuration:

```

user@switch# show
firewall {
  family ethernet-switching {
    filter egress-vlan-watch-employee {
      term employee-to-corp {
        from {
          destination-address 192.0.2.16/28
        }
        then {
          accept;
        }
      }
    }
  }
}

```

```
term employee-to-web {
  from {
    destination-port 80;
  }
  then {
    count employee-web-counter;
    analyzer employee-monitor;
  }
}
}
}
}
vlands {
  employee-vlan {
    description "filter at egress VLAN to count and analyze employee to Web traffic";
    filter {
      output egress-vlan-watch-employee;
    }
  }
}
```

### Configuring a VLAN Firewall Filter to Restrict Guest-to-Employee Traffic and Peer-to-Peer Applications on the Guest VLAN

---

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

#### CLI Quick Configuration

In the following example, the first filter term permits guests to talk with other guests but not employees on **employee-vlan**. The second filter term allows guests Web access but prevents them from using peer-to-peer applications on **guest-vlan**.

To quickly configure a VLAN firewall filter to restrict guest-to-employee traffic, blocking guests from talking with employees or employee hosts on **employee-vlan** or attempting to use peer-to-peer applications on **guest-vlan**, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family ethernet-switching filter ingress-vlan-limit-guest term guest-to-guest from
destination-address 192.0.2.33/28
set firewall family ethernet-switching filter ingress-vlan-limit-guest term guest-to-guest then
accept
set firewall family ethernet-switching filter ingress-vlan-limit-guest term
no-guest-employee-no-peer-to-peer from destination-mac-address 00.05.85.00.00.DF
set firewall family ethernet-switching filter ingress-vlan-limit-guest term
no-guest-employee-no-peer-to-peer then accept
set vlans guest-vlan description "restrict guest-to-employee traffic and peer-to-peer applications
on guest VLAN"
set vlans guest-vlan filter input ingress-vlan-limit-guest
```

**Step-by-Step Procedure** To configure and apply a VLAN firewall filter to restrict guest-to-employee traffic and peer-to-peer applications on **guest-vlan**:

1. Define the firewall filter **ingress-vlan-limit-guest**:  

```
[edit firewall]
set firewall family ethernet-switching filter ingress-vlan-limit-guest
```
2. Define the term **guest-to-guest** to permit guests on the **guest-vlan** to talk with other guests but not employees on the **employee-vlan**:  

```
[edit firewall family ethernet-switching filter ingress-vlan-limit-guest]
user@switch# set term guest-to-guest from destination-address 192.0.2.33/28
user@switch# set term guest-to-guest then accept
```
3. Define the term **no-guest-employee-no-peer-to-peer** to allow guests on **guest-vlan** Web access but prevent them from using peer-to-peer applications on the **guest-vlan**.



**NOTE:** The destination-mac-address is the default gateway, which for any host in a VLAN is the next-hop router.

- ```
[edit firewall family ethernet-switching filter ingress-vlan-limit-guest]
user@switch# set term no-guest-employee-no-peer-to-peer from destination-mac-address 00.05.85.00.00.DF
user@switch# set term no-guest-employee-no-peer-to-peer then accept
```
4. Apply the firewall filter **ingress-vlan-limit-guest** as an input filter to the interface for **guest-vlan** :  

```
[edit]
user@switch# set vlans guest-vlan description "restrict guest-to-employee traffic and peer-to-peer applications on guest VLAN"
user@switch# set vlans guest-vlan filter input ingress-vlan-limit-guest
```

**Results** Display the results of the configuration:

```
user@switch# show
firewall {
  family ethernet-switching {
    filter ingress-vlan-limit-guest {
      term guest-to-guest {
        from {
          destination-address 192.0.2.33/28;
        }
        then {
          accept;
        }
      }
      term no-guest-employee-no-peer-to-peer {
        from {
          destination-mac-address 00.05.85.00.00.DF;
        }
        then {
          accept;
        }
      }
    }
  }
}
```

```

    }
  }
}
vlands {
  guest-vlan {
    description "restrict guest-to-employee traffic and peer-to-peer applications on
    guest VLAN";
    filter {
      input ingress-vlan-limit-guest;
    }
  }
}
}

```

### Configuring a Router Firewall Filter to Give Priority to Egress Traffic Destined for the Corporate Subnet

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

#### CLI Quick Configuration

To quickly configure a firewall filter for a routed port (Layer 3 uplink module) to filter **employee-vlan** traffic, giving highest forwarding-class priority to traffic destined for the corporate subnet, copy the following commands and paste them into the switch terminal window:

```

[edit]
set firewall family inet filter egress-router-corp-class term corp-expedite from destination-address
192.0.2.16/28
set firewall family inet filter egress-router-corp-class term corp-expedite then forwarding-class
expedited-forwarding
set firewall family inet filter egress-router-corp-class term corp-expedite then loss-priority low
set firewall family inet filter egress-router-corp-class term not-to-corp then accept
set interfaces ge-0/1/0 description "filter at egress router to expedite destined for corporate
network"
set ge-0/1/0 unit 0 family inet address 103.104.105.1
set interfaces ge-0/1/0 unit 0 family inet filter output egress-router-corp-class

```

#### Step-by-Step Procedure

To configure and apply a firewall filter to a routed port (Layer 3 uplink module) to give highest priority to **employee-vlan** traffic destined for the corporate subnet:

1. Define the firewall filter **egress-router-corp-class**:

```

[edit]
user@switch# set firewall family inet filter egress-router-corp-class

```

2. Define the term **corp-expedite**:

```

[edit firewall]
user@switch# set family inet filter egress-router-corp-class term corp-expedite from
destination-address 192.0.2.16/28
user@switch# set family inet filter egress-router-corp-class term corp-expedite then
forwarding-class expedited-forwarding
user@switch# set family inet filter egress-router-corp-class term corp-expedite then
loss-priority low

```

3. Define the term **not-to-corp**:

```

[edit firewall]
user@switch# set family inet filter egress-router-corp-class term not-to-corp then accept

```

4. Apply the firewall filter **egress-router-corp-class** as an output filter for the port on the switch's uplink module, which provides a Layer 3 connection to a router:

```
[edit interfaces]
user@switch# set ge-0/1/0 description "filter at egress router to expedite employee traffic
destined for corporate network"
user@switch# set ge-0/1/0 unit 0 family inet address 103.104.105.1
user@switch# set ge-0/1/0 unit 0 family inet filter output egress-router-corp-class
```

**Results** Display the results of the configuration:

```
user@switch# show
firewall {
  family inet {
    filter egress-router-corp-class {
      term corp-expedite {
        from {
          destination-address 192.0.2.16/28;
        }
        then {
          forwarding-class expedited-forwarding;
          loss-priority low;
        }
      }
      term not-to-corp {
        then {
          accept;
        }
      }
    }
  }
}
interfaces {
  ge-0/1/0 {
    unit 0 {
      description "filter at egress router interface to expedite employee traffic destined
for corporate network";
      family inet {
        source-address 103.104.105.1
        filter {
          output egress-router-corp-class;
        }
      }
    }
  }
}
```

## Verification

To confirm that the firewall filters are working properly, perform the following tasks:

- [Verifying that Firewall Filters and Policers are Operational on page 4789](#)
- [Verifying that Schedulers and Scheduler-Maps are Operational on page 4790](#)

### *Verifying that Firewall Filters and Policers are Operational*

**Purpose** Verify the operational state of the firewall filters and policers that are configured on the switch.

**Action** Use the operational mode command:

```
user@switch> show firewall
Filter: ingress-port-voip-class-limit-tcp-icmp
Counters:
Name                                     Packets
icmp-counter                             0
tcp-counter                              0
Policers:
Name                                     Packets
icmp-connection-policer                  0
tcp-connection-policer                    0

Filter: ingress-vlan-rogue-block

Filter: egress-vlan-watch-employee
Counters:
Name                                     Packets
employee-web-counter                       0
```

**Meaning** The **show firewall** command displays the names of the firewall filters, policers, and counters that are configured on the switch. The output fields show byte and packet counts for all configured counters and the packet count for all policers.

#### *Verifying that Schedulers and Scheduler-Maps are Operational*

**Purpose** Verify that schedulers and scheduler-maps are operational on the switch.

**Action** Use the operational mode command:

```
user@switch> show class-of-service scheduler-map

Scheduler map: default, Index: 2

Scheduler: default-be, Forwarding class: best-effort, Index: 20
  Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent,
  Priority: low
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low            non-TCP   1      default-drop-profile
    Low            TCP       1      default-drop-profile
    High           non-TCP   1      default-drop-profile
    High            TCP       1      default-drop-profile

Scheduler: default-nc, Forwarding class: network-control, Index: 22
  Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,
  Priority: low
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low            non-TCP   1      default-drop-profile
    Low            TCP       1      default-drop-profile
    High           non-TCP   1      default-drop-profile
    High            TCP       1      default-drop-profile
Scheduler map: ethernet-diffsrv-cos-map, Index: 21657

Scheduler: best-effort, Forwarding class: best-effort, Index: 61257
  Transmit rate: remainder, Rate Limit: none, Buffer size: 75 percent,
  Priority: low
  Drop profiles:
```

Loss priority	Protocol	Index	Name
Low	non-TCP	1	<default-drop-profile>
Low	TCP	1	<default-drop-profile>
High	non-TCP	1	<default-drop-profile>
High	TCP	1	<default-drop-profile>

Scheduler: voice-high, Forwarding class: expedited-forwarding, Index: 3123  
 Transmit rate: remainder, Rate Limit: none, Buffer size: 15 percent,  
 Priority: high

Drop profiles:

Loss priority	Protocol	Index	Name
Low	non-TCP	1	<default-drop-profile>
Low	TCP	1	<default-drop-profile>
High	non-TCP	1	<default-drop-profile>
High	TCP	1	<default-drop-profile>

Scheduler: net-control, Forwarding class: network-control, Index: 2451  
 Transmit rate: remainder, Rate Limit: none, Buffer size: 10 percent,  
 Priority: high

Drop profiles:

Loss priority	Protocol	Index	Name
Low	non-TCP	1	<default-drop-profile>
Low	TCP	1	<default-drop-profile>
High	non-TCP	1	<default-drop-profile>
High	TCP	1	<default-drop-profile>

**Meaning** Displays statistics about the configured schedulers and schedulers-maps.

- Related Documentation**
- *Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches*
  - [Example: Configuring CoS on EX Series Switches on page 2075](#)
  - [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
  - [Configuring Firewall Filters \(J-Web Procedure\) on page 4813](#)
  - [Configuring Policers to Control Traffic Rates \(CLI Procedure\) on page 4818](#)
  - [Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704](#)
  - [\[edit firewall\] Configuration Statement Hierarchy on EX Series Switches on page 336](#)

## Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch

You can configure a firewall filter on a management interface on an EX Series switch to filter ingress or egress traffic on the management interface on the switch. You can use utilities such as SSH or Telnet to connect to the management interface over the network and then use management protocols such as SNMP to gather statistical data from the switch.

This example discusses how to configure a firewall filter on a management interface to filter SSH packets egressing from an EX Series switch:

- [Requirements on page 4792](#)
- [Overview and Topology on page 4792](#)

- [Configuration on page 4792](#)
- [Verification on page 4794](#)

## Requirements

This example uses the following hardware and software components:


- One EX Series switch and one management PC
- Junos OS Release 10.4 or later for EX Series switches

## Overview and Topology

In this example, a management PC establishes an SSH connection with the management interface on a switch to remotely manage the switch. The IP address configured for the management interface is 10.204.33.103/20. A firewall filter is configured on the management interface to count the number of packets egressing from a source SSH port on the management interface. When the management PC establishes the SSH session with the management interface, the management interface returns SSH packets to the management PC to confirm that the session is established. These SSH packets are filtered based on the match condition specified in the firewall filter before they are forwarded to the management PC. As these packets are generated from the source SSH port on the management interface, they fulfill the match condition specified for the management interface. The number of matched SSH packets provides a count of the number of packets that have traversed the management interface. A system administrator can use this information to monitor the management traffic and take any action if required.

[Figure 70 on page 4792](#) shows the topology for this example in which a management PC establishes an SSH connection with the switch.

**Figure 70: SSH Connection From a Management PC to an EX Series Switch**



## Configuration

To configure a firewall filter on a management interface, perform these tasks:

### CLI Quick Configuration

To quickly create and configure a firewall filter on the management interface to filter SSH packets egressing from the management interface, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family inet filter mgmt_fil1 term t1 from source-port ssh
set firewall family inet filter mgmt_fil1 term t1 then count c1
set firewall family inet filter mgmt_fil1 term t2 then accept
set interfaces me0 unit 0 family inet filter output mgmt_fil1
```



**Step-by-Step Procedure** To configure a firewall filter on the management interface to filter SSH packets:

1. Configure the firewall filter that matches SSH packets from the source port:

```
[edit]
user@switch# set firewall family inet filter (Firewall Filters) mgmt_fil1 term t1 from
source-port ssh
user@switch# set firewall family inet filter mgmt_fil1 term t1 then count c1
user@switch# set firewall family inet filter mgmt_fil1 term t2 then accept
```

These statements set a counter `c1` to count the number of SSH packets that egress from the source SSH interface on the management interface.

2. Set the firewall filter for the management interface:

```
[edit]
user@switch# set interfaces me0 unit 0 family inet filter output mgmt_fil1
```



**NOTE:** You can also set the firewall filter for a VME interface.

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
interfaces {
  me0 {
    unit 0 {
      family inet {
        filter {
          output mgmt_fil1;
        }
        address 10.93.54.6/24;
      }
    }
  }
}

firewall {
  family inet {
    filter mgmt_fil1 {
      term t1 {
        from {
          source-port ssh;
        }
        then count c1;
      }
      term t2 {
        then accept;
      }
    }
  }
}
```

## Verification

---

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Firewall Filter Is Configured on a Management Interface on page 4794](#)

### *Verifying That the Firewall Filter Is Configured on a Management Interface*

**Purpose** Verify that the firewall filter has been enabled on the management interface on the switch.

**Action** 1. Verify that the firewall filter is applied to the management interface:

```
[edit]
user@switch# show interfaces me0
unit 0 {
    family inet {
        filter {
            output mgmt_fil1;
        }
        address 10.204.33.103/20;
    }
}
```

2. Check the counter value that is associated with the firewall filter:

```
user@switch> show firewall
Filter: mgmt_fil1
Counters:
Name                               Bytes          Packets
c1                                 0              0
```

3. From the management PC, establish a secure shell session with the switch:

```
[user@management-pc ~]$ ssh user@10.204.33.103
```

4. Check counter values after SSH packets are generated from the switch in response to the secure shell session request by the management PC:

```
user@switch> show firewall
Filter: mgmt_fil1
Counters:
Name                               Bytes          Packets
c1                                 3533           23
```

**Meaning** The output indicates that the firewall filter has been applied to the management interface and the counter value indicates that 23 SSH packets were generated from the switch.

**Related Documentation**

- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)

## Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches

Administrators can configure filter-based forwarding on an EX Series switch by using a firewall filter to forward matched traffic to a specific virtual routing instance.

This example describes how to set up filter-based forwarding:

- [Requirements on page 4795](#)
- [Overview and Topology on page 4795](#)
- [Configuration on page 4795](#)
- [Verification on page 4797](#)

### Requirements

This example uses the following software and hardware components:

- One EX Series switch
- Junos OS Release 9.4 or later for EX Series switches

### Overview and Topology

In this example, traffic from one application server that is destined for a different application server is matched by a firewall filter based on the IP address. Any matching packets are routed to a particular virtual routing instance that first sends all traffic to a security device, then forwards it to the designated destination address.

### Configuration

To configure filter-based forwarding:

#### CLI Quick Configuration

To quickly create and configure filter-based forwarding, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.1/24
set interfaces ge-0/0/3 unit 0 family inet address 10.1.3.1/24
set firewall family inet filter fil term t1 from source-address 1.1.1.1/32
set firewall family inet filter fil term t1 from protocol tcp
set interfaces ge-0/0/0 unit 0 family inet filter input fil
set routing-instances vrf01 instance-type virtual-router
set routing-instances vrf01 interface ge-0/0/1.0
set routing-instances vrf01 interface ge-0/0/3.0
set routing-instances vrf01 routing-options static route 12.34.56.0/24 next-hop 10.1.3.254
set firewall family inet filter fil term t1 then routing-instance vrf01
```

#### Step-by-Step Procedure

To configure filter-based forwarding:

1. Create interfaces to the application servers:

```
[edit]
user@switch# set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.1/24
user@switch# set interfaces ge-0/0/3 unit 0 family inet address 10.1.3.1/24
```

2. Create a firewall filter that matches the correct source address:

- ```
[edit]
user@switch# set firewall family inet filter fil term t1 from source-address 1.1.1.1/32
user@switch# set firewall family inet filter fil term t1 from protocol tcp
```
3. Associate the filter with the source application server's interface:
- ```
[edit]
user@switch# set interfaces ge-0/0/0 unit 0 family inet filter input fil
```
4. Create a virtual router:
- ```
[edit]
user@switch# set routing-instances vrf01 instance-type virtual-router
```
5. Associate the interfaces with the virtual router:
- ```
[edit]
user@switch# set routing-instances vrf01 interface ge-0/0/1.0
user@switch# set routing-instances vrf01 interface ge-0/0/3.0
```
6. Configure the routing information for the virtual routing instance:
- ```
[edit]
user@switch# set routing-instances vrf01 routing-options static route 12.34.56.0/24
next-hop 10.1.3.254
```
7. Set the filter to forward packets to the virtual router you created:
- ```
[edit]
user@switch# set firewall family inet filter fil term t1 then routing-instance vrf01
```

### Results

Check the results of the configuration:

```
user@switch> show configuration
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        filter {
          input fil;
        }
        address 10.1.0.1/24;
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.3.1/24;
      }
    }
  }
}
firewall {
  family inet {
    filter fil {
      term t1 {
        from {
          source-address {
            1.1.1.1/32;
          }
        }
        protocol tcp;
      }
    }
  }
}
```

```

    }
    then {
        routing-instance vrf01;
    }
}
}
}
}
}
routing-instances {
    vrf01 {
        instance-type virtual-router;
        interface ge-0/0/1.0;
        interface ge-0/0/3.0;
        routing-options {
            static {
                route 12.34.56.0/24 next-hop 10.1.3.254;
            }
        }
    }
}
}
}

```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That Filter-Based Forwarding Was Configured on page 4797](#)

#### *Verifying That Filter-Based Forwarding Was Configured*

**Purpose** Verify that filter-based forwarding was properly enabled on the switch.

**Action** 1. Use the **show interfaces filters** command:

```

user@switch> show interfaces filters ge-0/0/0.0
Interface      Admin Link Proto Input Filter      Output Filter
ge-0/0/0.0     up    down inet  fil

```

2. Use the **show route forwarding-table** command:

```

user@switch> show route forwarding-table

```

Routing table: default.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	user	1	0:12:f2:21:cf:0	ucst	331	4	me0.0
default	perm	0		rjct	36	3	
0.0.0.0/32	perm	0		dscd	34	1	
10.1.0.0/24	ifdn	0		rslv	613	1	ge-0/0/0.0
10.1.0.0/32	iddn	0	10.1.0.0	recv	611	1	ge-0/0/0.0
10.1.0.1/32	user	0		rjct	36	3	
10.1.0.1/32	intf	0	10.1.0.1	loc1	612	2	
10.1.0.1/32	iddn	0	10.1.0.1	loc1	612	2	
10.1.0.255/32	iddn	0	10.1.0.255	bcst	610	1	ge-0/0/0.0
10.1.1.0/26	ifdn	0		rslv	583	1	vlan.0
10.1.1.0/32	iddn	0	10.1.1.0	recv	581	1	vlan.0
10.1.1.1/32	user	0		rjct	36	3	
10.1.1.1/32	intf	0	10.1.1.1	loc1	582	2	
10.1.1.1/32	iddn	0	10.1.1.1	loc1	582	2	
10.1.1.63/32	iddn	0	10.1.1.63	bcst	580	1	vlan.0

```

255.255.255.255/32 perm      0                               bcst      32      1

Routing table: vrf01.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm      0                               rjct     559     2
0.0.0.0/32       perm      0                               dscd     545     1
10.1.3.0/24      ifdn      0                               rslv     617     1 ge-0/0/3.0
10.1.3.0/32      iddn      0 10.1.3.0          recv     615     1 ge-0/0/3.0
10.1.3.1/32      user      0                               rjct     559     2
10.1.3.1/32      intf      0 10.1.3.1          locl     616     2
10.1.3.1/32      iddn      0 10.1.3.1          locl     616     2
10.1.3.255/32    iddn      0 10.1.3.255        bcst     614     1 ge-0/0/3.0
224.0.0.0/4      perm      0                               mdsc     546     1
224.0.0.1/32     perm      0 224.0.0.1         mcst     529     1
255.255.255.255/32 perm      0                               bcst     543     1

Routing table: default.iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm      0                               rjct     60      1

Routing table: vrf01.iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm      0                               rjct     600     1

```

**Meaning** The output indicates that the filter was created on the interface and that the virtual routing instance is forwarding matching traffic to the correct IP address.

- Related Documentation**
- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
  - [Configuring Static Routing \(CLI Procedure\) on page 3422](#)
  - [Configuring Static Routing \(J-Web Procedure\) on page 3422](#)
  - [Understanding Filter-Based Forwarding for EX Series Switches on page 4770](#)

## Example: Applying Firewall Filters to Multiple Supplicants on Interfaces Enabled for 802.1X or MAC RADIUS Authentication

On EX Series switches, firewall filters that you apply to interfaces enabled for 802.1X or MAC RADIUS authentication are dynamically combined with the per-user policies sent to the switch from the RADIUS server. The switch uses internal logic to dynamically combine the interface firewall filter with the user policies from the RADIUS server and create an individualized policy for each of the multiple users or nonresponsive hosts that are authenticated on the interface.

This example describes how dynamic firewall filters are created for multiple supplicants on an 802.1X-enabled interface (the same principles shown in this example apply to interfaces enabled for MAC RADIUS authentication):

- [Requirements on page 4799](#)
- [Overview and Topology on page 4799](#)

- [Configuration on page 4801](#)
- [Verification on page 4803](#)

---

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.5 or later for EX Series switches
- One EX Series switch
- One RADIUS authentication server. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you apply firewall filters to an interface for use with multiple supplicants, be sure you have:

- Set up a connection between the switch and the RADIUS server. See [“Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch” on page 1843](#).
- Configured 802.1X authentication on the switch, with the authentication mode for interface **ge-0/0/2** set to **multiple**. See [“Configuring 802.1X Interface Settings \(CLI Procedure\)” on page 1908](#) and [“Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on an EX Series Switch” on page 1852](#).
- Configured users on the RADIUS authentication server.

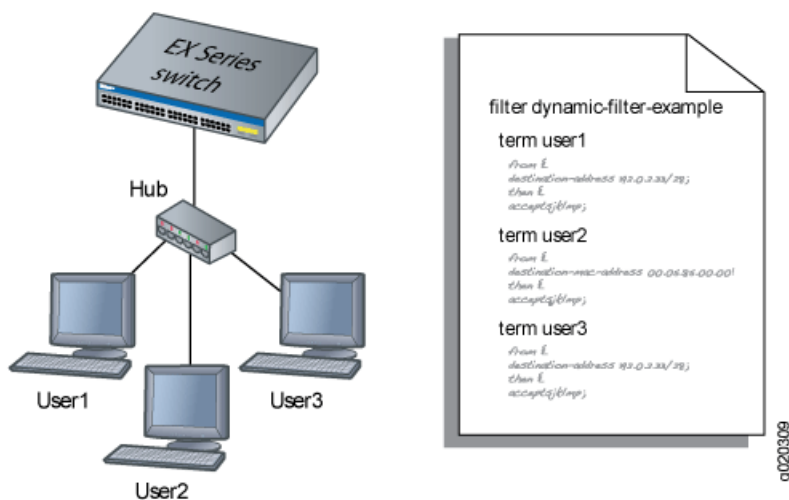
---

## Overview and Topology

When the 802.1X configuration on an interface is set to multiple supplicant mode, the system dynamically combines interface firewall filter with the user policies sent to the switch from the RADIUS server during authentication and creates separate terms for each user. Because there are separate terms for each user authenticated on the interface, you can, as shown in this example, use counters to view the activities of individual users that are authenticated on the same interface.

When a new user (or a nonresponsive host) is authenticated on an interface, the system adds a term to the firewall filter associated with the interface, and the term (policy) for each user is associated with the MAC address of the user. The term for each user is based on the user-specific filters set on the RADIUS server and the filters configured on the interface. For example, as shown in [Figure 18 on page 1895](#), when User1 is authenticated by the EX Series switch, the system creates the firewall filter **dynamic-filter-example**. When User2 is authenticated, another term is added to the firewall filter, and so on.

Figure 71: Conceptual Model: Dynamic Filter Updated for Each New User



This is a conceptual model of the internal process—you cannot access or view the dynamic filter.




**NOTE:** If the firewall filter on the interface is modified after the user (or nonresponsive host) is authenticated, the modifications are not reflected in the dynamic filter unless the user is reauthenticated.

In this example, you configure a firewall filter to count the requests made by each endpoint authenticated on interface **ge-0/0/2** to the file server, which is located on subnet **192.0.2.16/28**, and set policer definitions to rate limit the traffic. [Figure 19 on page 1896](#) shows the network topology for this example.



Figure 72: Multiple Supplicants on an 802.1X-Enabled Interface Connecting to a File Server



g020308

### Configuration

To configure firewall filters for multiple supplicants on 802.1X-enabled interfaces:

- [Configuring Firewall Filters on Interfaces with Multiple Supplicants on page 4801](#)

#### Configuring Firewall Filters on Interfaces with Multiple Supplicants

#### CLI Quick Configuration

To quickly configure firewall filters for multiple supplicants on an 802.1X-enabled interface copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols dot1x authenticator interface ge-0/0/2 supplicant multiple
set firewall family ethernet-switching filter filter1 term term1 from destination-address 192.0.2.16/28
set firewall policer p1 if-exceeding bandwidth-limit 1m
set firewall policer p1 if-exceeding burst-size-limit 1k
set firewall family ethernet-switching filter filter1 term term1 then count counter1
set firewall family ethernet-switching filter filter1 term term2 then policer p1
```

#### Step-by-Step Procedure

To configure firewall filters on an interface enabled for multiple supplicants:

1. Configure interface **ge-0/0/2** for multiple supplicant mode authentication:
 

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/2 supplicant multiple
```
2. Set policer definition:

```
user@switch# show policer p1 |display set
set firewall policer p1 if-exceeding bandwidth-limit 1m
set firewall policer p1 if-exceeding burst-size-limit 1k
set firewall policer p1 then discard
```

3. Configure a firewall filter to count packets from each user and a policer that limits the traffic rate. As each new user is authenticated on the multiple supplicant interface, this filter term will be included in the dynamically created term for the user:

```
[edit firewall family ethernet-switching]
user@switch# set filter filter1 term term1 from destination-address 192.0.2.16/28
user@switch# set filter filter1 term term1 then count counter1
user@switch# set filter filter1 term term2 then policer p1
```

**Results** Check the results of the configuration:

```
user@switch> show configuration
```

```
firewall {
  family ethernet-switching {
    filter filter1 {
      term term1 {
        from {
          destination-address {
            192.0.2.16/28;
          }
        }
        then count counter1;
      }
      term term2 {
        from {
          destination-address {
            192.0.2.16/28;
          }
        }
        then policer p1;
      }
    }
  }
}
policer p1 {
  if-exceeding {
    bandwidth-limit 1m;
    burst-size-limit 1k;
  }
  then discard;
}
protocols {
  dot1x {
    authenticator
    interface ge-0/0/2 {
      supplicant multiple;
    }
  }
}
```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying Firewall Filters on Interfaces with Multiple Supplicants on page 4803](#)

### *Verifying Firewall Filters on Interfaces with Multiple Supplicants*

**Purpose** Verify that firewall filters are functioning on the interface with multiple supplicants.

- Action**
1. Check the results with one user authenticated on the interface. In this case, the user is authenticated on **ge-0/0/2**:  
  

```
user@switch> show dot1x firewall
Filter: dot1x_ge-0/0/2
Counters
counter1_dot1x_ge-0/0/2_user1 100
```
  2. When a second user, User2, is authenticated on the same interface, **ge-0/0/2**, you can verify that the filter includes the results for both of the users authenticated on the interface:  
  

```
user@switch> show dot1x firewall
Filter: dot1x-filter-ge-0/0/0
Counters
counter1_dot1x_ge-0/0/2_user1 100
counter1_dot1x_ge-0/0/2_user2 400
```

**Meaning** The results displayed by the **show dot1x firewall** command output reflect the dynamic filter created with the authentication of each new user. User1 accessed the file server located at the specified destination address 100 times, while User2 accessed the same file server 400 times.

- Related Documentation**
- [Example: Applying a Firewall Filter to 802.1X-Authenticated Supplicants Using RADIUS Server Attributes on an EX Series Switch on page 1887](#)
  - [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
  - [Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 1910](#)

## Configuration Tasks

- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
- [Configuring Firewall Filters \(J-Web Procedure\) on page 4813](#)
- [Configuring Policers to Control Traffic Rates \(CLI Procedure\) on page 4818](#)
- [Assigning Multifield Classifiers in Firewall Filters to Specify Packet-Forwarding Behavior \(CLI Procedure\) on page 4821](#)
- [Configuring Routing Policies \(J-Web Procedure\) on page 4822](#)
- [Configuring Tricolor Marking Policers on page 4827](#)

## Configuring Firewall Filters (CLI Procedure)

You configure firewall filters on EX Series switches to control traffic that enters ports on the switch or enters and exits VLANs on the network and Layer 3 (routed) interfaces. To configure a firewall filter you must configure the filter and then apply it to a port, VLAN, or Layer 3 interface.

This topic describes:

- [Configuring a Firewall Filter on page 4804](#)
- [Configuring a Term Specifically for IPv4 or IPv6 Traffic on page 4808](#)
- [Applying a Firewall Filter to a Port on a Switch on page 4809](#)
- [Applying a Firewall Filter to a Management Interface on a Switch on page 4810](#)
- [Applying a Firewall Filter to a VLAN on a Network on page 4811](#)
- [Applying a Firewall Filter to a Layer 3 \(Routed\) Interface on page 4812](#)

### Configuring a Firewall Filter

---

Before you can apply a firewall filter to a port, VLAN, or Layer 3 interface, you must configure a firewall filter with the required details, such as type of family for the firewall filter, firewall filter name, and match conditions. A match condition in the firewall filter configuration can contain multiple terms that define the criteria for the match condition. For each term, you must specify an action to be performed if a packet matches the conditions in the term. For information on different match conditions and actions, see [“Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches” on page 4704](#).

To configure a firewall filter:

1. Configure the family address type for the firewall filter:

- For a firewall filter that is applied to a port or VLAN, specify the family address type **ethernet-switching** to filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets, for example:

```
[edit firewall]
user@switch# set family ethernet-switching
```

- For a firewall filter that is applied to a Layer 3 (routed) interface:

- To filter IPv4 packets, specify the family address type **inet**, for example:

```
[edit firewall]
user@switch# set family inet
```

- To filter IPv6 packets, specify the family address type **inet6**, for example:

```
[edit firewall]
user@switch# set family inet6
```



**NOTE:** You can configure firewall filters for both IPv4 and IPv6 traffic on the same Layer 3 interface.

2. Specify the filter name:

```
[edit firewall family ethernet-switching]
user@switch# set filter ingress-port-filter
```

The filter name can contain letters, numbers, and hyphens (-) and can have a maximum of 64 characters. Each filter name must be unique.

3. If you want to apply a firewall filter to multiple interfaces and name individual firewall counters specific to each interface, configure the **interface-specific** option:

```
[edit firewall family ethernet-switching filter ingress-port-filter]
user@switch# set interface-specific
```

4. Specify a term name:

```
[edit firewall family ethernet-switching filter ingress-port-filter]
user@switch# set term term-one
```

The term name can contain letters, numbers, and hyphens (-) and can have a maximum of 64 characters.

A firewall filter can contain one or more terms. Each term name must be unique within a filter.

**NOTE:**

The maximum number of terms allowed per firewall filter for EX Series switches is:

- 512 for EX2200 switches
- 1,436 for EX3300 switches



**NOTE:** On EX3300 switches, if you add and delete filters with a large number of terms (on the order of 1000 or more) in the same commit operation, not all the filters are installed. You must add filters in one commit operation, and delete filters in a separate commit operation.

- 7,168 for EX3200 and EX4200 switches
- On EX4300 switches, following are the number of terms supported for ingress and egress traffic, for firewall filters configured on a port, VLAN and Layer 3 interface:
  - For ingress traffic:
    - 3,500 terms for firewall filters configured on a port
    - 3,500 terms for firewall filters configured on a VLAN
    - 7,000 terms for firewall filters configured on Layer 3 interfaces for IPv4 traffic
    - 3,500 terms for firewall filters configured on Layer 3 interfaces for IPv6 traffic
  - For egress traffic:
    - 512 terms for firewall filters configured on a port
    - 256 terms for firewall filters configured on a VLAN
    - 512 terms for firewall filters configured on Layer 3 interfaces for IPv4 traffic
    - 512 terms for firewall filters configured on Layer 3 interfaces for IPv6 traffic



**NOTE:** You can configure these maximum number of terms only when you configure one type of firewall filter (Port, VLAN, or Router (Layer 3) firewall filter) on the switch, and when storm control is not enabled on all interfaces in the switch.

- 1,200 for EX4500 and EX4550 switches

- 1,400 for EX6200 switches
- 32,768 for EX8200 switches

If you attempt to configure a firewall filter that exceeds these limits, the switch returns an error message when you commit the configuration.

5. In each firewall filter term, specify the match conditions to use to match components of a packet.

To specify match conditions to match on packets that contain a specific source address and source port—for example:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one]
user@switch# set from source-address 192.0.2.14
user@switch# set from source-port 80
```

You can specify one or more match conditions in a single **from** statement. For a match to occur, the packet must match all the conditions in the term.

The **from** statement is optional, but if included in a term, the **from** statement cannot be empty. If you omit the **from** statement, all packets are considered to match.

6. In each firewall filter term, specify the action to take if the packet matches all the conditions in that term.

You can specify an action and/or action modifiers:

- To specify a filter action, for example, to discard packets that match the conditions of the filter term:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one]
user@switch# set then discard
```

You can specify no more than one action per filter term.

- To specify an action modifier, for example, to count and classify packets in a forwarding class:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one]
user@switch# set then count counter-one
user@switch# set then forwarding-class expedited-forwarding
```

In a **then** statement, you can specify the following action modifiers:

- **analyzer analyzer-name**—Mirror port traffic to a specified destination port or VLAN that is connected to a protocol analyzer application. An **analyzer** must be configured under the **ethernet-switching** family address type. See *Configuring Port Mirroring to Analyze Traffic (CLI Procedure)*.
- **count counter-name**—Count the number of packets that pass this filter term.



**NOTE:** We recommend that you configure a counter for each term in a firewall filter, so that you can monitor the number of packets that match the conditions specified in each filter term.

- **forwarding-class *class***—Classify packets in a forwarding class.
- **loss-priority *priority***—Set the priority for dropping a packet.
- **policer *policer-name***—Apply rate limiting to the traffic.
- **interface *interface-name***—Forward the traffic to the specified interface, bypassing the switching lookup.
- **log**—Log the packet's header information in the Routing Engine.

If you omit the **then** statement or do not specify an action, packets that match all the conditions in the **from** statement are accepted. However, you must always explicitly configure an action and/or action modifier in the **then** statement. You can include no more than one action, but you can use any combination of action modifiers. For an action or action modifier to take effect, all conditions in the **from** statement must match.



**NOTE:** Implicit discard is also applicable to a firewall filter applied to the loopback interface, lo0.

On Juniper Networks EX8200 Ethernet Switches, if an implicit or explicit discard action is configured on a loopback interface for IPv4 traffic, next hop resolve packets are accepted and allowed to pass through the switch. However, for IPv6 traffic, you must explicitly configure a rule to allow the next hop IPv6 resolve packets to pass through the switch.

---

### Configuring a Term Specifically for IPv4 or IPv6 Traffic

---

To configure a term in a firewall filter configuration specifically for IPv4 traffic:

1. Verify that neither **ether-type ipv6** nor **ip-version ipv6** is specified in the term in the configuration. By default, a configuration that does not contain either **ether-type ipv6** or **ip-version ipv6** in a term applies to IPv4 traffic.
2. (Optional) Perform one of these tasks:
  - Define **ether-type ipv4** in a term in the configuration.
  - Define **ip-version ipv4** in a term in the configuration.
  - Define both **ether-type ipv4** and **ip-version ipv4** in a term in the configuration.
3. Verify that neither **ether-type ipv6** nor **ip-version ipv6** is specified in a term in the configuration—by default, a configuration that does not contain either **ether-type ipv6** or **ip-version ipv6** in a term applies to IPv4 traffic if it does not contain **ether-type ipv6** or **ip-version ipv6**.



3. Ensure that other match conditions in the term are valid for IPv4 traffic.

To configure a term in a firewall filter configuration specifically for IPv6 traffic:

1. Perform one of these tasks:
  - Define **ether-type ipv6** in a term in the configuration.
  - Define **ip-version ipv6** in a term in the configuration.
  - Define both **ether-type ipv6** and **ip-version ipv4** in a term in the configuration.



**NOTE:** By default, a configuration that does not contain either **ether-type ipv6** or **ip-version ipv6** in a term applies to IPv4 traffic.

2. Ensure that other match conditions in the term are valid for IPv6 traffic.



**NOTE:** If the term contains either of the match conditions **ether-type ipv6** or **ip-version ipv6**, with no other IPv6 match condition specified, all IPv6 traffic is matched.



**NOTE:** To configure a firewall filter for both IPv4 and IPv6 traffic, you must include two separate terms, one for IPv4 traffic and the other for IPv6 traffic.

### Applying a Firewall Filter to a Port on a Switch

You can apply a firewall filter to a port on a switch to filter ingress or egress traffic on the switch. When you configure the firewall filter, you can specify any match condition, action, and action modifiers specified in “[Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches](#)” on page 4704. The action specified in the match condition indicates the action for the matched packets in the ingress or egress traffic.

To apply a firewall filter to a port to filter ingress or egress traffic:



**NOTE:** For applying a firewall filter to a management interface, see “[Applying a Firewall Filter to a Management Interface on a Switch](#)” on page 4810

1. Specify the interface name and provide a meaningful description of the firewall filter and the interface to which the filter is applied:

```
[edit interfaces]
user@switch# set ge-0/0/1 description "filter to limit tcp traffic filter at trunk port for
employee-vlan and voice-vlan applied on the interface"
```



**NOTE:** Providing the description is optional.

2. Specify the unit number and family address type for the interface:

```
[edit interfaces]
user@switch# set ge-0/0/1 unit 0 family ethernet-switching
```

For firewall filters that are applied to ports, the family address type must be **ethernet-switching**.

3. To apply a firewall filter to filter packets that are entering a port:

```
[edit interfaces]
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input ingress-port-filter
```

To apply a firewall filter to filter packets that are exiting a port:

```
[edit interfaces]
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter output egress-port-filter
```



**NOTE:** You can apply no more than one firewall filter per port, per direction.

---

### Applying a Firewall Filter to a Management Interface on a Switch

You can configure and apply a firewall filter to a management interface to control traffic that is entering or exiting the interface on a switch. You can use utilities such as SSH or Telnet to connect to the management interface over the network and then use management protocols such as SNMP to gather statistical data from the switch. Similar to configuring a firewall filter on other types of interfaces, you can configure a firewall filter on a management interface using any match condition, action, and action modifier specified in [“Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches”](#) on page 4704 except for the following action modifiers:

- **loss-priority**
- **forwarding-class**

You can apply a firewall filter to the management Ethernet interface on any EX Series switch. You can also apply a firewall filter to the virtual management Ethernet (VME) interface on the EX4200 switch. For more information on the management Ethernet interface and the VME interface, see [“EX Series Switches Interfaces Overview”](#) on page 2577.

To apply a firewall filter on the management interface to filter ingress or egress traffic:

1. Specify the interface name and provide a meaningful description of the firewall filter and the interface to which the filter is applied:

```
[edit interfaces]
user@switch# set me0 description "filter to limit tcp traffic filter at management interface"
```



**NOTE:** Providing the description is optional.

2. Specify the unit number and family address type for the management interface:

```
[edit interfaces]
user@switch# set me0 unit 0 family inet
```



**NOTE:** For firewall filters that are applied to management interfaces, the family address type can be either `inet` or `inet6`.

3. To apply a firewall filter to filter packets that are entering a management interface:

```
[edit interfaces]
user@switch# set me0 unit 0 family inet filter input ingress-port-filter
```

To apply a firewall filter to filter packets that are exiting a management interface:

```
[edit interfaces]
user@switch# set me0 unit 0 family inet filter output egress-port-filter
```



**NOTE:** You can apply no more than one firewall filter per management interface, per direction.

### Applying a Firewall Filter to a VLAN on a Network

You can apply a firewall filter to a VLAN on a network to filter ingress or egress traffic on the network. To apply a firewall filter to a VLAN, specify the VLAN name and ID, and then apply the firewall filter to the VLAN. When you configure the firewall filter, you can specify any match condition, action, and action modifiers specified in [“Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches” on page 4704](#). The action specified in the match condition indicates the action for the matched packets in the ingress or egress traffic.

To apply a firewall filter to a VLAN:

1. Specify the VLAN name and VLAN ID and provide a meaningful description of the firewall filter and the VLAN to which the filter is applied:

```
[edit vlans]
user@switch# set employee-vlan vlan-id (802.1Q Tagging) 20 vlan-description "filter to rate limit traffic applied on employee-vlan"
```



**NOTE:** Providing the description is optional.

2. Apply firewall filters to filter packets that are entering or exiting the VLAN:

- To apply a firewall filter to filter packets that are entering the VLAN:

```
[edit vlans]
user@switch# set employee-vlan vlan-id 20 filter input ingress-vlan-filter
```

(On EX4300 switches) To apply a firewall filter to filter packets that are entering the VLAN:

```
[edit vlans]
user@switch# set employee-vlan vlan-id 20 forwarding-options input ingress-vlan-filter
```

- To apply a firewall filter to filter packets that are exiting the VLAN:

```
[edit vlans]
user@switch# set employee-vlan vlan-id 20 filter output egress-vlan-filter
```

(On EX4300 switches) To apply a firewall filter to filter packets that are exiting the VLAN:

```
[edit vlans]
user@switch# set employee-vlan vlan-id 20 forwarding-options output egress-vlan-filter
```



**NOTE:** You can apply no more than one firewall filter per VLAN, per direction.

---

### Applying a Firewall Filter to a Layer 3 (Routed) Interface

You can apply a firewall filter to a Layer 3 (routed) interface to filter ingress or egress traffic on the switch. When you configure the firewall filter, you can specify any match condition, action, and action modifiers specified in [“Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches” on page 4704](#). The action specified in the match condition indicates the action for the matched packets in the ingress or egress traffic.

To apply a firewall filter to a Layer 3 interface on a switch:

1. Specify the interface name and provide a meaningful description of the firewall filter and the interface to which the filter is applied:

```
[edit interfaces]
user@switch# set ge-0/1/0 description "filter to count and monitor employee-vlan traffic
applied on layer 3 interface"
```



**NOTE:** Providing the description is optional.

2. Specify the unit number, family address type, and address for the interface:

```
[edit interfaces]
user@switch# set ge-0/1/0 unit 0 family inet address 10.10.10.1/24
```

For firewall filters applied to Layer 3 interfaces, the family address type must be **inet** (for IPv4 traffic) or **inet6** (for IPv6 traffic).

3. You can apply firewall filters to filter packets that are entering or exiting a Layer 3 (routed) interface:

- To apply a firewall filter to filter packets that are entering a Layer 3 interface:

```
[edit interfaces]
user@switch# set ge-0/1/0 unit 0 family inet address 10.10.10.1/24 filter input
ingress-router-filter
```

- To apply a firewall filter to filter packets that are exiting a Layer 3 interface:

```
[edit interfaces]
user@switch# set ge-0/1/0 unit 0 family inet address 10.10.10.1/24 filter output
egress-router-filter
```



**NOTE:** You can apply no more than one firewall filter per Layer 3 interface, per direction.

#### Related Documentation

- [Configuring Firewall Filters \(J-Web Procedure\) on page 4813](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 4795](#)
- [Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch on page 4791](#)
- [Verifying That Firewall Filters Are Operational on page 4873](#)
- [Monitoring Firewall Filter Traffic on page 4875](#)
- [Configuring Policers to Control Traffic Rates \(CLI Procedure\) on page 4818](#)

## Configuring Firewall Filters (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

You configure firewall filters on EX Series switches to control traffic that enters ports on the switch or enters and exits VLANs on the network and Layer 3 (routed) interfaces. To configure a firewall filter, you must configure the filter and then apply it to a port, VLAN, or Layer 3 interface.

To configure firewall filter settings by using the J-Web interface:

1. Select **Configure > Security > Filters**.

The Firewall Filter Configuration page displays a list of all configured ports or VLANs or router filters and the ports or VLANs associated with a particular filter.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—Select this option to create a new filter. Enter information as specified in [Table 532 on page 4814](#).
- **Edit**—Select this option to edit an existing filter. Enter information as specified in [Table 532 on page 4814](#).

- **Delete**—Select this option to delete a filter.
- **Term Up**—Select this option to move a term up in the filter term list.
- **Term Down**—Select this option to move a term down in the filter term list.

**Table 532: Create a New Filter**

Field	Function	Your Action
Filter tab		
Filter type	Specifies the filter type: port or VLAN firewall filter or router firewall filter.	Select the filter type.
Filter name	Specifies the name for the filter.	Enter a name.
Select terms to be part of the filter	Specifies the terms to be associated with the filter. Add new terms or edit existing terms.	Click <b>Add</b> to add new terms. Enter information as specified in <a href="#">Table 533 on page 4814</a> and <a href="#">Table 534 on page 4815</a> .
Association tab		
Port Associations	Specifies the ports with which the filter is associated.  <b>NOTE:</b> For a port or VLAN filter type, only Ingress direction is supported for port association.	<ol style="list-style-type: none"> <li>1. Click <b>Add</b>.</li> <li>2. Select the direction: Ingress or Egress.</li> <li>3. Select the ports. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the available ports from the list.</li> <li>4. Click <b>OK</b>.</li> </ol>
VLAN Associations	Specifies the VLANs with which the filter is associated.  <b>NOTE:</b> Because router firewall filters can be associated with ports only, this section is not displayed for a router firewall filter.	<ol style="list-style-type: none"> <li>1. Click <b>Add</b>.</li> <li>2. Select the direction: Ingress or Egress.</li> <li>3. Select the VLANs.</li> <li>4. Click <b>OK</b>.</li> </ol>

**Table 533: Create a New Term**

Field	Function	Your Action
Term Name	Specifies the name of the term.	Enter a name.
Protocols	Specifies the protocols to be associated with the term.	<ol style="list-style-type: none"> <li>1. Click <b>Add</b>.</li> <li>2. Select the protocols.</li> <li>3. Click <b>OK</b>.</li> </ol>

Table 533: Create a New Term (*continued*)

Field	Function	Your Action
Source	Specifies the source IP address, MAC address, and available ports.  <b>NOTE:</b> MAC address is specified only for port or VLAN filters.	To specify the IP address, click <b>Add &gt; IP</b> and enter the IP address.  To specify the MAC address, click <b>Add &gt; MAC</b> and enter the MAC address.  To specify the ports (interfaces), click <b>Add &gt; Ports</b> and enter the port number.  To delete the IP address, MAC address, or port details, select it and click <b>Remove</b> .
Destination	Specifies the destination IP address, MAC address, and available ports.  <b>NOTE:</b> MAC address is specified only for port or VLAN filters.	To specify the IP address, click <b>Add &gt; IP</b> and enter the IP address.  To specify the MAC address, click <b>Add &gt; MAC</b> and enter the MAC address.  To specify the ports (interfaces), click <b>Add &gt; Ports</b> and enter the port number.  To delete the IP address, MAC address, or port details, select it and click <b>Remove</b> .
Action	Specifies the packet action for the term.	Select one of the following options: <ul style="list-style-type: none"> <li>• Accept</li> <li>• Discard</li> </ul>
More	Specifies advanced configuration options for the filter.	Select the match conditions as specified in <a href="#">Table 534 on page 4815</a> .  Select the packet action for the term as specified in <a href="#">Table 534 on page 4815</a> .

Table 534: Advanced Options for Terms

Table	Function	Your Action
ICMP Type	Specifies the ICMP packet type field. Typically, you specify this match condition in conjunction with the protocol match condition to determine which protocol is being used on the port.	Select the option from the list.
ICMP Code	Specifies more specific information than the ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify <b>icmp-type</b> along with <b>icmp-code</b> . The keywords are grouped by the ICMP type with which they are associated.	Select a value from the list.
DSCP	Specifies the Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP.	Select the DSCP number from the list.

Table 534: Advanced Options for Terms (*continued*)

Table	Function	Your Action
Precedence	Specifies the IP precedence.  <b>NOTE:</b> The IP precedence and the DSCP number cannot be specified together for the same term.	Select the option from the list.
IP Options	Specifies the presence of the options field in the IP header.	Select the option from the list.
Interface	Specifies the interface on which the packet is received.	Select the interface from the list.
Ether type  <b>NOTE:</b> This option is not supported on EX4300 switches.	Specifies the Ethernet type field of a packet.  <b>NOTE:</b> This option is not applicable for a routing filter.	Select a value from the list.
Dot 1q user priority  <b>NOTE:</b> This option is not supported on EX4300 switches.	Specifies the user-priority field of the tagged Ethernet packet. User-priority values can be 0–7.  In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed) :  <ul style="list-style-type: none"> <li>background (1)—Background</li> <li>best-effort (0)—Best effort</li> <li>controlled-load (4)—Controlled load</li> <li>excellent-load (3)—Excellent load</li> <li>network-control (7)—Network control reserved traffic</li> <li>standard (2)—Standard or spare</li> <li>video (5)—Video</li> <li>voice (6)—Voice</li> </ul> <b>NOTE:</b> This option is not applicable for a routing filter.	Select a value from the list.
VLAN  <b>NOTE:</b> This option is not supported on EX4300 switches.	Specifies the VLAN to be associated with the packet.  <b>NOTE:</b> This option is not applicable for a routing filter.	Select the VLAN from the list.
TCP Flags	Specifies one or more TCP flags.  <b>NOTE:</b> TCP flags are supported on ingress ports, VLANs, and router interfaces.	Select the option <b>TCP Initial</b> or enter a combination of TCP flags.
Fragmentation Flags	Specifies the IP fragmentation flags.  <b>NOTE:</b> Fragmentation flags are supported on ingress ports, VLANs, and router interfaces.	Select either the option <b>is-fragment</b> or enter a combination of fragment action flags.



Table 534: Advanced Options for Terms (*continued*)

Table	Function	Your Action
Dot1q tag <b>NOTE:</b> This option is not supported on EX4300 switches.	Specifies the value for the tag field in the Ethernet header. The value can be from 1 through 4095. <b>NOTE:</b> This option is not applicable for a routing filter.	Enter the value.
User Vlan Id <b>NOTE:</b> This option is supported only on EX4300 switches.	Specifies the value of the VLAN ID. The value can be from 0 through 4095 or a range of values.	Enter a value.
User Vlan IP Priority <b>NOTE:</b> This option is supported only on EX4300 switches.	Specifies the priority value. The values can be from 0 through 7.	Enter a value.
Learn Vlan Id <b>NOTE:</b> This option is supported only on EX4300 switches.	Specifies the value of the learnt VLAN ID. The value can be from 0 through 4095 or a range of values.	Enter a value.
Action		
Counter name	Specifies the count of the number of packets that pass this filter, term, or policer.	Enter a value.
Forwarding class	Classifies the packet into one of the following forwarding classes: <ul style="list-style-type: none"> <li>assured-forwarding</li> <li>best-effort</li> <li>expedited-forwarding</li> <li>network-control</li> <li>None</li> </ul>	Select the option from the list.
Loss priority	Specifies the packet loss priority. <b>NOTE:</b> Forwarding class and loss priority must be specified together for the same term.	Enter the value.
Analyzer <b>NOTE:</b> This option is not supported on EX4300 switches.	Specifies whether to perform port mirroring on packets. Port mirroring copies all packets entering one switch port to a network- monitoring connection on another switch port.	Select the analyzer (port mirroring configuration) from the list.

Table 534: Advanced Options for Terms (*continued*)

Table	Function	Your Action
Port Mirror Instance  <b>NOTE:</b> This option is supported only on EX4300 switches.	Specifies whether to perform port mirroring on packets. Port mirroring copies all packets entering one switch port to a network- monitoring connection on another switch port.	Select the port mirroring instance from the list. <b>Default</b> is selected by default.

**Related Documentation**

- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Verifying That Firewall Filters Are Operational on page 4873](#)
- [Firewall Filters for EX Series Switches Overview on page 4696](#)
- [Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704](#)

## Configuring Policers to Control Traffic Rates (CLI Procedure)

You can configure policers to rate limit traffic on EX Series switches. After you configure a policer, you can include it in an ingress firewall filter configuration.

When you configure a firewall filter, you can specify a policer action for any term or terms within the filter. All traffic that matches a term that contains a policer action goes through the policer that the term references. Each policer that you configure includes an implicit counter. To get term-specific packet counts, you must configure a separate policer for each filter term that requires policing.



**NOTE:** On all EX Series switches except EX8200 switches, each policer that you configure includes an implicit counter. To ensure term-specific packet counts, configure a policer for each term in the filter that requires policing. For EX8200 switches, configure a policer and associate it with a global management counter using the *counter* option.

The following policer limits apply on a switch:

- A maximum of 512 policers can be configured for port firewall filters.
- A maximum of 512 policers can be configured for VLAN and Layer 3 firewall filters.

If the number of policers in the firewall filter configuration exceeds these limits, the switch returns the following message when you commit the configuration:

```
Cannot assign policers: Max policer limit reached
```

This topic includes these tasks:

1. [Configuring Policers on page 4819](#)
2. [Specifying Policers in a Firewall Filter Configuration on page 4820](#)
3. [Applying a Firewall Filter That Is Configured with a Policer on page 4820](#)

## Configuring Policers

To configure a policer:

1. Specify the name of the policer:

```
[edit firewall]
user@switch# set policer policer-one
```

The policer name can include letters, numbers, and hyphens (-) and can contain up to 64 characters.

2. Specify the **filter-specific** statement to configure a policer to act as a filter-specific policer; else proceed to step 3:

```
[edit firewall]
user@switch# set policer policer-one filter-specific
```

If you do not specify the **filter-specific** statement, the policer acts as a term-specific policer by default.

3. Configure rate limiting for the policer:

- a. Specify the bandwidth limit in bits per second (bps) to control the traffic rate on an interface:

```
[edit firewall policer policer-one]
user@switch# set if-exceeding bandwidth-limit 300k
```

The range for the bandwidth limit is 1k through 102.3g bps.

- b. Specify the burst-size limit (the maximum allowed burst size in bytes) to control the amount of traffic bursting:

```
[edit firewall policer policer-one]
user@switch# set if-exceeding burst-size-limit 500k
```

To determine the value for the burst-size limit, multiply the bandwidth of the interface on which the filter is applied by the amount of time to allow a burst of traffic at that bandwidth to occur:

$$\text{burst size} = (\text{bandwidth}) * (\text{allowable time for burst traffic})$$

The range for the burst-size limit is 1 through 2,147,450,880 bytes.

4. Specify the policer action **discard** to discard packets that exceed the rate limits:

```
[edit firewall policer]
user@switch# set policer-one then (Policer Action) discard
```

Discard is the only supported policer action.

5. On EX8200 switches, you must assign a global management counter to the policer to obtain policer statistics:

```
[edit firewall policer]
user@switch# set policer-one counter counter-id 0
```

In this sample statement, the global management counter ID is **0**. You can assign any number of policers to the global management counter. The policer statistics displayed for each counter are the collective statistics of all policers assigned to that counter.

### [Specifying Policers in a Firewall Filter Configuration](#)

---

To reference a policer for a single firewall, configure a filter term that includes the policer action:

```
[edit firewall family ethernet-switching]
user@switch# set filter limit-hosts term term-one from source-address 192.0.2.16/28
users@witch# set filter limit-hosts term term-one then policer policer-one
```

### [Applying a Firewall Filter That Is Configured with a Policer](#)

---

A firewall filter that is configured with one or more policer actions, like any other firewall filter, must be applied to a port, VLAN, or Layer 3 interface. For information about applying firewall filters, see the sections on applying firewall filters in [“Configuring Firewall Filters \(CLI Procedure\)”](#) on page 4804.

#### **Related Documentation**

- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches](#) on page 4773
- [Configuring Firewall Filters \(CLI Procedure\)](#) on page 4804
- [Configuring Firewall Filters \(J-Web Procedure\)](#) on page 4813
- [Verifying That Policers Are Operational](#) on page 4874
- [Understanding the Use of Policers in Firewall Filters](#) on page 4767

## Assigning Multifield Classifiers in Firewall Filters to Specify Packet-Forwarding Behavior (CLI Procedure)

You can configure firewall filters with multifield classifiers to classify packets transiting a port, VLAN, or Layer 3 interface on an EX Series switch.

You specify multifield classifiers in a firewall filter configuration to set the forwarding class and packet loss priority (PLP) for incoming or outgoing packets. By default, the data traffic that is not classified is assigned to the **best-effort** class associated with queue 0.

You can specify any of the following default forwarding classes:

Forwarding class	Queue
best-effort	0
assured-forwarding	1
expedited-forwarding	5
network-control	7

To assign multifield classifiers in firewall filters:

1. Configure the family name and filter name for the filter at the **[edit firewall]** hierarchy level, for example:

```
[edit firewall]
user@switch# set family ethernet-switching
user@switch# set family ethernet-switching filter ingress-filter
```

2. Configure the terms of the filter, including the **forwarding-class** and **loss-priority** action modifiers as appropriate. When you specify a forwarding class you must also specify the packet loss priority. For example, each of the following terms examines different packet header fields and assigns an appropriate classifier and the packet loss priority:

- The term **voice-traffic** matches packets on the **voice-vlan** and assigns the forwarding class **expedited-forwarding** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term voice-traffic from vlan-id voice-vlan
user@switch# set term voice-traffic then forwarding-class expedited-forwarding
user@switch# set term voice-traffic then loss-priority low
```

- The term **data-traffic** matches packets on **employee-vlan** and assigns the forwarding class **assured-forwarding** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term data-traffic from vlan-id employee-vlan
user@switch# set term data-traffic then forwarding-class assured-forwarding
user@switch# set term data-traffic then loss-priority low
```

- Because loss of network-generated packets can jeopardize proper network operation, delay is preferable to discard of packets. The following term,

**network-traffic**, assigns the forwarding class **network-control** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term network-traffic from precedence net-control
user@switch# set term network-traffic then forwarding-class network
user@switch# set term network-traffic then loss-priority low
```

- The last term **accept-traffic** matches any packets that did not match on any of the preceding terms and assigns the forwarding class **best-effort** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term accept-traffic from precedence net-control
user@switch# set term accept-traffic then forwarding-class best-effort
user@switch# set term accept-traffic then loss-priority low
```

3. Apply the filter **ingress-filter** to a port, VLAN or Layer 3 interface. For information about applying the filter, see [“Configuring Firewall Filters \(CLI Procedure\)” on page 4804](#).

#### Related Documentation

- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Verifying That Firewall Filters Are Operational on page 4873](#)
- [Monitoring Firewall Filter Traffic on page 4875](#)
- [Defining CoS Classifiers \(CLI Procedure\) on page 2103](#)
- [Defining CoS Classifiers \(J-Web Procedure\) on page 2105](#)
- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
- [Configuring Firewall Filters \(J-Web Procedure\) on page 4813](#)

## Configuring Routing Policies (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

All routing protocols use the Junos OS routing table to store the routes that they learn and to determine which routes are advertised in the protocol packets. Routing policy allows you to control which routes the routing protocols store in and retrieve from the routing table on the routing device.

To configure routing policies for an EX Series switch using the J-Web interface:

1. Select **Configure > Routing > Policies**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Global Options**—Configures global options for policies. Enter information into the configuration page as described in [Table 535 on page 4823](#).
- **Add**—Configures a new policy. Select **New** and specify a policy name. To add terms, enter information into the configuration page as described in [Table 536 on page 4824](#). Select **Clone** to create a copy of an existing policy.
- **Edit**—Edits an existing policy. To modify an existing term, enter information into the configuration page as described in [Table 536 on page 4824](#).
- **Term Up**—Moves a term up in the list.
- **Term Down**—Moves a term down in the list.
- **Delete**—Deletes the selected policy.
- **Test Policy**—Tests the policy. Use this option to check whether the policy produces the results that you expect.

**Table 535: Policies Global Configuration Parameters**

Field	Function	Your Action
Prefix List	Specifies a list of IPv4 address prefixes for use in a routing policy statement.	<p>To add a prefix list:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b>.</li> <li>2. Enter a name for the prefix list.</li> <li>3. To add an IP address, click <b>Add</b>.</li> <li>4. Enter the IP address and the subnet mask and click <b>OK</b>.</li> <li>5. Click <b>OK</b>.</li> </ol> <p>To edit a prefix list, click <b>Edit</b>. Edit the settings and click <b>OK</b>.</p> <p>To delete a prefix list, select it and click <b>Delete</b>.</p>
BGP Community	Specifies a BGP community.	<p>To add a BGP community:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b>.</li> <li>2. Enter a name for the community.</li> <li>3. To add a community, click <b>Add</b>.</li> <li>4. Enter the community ID and click <b>OK</b>.</li> <li>5. Click <b>OK</b>.</li> </ol> <p>To edit a BGP community, click <b>Edit</b>. Edit the settings and click <b>OK</b>.</p> <p>To delete a BGP community, select it and click <b>Delete</b>.</p>

Table 535: Policies Global Configuration Parameters (*continued*)

Field	Function	Your Action
AS Path	Specifies an AS path. This is applicable to BGP only.	<p>To add an AS path:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b>.</li> <li>2. Enter the AS path name.</li> <li>3. Enter the regular expression and click <b>OK</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> <p>To edit an AS path, click <b>Edit</b>. Edit the settings and click <b>OK</b>.</p> <p>To delete an AS path, select it and click <b>Delete</b>.</p>

Table 536: Terms Configuration Parameters

Field	Function	Your Action
Term Name	Specifies a term name.	Type or select and edit the name.
<b>Source tab</b>		
Family	Specifies an address family protocol.	Select a value from the list.
Routing Instance	Specifies a routing instance.	Select a value from the list.
RIB	Specifies the name of a routing table.	Select a value from the list
Preference	Specifies the individual preference value for the route.	Type or select and edit the value.
Metric	Specifies a metric value. You can specify up to four metric values.	Type or select and edit the value.
Interface	Specifies a name or IP address of one or more routing device interfaces. Do not use this qualifier with protocols that are not interface-specific, such as internal BGP (IBGP).	<p>To add an interface, select <b>Add &gt; Interface</b>. Select the interface from the list. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.</p> <p>To add an address, select <b>Add &gt; Address</b>. Select the address from the list.</p> <p>To remove an interface, select it and click <b>Remove</b>.</p>
Prefix List	Specifies a named list of IP addresses. You can specify an exact match with incoming routes.	<p>Click <b>Add</b>. Select the prefix list from the list and click <b>OK</b>.</p> <p>To remove a prefix list, select it and click <b>Remove</b>.</p>
Protocol	Specifies the name of the protocol from which the route was learned or to which the route is being advertised.	<p>Click <b>Add</b> and select the protocol from the list.</p> <p>To remove a protocol, select it and click <b>Remove</b>.</p>



Table 536: Terms Configuration Parameters (*continued*)

Field	Function	Your Action
Policy	Specifies the name of a policy to evaluate as a subroutine.	Click <b>Add</b> . Select the policy from the list.  To remove a policy, select it and click <b>Remove</b> .
More	Specifies advanced configuration options for policies.	Click <b>More</b> for advanced configuration.
OSPF Area ID	Specifies the area identifier.	Type the IP address.
BGP Origin	Specifies the origin of the AS path information.	Select a value from the list.
Local Preference	Specifies the BGP local preference.	Type a value.
Route	Specifies the type of route.	Select <b>External</b> .  Select the OSPF type from the list.
AS Path	Specifies the name of an AS path regular expression.	Click <b>Add</b> . Select the AS path from the list.
Community	Specifies the name of one or more communities.	Click <b>Add</b> . Select the community from the list.
<b>Destination tab</b>		
Family	Specifies an address family protocol.	Select a value from the list.
Routing Instance	Specifies a routing instance.	Select a value from the list.
RIB	Specifies the name of a routing table.	Select a value from the list.
Preference	Specifies the individual preference value for the route.	Type a value.
Metric	Specifies a metric value.	Type a value.
Interface	Specifies a name or IP address of one or more routing device interfaces. Do not use this qualifier with protocols that are not interface-specific, such as internal BGP (IBGP).	To add an interface, select <b>Add &gt; Interface</b> . Select the interface from the list. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.  To add an address, select <b>Add &gt; Address</b> . Select the address from the list.  To delete an interface, select it and click <b>Remove</b> .
Protocol	Specifies the name of the protocol from which the route was learned or to which the route is being advertised.	Click <b>Add</b> and select the protocol from the list.  To delete a protocol, select it and click <b>Remove</b> .

Table 536: Terms Configuration Parameters (*continued*)

Field	Function	Your Action
<b>Action tab</b>		
Action	Specifies the action to take if the conditions match.	Select a value from the list.
Default Action	Specifies that any action that is intrinsic to the protocol is overridden. This action is also nonterminating, so that various policy terms can be evaluated before the policy is terminated.	Select a value from the list.
Next	Specifies the default control action if a match occurs, and there are no further terms in the current routing policy.	Select a value from the list.
Priority	Specifies a priority for prefixes included in an OSPF import policy. Prefixes learned through OSPF are installed in the routing table based on the priority assigned to the prefixes.	Select a value from the list.
BGP Origin	Specifies the BGP origin attribute.	Select a value from the list.
AS Path Prepend	Affixes an AS number at the beginning of the AS path. The AS numbers are added after the local AS number has been added to the path. This action adds an AS number to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the affixed AS number is placed within a confederation sequence. Otherwise, the affixed AS number is placed with a nonconfederation sequence.	Enter a value.
AS Path Expand	Extracts the last AS number in the existing AS path and affixes that AS number to the beginning of the AS path $n$ times, where $n$ is a number from 1 through 32. The AS number is added before the local AS number has been added to the path. This action adds AS numbers to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the affixed AS numbers are placed within a confederation sequence. Otherwise, the affixed AS numbers are placed within a nonconfederation sequence. This option is typically used in non-IBGP export policies.	Select the type and type a value.
Load Balance Per Packet	Specifies that all next-hop addresses in the forwarding table must be installed and have the forwarding table perform per-packet load balancing. This policy action allows you to optimize VPLS traffic flows across multiple paths.	Select the check box to enable the option.
Tag	Specifies the tag value. The tag action sets the 32-bit tag field in OSPF external link-state advertisement (LSA) packets.	Select the action and type a value.
Metric	Changes the metric (MED) value by the specified negative or positive offset. This action is useful only in an external BGP (EBGP) export policy.	Select the action and type a value.
Route	Specifies whether the route is external.	Select the <b>External</b> check box to enable the option, and select the OSPF type.

Table 536: Terms Configuration Parameters (*continued*)

Field	Function	Your Action
Preference	Specifies the preference value.	Select the preference action and type a value.
Local Preference	Specifies the BGP local preference attribute.	Select the action and type a value.
Class of Service	Specifies and applies the class-of-service parameters to routes installed into the routing table. <ul style="list-style-type: none"> <li>Source class The value entered here maintains the packet counts for a route passing through your network, based on the source address.</li> <li>Destination class The value entered here maintains packet counts for a route passing through your network, based on the destination address in the packet.</li> <li>Forwarding class</li> </ul>	Type the source class. Type the destination class. Type the forwarding class.

**Related Documentation**

- [Configuring BGP Sessions \(J-Web Procedure\) on page 2943](#)
- [Configuring an OSPF Network \(J-Web Procedure\) on page 3207](#)
- [Configuring a RIP Network \(J-Web Procedure\) on page 3347](#)
- [Configuring Static Routing \(J-Web Procedure\) on page 3422](#)
- [Layer 3 Protocols Supported on EX Series Switches on page 2939](#)

## Configuring Tricolor Marking Policers

You can rate-limit traffic on EX Series switches by configuring a policer and specifying it as an action modifier for a term in a firewall filter. By default, if you specify the same policer in multiple terms, Junos OS creates a separate policer instance for each term and applies rate limiting separately for each instance. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, each policer instance enforces a 1-Gbps limit. In this case, the total bandwidth allowed by the filter is 3 Gbps.

You can also configure a policer to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps.

This topic describes how to configure single-rate and two-rate tricolor marking (TCM) policers, also known as single-rate and two-rate three-color policers. If you want to

configure a single-rate two-color policer (also known just as a "policer"), see ["Configuring Policers to Control Traffic Rates \(CLI Procedure\)" on page 4818](#).

This topic includes:

- [Configuring a Tricolor Marking Policer on page 4828](#)
- [Applying Tricolor Marking Policers to Firewall Filters on page 4828](#)

---

### Configuring a Tricolor Marking Policer

A tricolor marking policer polices traffic on the basis of metering rates, including the configured information rate (CIR), the peak information rate (PIR), their associated burst sizes, and any policing actions configured for the traffic. With tri-color marking, you can configure traffic policing according to two separate modes—color-blind and color-aware. In color-blind mode, the current packet loss priority (PLP) value is ignored. In color-aware mode, the current PLP values are considered by the policer, and the policer can increase those values but cannot decrease them.

To configure a tricolor marking (TCM) policer:

1. Specify the name of the policer and (optionally) whether to automatically discard packets with high loss priority (PLP):

```
[edit firewall]
user@switch# set three-color-policer policer-name
user@switch# set three-color-policer policer-name action loss-priority high then discard
```

2. Specify the policer as either single-rate or two-rate and as color-aware or color-blind:

```
[edit firewall three-color-policer policer-name]
user@switch# set rate mode
```

For example:

```
[edit firewall three-color-policer srTCm1a]
user@switch# set single-rate color-aware
[edit firewall three-color-policer trTCM2-cb]
user@switch# set two-rate color-blind
```

3. For a single-rate TCM policer, configure the CIR, committed burst size (CBS), and excess burst size (EBS):

```
[edit firewall three-color-policer policer-name single-rate]
user@switch# set committed-information-rate bps
user@switch# set committed-burst-size bytes
user@switch# set excess-burst-size bytes
```

4. For a two-rate TCM policer, configure the CIR, CBS, PIR, and peak burst size (PBS):

```
[edit firewall three-color-policer policer-name single-rate]
user@switch# set committed-information-rate bps
user@switch# set committed-burst-size bytes
user@switch# set peak-information-rate bps
user@switch# set peak-burst-size bytes
```

---

### Applying Tricolor Marking Policers to Firewall Filters

To rate-limit traffic by applying a tricolor marking (TCM) policer to a firewall filter:

```
[edit firewall family family filter filter-name term term-name then]
user@switch# set three-color-policer rate stTCM1-ca
```

For example:

```
[edit firewall family inet filter test1 term term1 then]
user@switch# set three-color-policer single-rate policer1
```

You must include either the **single-rate** statement or the **two-rate** statement in the reference to the policer in the firewall filter configuration, and this statement must match the configured TCM policer. Otherwise, an error message appears in the configuration listing.

For example, if you configure **srTCM1-ca** as a single-rate TCM policer and try to apply it as a two-rate policer, the following message appears:

```
[edit firewall]
user@switch# show three-color-policer srTCM1-ca
single-rate {
  color-aware;
  ...
}
user@switch# show filter TESTER
term A {
  then {
    three-color-policer {
      ##
      ## Warning: Referenced two-rate policer does not exist
      ##
      two-rate srTCM;
    }
  }
}
```

#### Related Documentation

- [Understanding Tricolor Marking Architecture on page 4770](#)
- [Understanding the Use of Policers in Firewall Filters on page 4767](#)

## Configuration Statements

- [\[edit firewall\] Configuration Statement Hierarchy on EX Series Switches on page 4830](#)
- [Firewall Filter Configuration Statements Supported by Junos OS for EX Series Switches on page 4832](#)
- [action \(TCM Policers\) on page 4835](#)
- [apply-path on page 4835](#)
- [as-path \(Policy Options\) on page 4836](#)
- [as-path-group on page 4837](#)
- [bandwidth-limit on page 4838](#)
- [burst-size-limit on page 4839](#)
- [color-aware on page 4840](#)
- [color-blind on page 4840](#)
- [committed-burst-size on page 4841](#)

- [committed-information-rate](#) on page 4842
- [community \(Policy Options\)](#) on page 4843
- [condition](#) on page 4846
- [damping \(Policy Options\)](#) on page 4847
- [dynamic-db](#) on page 4848
- [excess-burst-size](#) on page 4849
- [family \(Firewall Filter\)](#) on page 4850
- [filter \(Firewall Filters\)](#) on page 4851
- [filter \(VLANs\)](#) on page 4852
- [filter-specific](#) on page 4853
- [firewall](#) on page 4854
- [from](#) on page 4855
- [if-exceeding](#) on page 4856
- [interface-specific](#) on page 4857
- [loss-priority high then discard \(Three-Color Policer\)](#) on page 4857
- [peak-burst-size](#) on page 4858
- [policer](#) on page 4859
- [policy-options](#) on page 4860
- [policy-statement](#) on page 4861
- [prefix-list](#) on page 4865
- [routing-instance](#) on page 4866
- [single-rate](#) on page 4867
- [term](#) on page 4868
- [then \(Firewall Filters\)](#) on page 4869
- [then \(Policer Action\)](#) on page 4870
- [three-color-policer \(Configuring\)](#) on page 4871
- [two-rate](#) on page 4872

## **[edit firewall]** Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit firewall]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit firewall\] Hierarchy Level on page 4831](#)
- [Unsupported Statements in the \[edit firewall\] Hierarchy Level on page 4832](#)

### Supported Statements in the [edit firewall] Hierarchy Level

The following hierarchy shows the **[edit firewall]** configuration statements supported on EX Series switches:

```

firewall {
  family family-name {
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
  policer policer-name {
    filter-specific;
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      policer-action;
    }
  }
  three-color-policer policer-name {
    action {
      loss-priority high then discard;
    }
    filter-specific;
    single-rate {
      (color-aware | color-blind);
      committed-burst-size bytes;
      committed-information-rate bps;
      excess-burst-size bytes;
    }
    two-rate {
      (color-aware | color-blind);
      committed-burst-size bytes;
      committed-information-rate bps;
      peak-information-rate bps;
      peak-burst-size bytes;
    }
  }
}

```

### Unsupported Statements in the [edit firewall] Hierarchy Level

All statements in the [edit firewall] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
- [Configuring Policers to Control Traffic Rates \(CLI Procedure\) on page 4818](#)
- [Firewall Filter Configuration Statements Supported by Junos OS for EX Series Switches on page 4832](#)
- [Firewall Filters for EX Series Switches Overview on page 4696](#)

## Firewall Filter Configuration Statements Supported by Junos OS for EX Series Switches

You configure firewall filters to filter packets based on their components and to perform an action on packets that match the filter.

[Table 537 on page 4832](#) lists the options that are supported for the firewall statement in Junos OS for EX Series switches.

**Table 537: Supported Options for Firewall Filter Statements**

Statement and Option	Description
<code>family <i>family-name</i> {</code> <code>}</code>	<p>The <b><i>family-name</i></b> option specifies the version or type of addressing protocol:</p> <ul style="list-style-type: none"> <li>• <b>any</b>—Filter packets based on protocol-independent match conditions.</li> <li>• <b>ethernet-switching</b>—Filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets</li> <li>• <b>inet</b>—Filter IPv4 packets</li> <li>• <b>inet6</b>—Filter IPv6 packets</li> </ul>
<code>filter <i>filter-name</i> {</code> <code>}</code>	<p>The <b><i>filter-name</i></b> option identifies the filter. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the name in quotation marks (" ").</p>
<code>interface-specific</code>	<p>The <b>interface-specific</b> statement configures unique names for individual firewall counters specific to each interface.</p>
<code>term <i>term-name</i> {</code> <code>}</code>	<p>The <b><i>term-name</i></b> option identifies the term. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). Each term name must be unique within a filter.</p>



Table 537: Supported Options for Firewall Filter Statements (*continued*)

Statement and Option	Description
<pre> from {   match-conditions; } </pre>	<p>The <b>from</b> statement is optional. If you omit it, all packets are considered to match.</p>
<pre> then {   action;   action-modifiers; } </pre>	<p>For information about the <b>action</b> and <b>action-modifiers</b> options, see <a href="#">“Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches”</a> on page 4704.</p>
<pre> policer policer-name { } </pre>	<p>The <b>policer-name</b> option identifies the policer. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the name in quotation marks (" ").</p>
filter-specific	<p>The <b>filter-specific</b> statement configures policers and counters for a specific filter name.</p>
<pre> if-exceeding {   bandwidth-limit bps   burst-size-limit bytes } </pre>	<p>The <b>bandwidth-limit bps</b> option specifies the traffic rate in bits per second (bps).</p> <p>You can specify <b>bps</b> as a decimal value or as a decimal number followed by one of the following abbreviations:</p> <ul style="list-style-type: none"> <li>• k (thousand)</li> <li>• m (million)</li> <li>• g (billion, which is also called a thousand million)</li> </ul> <p><b>Range:</b> 1000 (1k) through 102,300,000,000 (102.3g) bps</p> <p>The <b>burst-size-limit bytes</b> option specifies the maximum allowed burst size to control the amount of traffic bursting. To determine the value for the burst-size limit, you can multiply the bandwidth of the interface on which the filter is applied by the amount of time (in seconds) to allow a burst of traffic at that bandwidth to occur:</p> <p>burst size = bandwidth * allowable time for burst traffic</p> <p>You can specify a decimal value or a decimal number followed by k (thousand) or m (million).</p> <p><b>Range:</b> 1 through 2,147,450,880 bytes</p>
<pre> then {   policer-action } </pre>	<p>Use the <b>policer-action</b> option to specify <b>discard</b> to discard traffic that exceeds the rate limits.</p>

Junos OS for EX Series switches does not support some of the firewall filter statements that are supported by other Junos OS packages. [Table 538 on page 4834](#) shows the firewall filter statements that are not supported by Junos OS for EX Series switches.

Table 538: Firewall Filter Statements That Are Not Supported by Junos OS for EX Series Switches

Statements Not Supported	Statement Hierarchy Level
<ul style="list-style-type: none"> <li>• <code>interface-set <i>interface-set-name</i> {</code>     <code>}</code></li> <li>• <code>load-balance-group <i>group-name</i> {</code>     <code>}</code></li> <li>• <code>three-color-policer <i>name</i> {</code>     <code>}</code></li> <li>• <code>logical-interface-policer;</code></li> <li>• <code>single-rate {</code>     <code>}</code></li> <li>• <code>two-rate {</code>     <code>}</code></li> </ul>	[edit firewall]
<ul style="list-style-type: none"> <li>• <code>prefix-action <i>name</i> {</code>     <code>}</code></li> <li>• <code>prefix-policer {</code>     <code>}</code></li> <li>• <code>service-filter <i>filter-name</i> {</code>     <code>}</code></li> <li>• <code>simple-filter <i>simple-filter-name</i> {</code>     <code>}</code></li> </ul>	[edit firewall family <i>family-name</i> ]
<ul style="list-style-type: none"> <li>• <code>accounting-profile <i>name</i>;</code></li> </ul>	[edit firewall family <i>family-name</i> filter <i>filter-name</i> ]
<ul style="list-style-type: none"> <li>• <code>logical-bandwidth-policer;</code></li> <li>• <code>logical-interface-policer;</code></li> </ul>	[edit firewall policer <i>policer-name</i> ]
<code>bandwidth-percent <i>number</i>;</code>	[edit firewall policer <i>policer-name</i> if-exceeding]

**Related  
Documentation**

- [Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
- [Configuring Policers to Control Traffic Rates \(CLI Procedure\) on page 4818](#)
- [Firewall Filters for EX Series Switches Overview on page 4696](#)

## action (TCM Policers)

---

<b>Syntax</b>	action { loss-priority high then discard; }
<b>Hierarchy Level</b>	[edit firewall three-color-policer <i>name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2 for EX Series switches.
<b>Description</b>	Discard traffic on a logical interface using tricolor marking policing.  The remaining statement is explained separately.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Tricolor Marking Policers on page 4827</a></li> </ul>

## apply-path

---

<b>Syntax</b>	apply-path <i>path</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> policy-options <i>prefix-list name</i> ], [edit policy-options <i>prefix-list name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Expand a prefix list to include all prefixes pointed to by a defined path.
<b>Options</b>	<i>path</i> —String of elements composed of identifiers or configuration keywords that points to a set of prefixes. You can include wildcards (enclosed in angle brackets) to match more than one identifier. You cannot add a path element, including wildcards, after a leaf statement. Path elements, including wildcards, can only be used after a container statement.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Prefix Lists</a></li> <li>• <a href="#">Example: Configuring Routing Policy Prefix Lists</a></li> <li>• <a href="#">Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List</a></li> </ul>

## as-path (Policy Options)

---

<b>Syntax</b>	<code>as-path name regular-expression;</code>
<b>Hierarchy Level</b>	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Define an autonomous system (AS) path regular expression for use in a routing policy match condition.
<b>Options</b>	<b>name</b> —Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 65,536 characters long. To include spaces in the name, enclose it in quotation marks (" ").  <b>regular-expression</b> —One or more regular expressions used to match the AS path.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Understanding AS Path Regular Expressions for Use as Routing Policy Match Conditions</i></li><li>• <i>Example: Using AS Path Regular Expressions</i></li><li>• <a href="#">dynamic-db on page 4848</a></li></ul>

## as-path-group

---

<b>Syntax</b>	<code>as-path-group <i>group-name</i> {     as-path <i>name</i> <i>regular-expression</i>; }</code>
<b>Hierarchy Level</b>	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for dynamic database configuration introduced in Junos OS Release 9.5. Support for dynamic database configuration introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Define a group containing multiple AS path regular expressions for use in a routing policy match condition.
<b>Options</b>	<p><b><i>group-name</i></b>—Name that identifies the AS path group. One or more AS path regular expressions must be listed below the <b>as-path-group</b> hierarchy.</p> <p><b><i>name</i></b>—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p><b><i>regular-expression</i></b>—One or more regular expressions used to match the AS path.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding AS Path Regular Expressions for Use as Routing Policy Match Conditions</i></li> <li>• <a href="#">dynamic-db on page 4848</a></li> </ul>

## bandwidth-limit

---

<b>Syntax</b>	<code>bandwidth-limit <i>bps</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">firewall policer <i>policer-name</i> if-exceeding</a> ] [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i> if-exceeding]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Logical systems support introduced in Junos OS Release 9.3.
<b>Description</b>	Specify the traffic rate in bits per second.
<b>Options</b>	<p><b><i>bps</i></b> —Traffic rate to be specified in bits per second. Specify <b><i>bps</i></b> as a decimal value or as a decimal number followed by one of the following abbreviations:</p> <ul style="list-style-type: none"><li>• k (thousand)</li><li>• m (million)</li><li>• g (billion, which is also called a thousand million)</li></ul> <p><b>Range:</b></p> <ul style="list-style-type: none"><li>• 1000 (1k) through 102,300,000,000 (102.3g) bps (EX Series switches)</li><li>• 8000 (8k) through 40,000,000,000 (40g) bps (routers)</li></ul>
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li><li>• <a href="#">Configuring Policers to Control Traffic Rates (CLI Procedure) on page 4818</a></li><li>• <a href="#">Understanding the Use of Policers in Firewall Filters on page 4767</a></li><li>• <i>Basic Single-Rate Two-Color Policers</i></li></ul>

## burst-size-limit

---

<b>Syntax</b>	<code>burst-size-limit bytes;</code>
<b>Hierarchy Level</b>	[edit <code>firewall policer policer-name if-exceeding</code> ] [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i> if-exceeding]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Logical systems support introduced in Junos OS Release 9.3.
<b>Description</b>	Specify the maximum allowed burst size to control the amount of traffic bursting.
<b>Options</b>	<b>bytes</b> —Decimal value or a decimal number followed by k (thousand) or m (million). <b>Range:</b> <ul style="list-style-type: none"> <li>1 through 2,147,450,880 bytes (EX Series switches)</li> <li>1500 through 1,00,000,000,000 bytes (routers)</li> </ul>
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li> <li>• <a href="#">Configuring Policers to Control Traffic Rates (CLI Procedure) on page 4818</a></li> <li>• <a href="#">Understanding the Use of Policers in Firewall Filters on page 4767</a></li> <li>• <i>Basic Single-Rate Two-Color Policers</i></li> </ul>

## color-aware

---

<b>Syntax</b>	color-aware;
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate] [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2 for EX Series switches.
<b>Description</b>	Configure the way preclassified packets are metered. In color-aware mode, the local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority.
<b>Default</b>	If you omit the <b>color-aware</b> statement, the default behavior is color-aware mode.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Tricolor Marking Policers on page 4827</a></li></ul>


## color-blind

---

<b>Syntax</b>	color-blind;
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate] [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2 for EX Series switches.
<b>Description</b>	Configure the way preclassified packets are metered. In color-blind mode, the local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority.
<b>Default</b>	If you omit the <b>color-blind</b> statement, the default behavior is color-aware mode.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Tricolor Marking Policers on page 4827</a></li></ul>




## committed-burst-size

<b>Syntax</b>	<code>committed-burst-size <i>bytes</i>;</code>
<b>Hierarchy Level</b>	[edit <code>firewall three-color-policer <i>policer-name</i> single-rate</code> ], [edit <code>firewall three-color-policer <i>policer-name</i> two-rate</code> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2 for EX Series switches.
<b>Description</b>	Configure the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green).
<div>  <p><b>NOTE:</b> When you include the <code>committed-burst-size</code> statement in the configuration, you must also include the <code>committed-information-rate</code> statement at the same hierarchy level.</p> </div>	
<b>Options</b>	<p><b>bytes</b>—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1500 through 100,000,000,000 bytes</p>
<b>Required Privilege Level</b>	<p><code>firewall</code>—To view this statement in the configuration.</p> <p><code>firewall-control</code>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Tricolor Marking Policers on page 4827</a></li> </ul>

## committed-information-rate

---

<b>Syntax</b>	<code>committed-information-rate <i>bps</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate], [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2 for EX Series switches.
<b>Description</b>	Configure the guaranteed bandwidth under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green).
<div> <b>NOTE:</b> When you include the <code>committed-information-rate</code> statement in the configuration, you must also include the <code>committed-burst-size</code> statement at the same hierarchy level.</div>	
<b>Options</b>	<b><i>bps</i></b> —Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000). <b>Range:</b> 32,000 through 40,000,000,000 bps
<b>Required Privilege Level</b>	<b>firewall</b> —To view this statement in the configuration. <b>firewall-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Tricolor Marking Policers on page 4827</a></li></ul>

## community (Policy Options)

<b>Syntax</b>	<pre>community <i>name</i> {     invert-match;     members [ <i>community-ids</i> ]; }</pre>
<b>Hierarchy Level</b>	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Define a community or extended community for use in a routing policy match condition.
<b>Options</b>	<p><b><i>name</i></b>—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters. To include spaces in the name, enclose it in quotation marks (" ").</p> <p><b><i>invert-match</i></b>—Invert the results of the community expression matching. The <b>community</b> match condition defines a regular expression and if it matches the community attribute of the received prefix, Junos OS returns a TRUE result. If not, Junos OS returns a FALSE result. The <b><i>invert-match</i></b> statement makes Junos OS behave to the contrary. If there is a match, Junos OS returns a FALSE result. If there is no match, Junos OS returns a TRUE result.</p> <p><b><i>members community-ids</i></b>—One or more community members. If you specify more than one member, you must enclose all members in brackets.</p> <p>The format for <b><i>community-ids</i></b> is:</p> <p style="padding-left: 40px;"><b><i>as-number:community-value</i></b></p> <p><b><i>as-number</i></b> is the AS number and can be a value in the range from 0 through 65,535.  <b><i>community-value</i></b> is the community identifier and can be a number in the range from 0 through 65,535.</p> <p>You also can specify <b><i>community-ids</i></b> for communities as one of the following well-known community names, which are defined in RFC 1997, <i>BGP Communities Attribute</i>:</p> <ul style="list-style-type: none"> <li>• <b><i>no-export</i></b>—Routes containing this community name are not advertised outside a BGP confederation boundary.</li> <li>• <b><i>no-advertise</i></b>—Routes containing this community name are not advertised to other BGP peers.</li> <li>• <b><i>no-export-subconfed</i></b>—Routes containing this community name are not advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation.</li> </ul>

You can explicitly exclude BGP community information with a static route using the **none** option. Include **none** when configuring an individual route in the **route** portion of the **static** statement to override a **community** option specified in the **defaults** portion of the statement.

The format for extended **community-ids** is the following:

*type:administrator:assigned-number*

**type** is the type of extended community and can be either a **bandwidth**, **target**, **origin**, **domain-id**, **src-as**, or **rt-import** community or a 16-bit number that identifies a specific BGP extended community. The **target** community identifies the destination to which the route is going. The **origin** community identifies where the route originated. The **domain-id** community identifies the OSPF domain from which the route originated. The **src-as** community identifies the autonomous system from which the route originated. The **rt-import** community identifies the route to install in the routing table.



**NOTE:** For **src-as**, you can specify only an AS number and not an IP address. For **rt-import**, you can specify only an IP address and not an AS number.

**administrator** is the administrator. It is either an AS number or an IPv4 address prefix, depending on the type of extended community.

**assigned-number** identifies the local provider.

The format for linking a bandwidth with an AS number is:

*bandwidth:as-number:bandwidth*

**as-number** specifies the AS number and **bandwidth** specifies the bandwidth in bytes per second.



**NOTE:** In Junos OS Release 9.1 and later, you can specify 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*, as well as the 2-byte AS numbers that are supported in earlier releases of the Junos OS. In plain-number format, you can configure a value in the range from 1 through 4,294,967,295. To configure a **target** or **origin** extended community that includes a 4-byte AS number in the plain-number format, append the letter “L” to the end of number. For example, a target community with the 4-byte AS number 334,324 and an assigned number of 132 is represented as **target:334324L:132**.

In Junos OS Release 9.2 and later, you can also use AS-dot notation when defining a 4-byte AS number for the **target** and **origin** extended communities. Specify two integers joined by a period: *16-bit high-order value in decimal.16-bit low-order value in decimal*. For example, the 4-byte AS number represented in plain-number format as 65546 is represented in AS-dot notation as 1.10.



<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Understanding BGP Communities and Extended Communities as Routing Policy Match Conditions</i></li><li>• <i>Understanding How to Define BGP Communities and Extended Communities</i></li><li>• <a href="#">dynamic-db on page 4848</a></li></ul>

## condition

---

<b>Syntax</b>	<pre>condition <i>condition-name</i> {     <b>dynamic-db</b>;     if-route-exists{         <i>address</i>;         address-family {             inet {                 <i>address</i>;                 table <i>table-name</i>;             }             ccc {                 <i>interface-name</i>;                 standby;                 peer-unit <i>unit-number</i>;                 table <i>table-name</i>;             }         }         table <i>table-name</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches. Support for the address families introduced in Junos OS Release 13.2.
<b>Description</b>	Define a policy condition based on the existence of routes in specific tables for use in BGP export policies.
<b>Options</b>	<b><i>condition-name</i></b> —Name of the condition.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Understanding Conditional Installation of Prefixes in a Routing Table</i></li><li>• <i>Example: Configuring Pseudowire Redundancy for Mobile Backhaul Scenarios</i></li><li>• <a href="#">dynamic-db on page 4848</a></li></ul>

## damping (Policy Options)

<b>Syntax</b>	<pre>damping <i>name</i> {   disable;   half-life <i>minutes</i>;   max-suppress <i>minutes</i>;   reuse <i>number</i>;   suppress <i>number</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Define route flap damping properties to set on BGP routes.
<b>Options</b>	<p><b>disable</b>—Disable damping on a per-prefix basis. Any damping state that is present in the routing table for a prefix is deleted if damping is disabled.</p> <p><b>half-life <i>minutes</i></b>—Decay half-life. <i>minutes</i> is the interval after which the accumulated figure-of-merit value is reduced by half if the route remains stable.</p> <p><b>Range:</b> 1 through 45</p> <p><b>Default:</b> 15 minutes</p> <hr/> <p> <b>NOTE:</b> For the half-life, configure a value that is less than the max-suppress. If you do not, the configuration is rejected.</p> <hr/> <p><b>max-suppress <i>minutes</i></b>—Maximum hold-down time. <i>minutes</i> is the maximum time that a route can be suppressed no matter how unstable it has been.</p> <p><b>Range:</b> 1 through 720</p> <p><b>Default:</b> 60 minutes</p> <hr/> <p> <b>NOTE:</b> For the max-suppress, configure a value that is greater than the half-life. If you do not, the configuration is rejected.</p> <hr/> <p><b><i>name</i></b>—Name that identifies the set of damping parameters. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p><b>reuse <i>number</i></b>—Reuse threshold. <i>number</i> is the figure-of-merit value below which a suppressed route can be used again.</p> <p><b>Range:</b> 1 through 20,000</p> <p><b>Default:</b> 750 (unitless)</p>

**suppress *number***—Cutoff (suppression) threshold. *number* is the figure-of-merit value above which a route is suppressed for use or inclusion in advertisements.

**Range:** 1 through 20,000

**Default:** 3000 (unitless)

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring BGP Flap Damping Parameters*
- *Example: Configuring BGP Route Flap Damping Parameters*
- *Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family*

---

## dynamic-db

**Syntax** dynamic-db;

**Hierarchy Level** [edit logical-systems *logical-system-name* policy-options **as-path** *path-name*],  
[edit logical-systems *logical-system-name* policy-options **as-path-group** *group-name*],  
[edit logical-systems *logical-system-name* policy-options **community** *community-name*],  
[edit logical-systems *logical-system-name* policy-options **condition** *condition-name*],  
[edit logical-systems *logical-system-name* policy-options **policy-statement** *policy-statement-name*],  
[edit logical-systems *logical-system-name* policy-options **prefix-list** *prefix-list-name*],  
[edit policy-options **as-path** *path-name*],  
[edit policy-options **as-path-group** *group-name*],  
[edit policy-options **community** *community-name*],  
[edit policy-options **condition** *condition-name*],  
[edit policy-options **policy-statement** *policy-statement-name*],  
[edit policy-options **prefix-list** *prefix-list-name*]

**Release Information** Statement introduced in Junos OS Release 9.5.  
Statement introduced in Junos OS Release 9.5 for EX Series switches.

**Description** Define routing policies and policy objects that reference policies configured in the dynamic database at the **[edit dynamic]** hierarchy level.


**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control-level—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Dynamic Routing Policies*



## excess-burst-size

<b>Syntax</b>	<code>excess-burst-size bytes;</code>
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2 for EX Series switches.
<b>Description</b>	Configure the maximum number of bytes allowed for incoming packets to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red).
<div>  <p><b>NOTE:</b> When you include the <code>excess-burst-size</code> statement in the configuration, you must also include the <code>committed-burst-size</code> and <code>committed-information-rate</code> statements at the same hierarchy level.</p> </div>	
<b>Options</b>	<p><b>bytes</b>—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1500 through 100,000,000,000 bytes</p>
<b>Required Privilege Level</b>	<p><code>firewall</code>—To view this statement in the configuration.</p> <p><code>firewall-control</code>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Tricolor Marking Policers on page 4827</a></li> </ul>

## family (Firewall Filter)

---

**Syntax**    family *family-name* {  
              filter *filter-name* {  
                  interface-specific;  
                  term *term-name* {  
                      from {  
                        match-conditions;  
                      }  
                      then {  
                        action;  
                        action-modifiers;  
                      }  
                  }  
              }  
          }

**Hierarchy Level**    [edit [firewall](#)]

**Release Information**    Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Option [interface-specific](#) introduced in Junos OS Release 9.5 for EX Series switches.

**Description**    Configure a firewall filter for IP version 4 or IP version 6.

**Options**    *family-name*—Version or type of addressing protocol:

- **any**—Filter packets based on protocol-independent match conditions.
- **ethernet-switching**—Filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets.
- **inet**—Filter IPv4 packets.
- **inet6**—Filter IPv6 packets.

The remaining statements are explained separately.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**

- [Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
- [Configuring Firewall Filters \(J-Web Procedure\) on page 4813](#)
- [Firewall Filters for EX Series Switches Overview on page 4696](#)

## filter (Firewall Filters)

<b>Syntax</b>	<pre>filter <i>filter-name</i> {     <i>interface-specific</i>;     term <i>term-name</i> {         from {             <i>match-conditions</i>;         }         then {             <i>action</i>;             <i>action-modifiers</i>;         }     } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">firewall family</a> <i>family-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Option <a href="#">interface-specific</a> introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Configure firewall filters.
<b>Options</b>	<p><b><i>filter-name</i></b>—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704</a></li> <li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li> <li>• <a href="#">Configuring Firewall Filters (CLI Procedure) on page 4804</a></li> <li>• <a href="#">Configuring Firewall Filters (J-Web Procedure) on page 4813</a></li> <li>• <a href="#">Firewall Filters for EX Series Switches Overview on page 4696</a></li> </ul>

## filter (VLANs)

---

<b>Syntax</b>	<code>filter (input   output) <i>filter-name</i>;</code>
<b>Hierarchy Level</b>	[edit <b>vlan</b> <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Apply a firewall filter to traffic coming into or exiting from the VLAN.
<b>Default</b>	All incoming traffic is accepted unmodified to the VLAN, and all outgoing traffic is sent unmodified from the VLAN.
<b>Options</b>	<p><b><i>filter-name</i></b> —Name of a firewall filter defined in a <b>filter</b> statement.</p> <ul style="list-style-type: none"><li>• <b>input</b>—Apply a firewall filter to VLAN ingress traffic.</li><li>• <b>output</b>—Apply a firewall filter to VLAN egress traffic.</li></ul>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li><li>• <a href="#">Configuring Firewall Filters (CLI Procedure) on page 4804</a></li><li>• <a href="#">Configuring Firewall Filters (J-Web Procedure) on page 4813</a></li><li>• <a href="#">Firewall Filters for EX Series Switches Overview on page 4696</a></li><li>• <a href="#">Configuring VLANs for EX Series Switches (CLI Procedure) on page 2337</a></li></ul>

---

## filter-specific

---

<b>Syntax</b>	filter-specific;
<b>Hierarchy Level</b>	[edit <a href="#">firewall policer</a> <i>policer-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Configure a policer to act as a filter-specific policer. If you do not specify the <b>filter-specific</b> statement, the policer acts as a term-specific policer by default.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li><li>• <a href="#">Configuring Policers to Control Traffic Rates (CLI Procedure) on page 4818</a></li><li>• <a href="#">Understanding the Use of Policers in Firewall Filters on page 4767</a></li></ul>

## firewall

---

```
Syntax  firewall {
        family family-name {
            filter filter-name {
                interface-specific;
                term term-name {
                    from {
                        match-conditions;
                    }
                    then {
                        action;
                        action-modifiers;
                    }
                }
            }
        }
        policer policer-name {
            filter-specific;
            if-exceeding {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
            then {
                policer-action;
            }
        }
    }
    three-color-policer policer-name {
        action {
            loss-priority high then discard;
        }
        single-rate {
            (color-aware | color-blind);
            committed-information-rate bps;
            committed-burst-size bytes;
            excess-burst-size bytes;
        }
        two-rate {
            (color-aware | color-blind);
            committed-information-rate bps;
            committed-burst-size bytes;
            peak-information-rate bps;
            peak-burst-size bytes;
        }
    }
```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Options **interface-specific** and **filter-specific** introduced in Junos OS Release 9.5 for EX Series switches.

<b>Description</b>	Configure firewall filters and policers.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704</a></li> <li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li> <li>• <a href="#">Configuring Firewall Filters (CLI Procedure) on page 4804</a></li> <li>• <a href="#">Configuring Policers to Control Traffic Rates (CLI Procedure) on page 4818</a></li> <li>• <a href="#">Firewall Filters for EX Series Switches Overview on page 4696</a></li> </ul>

---

## from

---

<b>Syntax</b>	from { <i>match-conditions</i> ; }
<b>Hierarchy Level</b>	[edit <b>firewall</b> <i>family family-name filter filter-name term term-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Match packet fields to values specified in a match condition. If the <b>from</b> statement is not included in a firewall filter configuration, all packets are considered to match and the actions and action modifiers in the <b>then</b> statement are taken.
<b>Options</b>	<b>match-conditions</b> —Conditions that define the values or fields that the incoming or outgoing packets must contain for a match. You can specify one or more match conditions. If you specify more than one, they all must match for a match to occur and for the action in the <b>then</b> statement to be taken.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704</a></li> <li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li> <li>• <a href="#">Configuring Firewall Filters (CLI Procedure) on page 4804</a></li> <li>• <a href="#">Configuring Firewall Filters (J-Web Procedure) on page 4813</a></li> <li>• <a href="#">Understanding Firewall Filter Match Conditions on page 4762</a></li> </ul>

## if-exceeding

---

<b>Syntax</b>	<pre>if-exceeding {     bandwidth-limit <i>bps</i>;     bandwidth-percent <i>percent</i>     burst-size-limit <i>bytes</i>; }</pre>
<b>Hierarchy Level</b>	[edit <b>firewall</b> <b>policer</b> <i>policer-name</i> ] [edit logical-systems logical-system-name firewall policer policer-name]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Logical systems support introduced in Junos OS Release 9.3.
<b>Description</b>	Configure policer rate limits.  The <b>bandwidth-percent</b> statement is supported on routers only.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li><li>• <a href="#">Configuring Policers to Control Traffic Rates (CLI Procedure) on page 4818</a></li><li>• <a href="#">Understanding the Use of Policers in Firewall Filters on page 4767</a></li><li>• <i>Basic Single-Rate Two-Color Policers</i></li></ul>



## interface-specific


<b>Syntax</b>	interface-specific;
<b>Hierarchy Level</b>	[edit <b>firewall family</b> <i>family-name</i> <b>filter</b> <i>filter-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Configure firewall counters that are interface-specific. You can configure an interface-specific firewall filter only on a port or a Layer 3 interface as an interface-specific firewall filter is not supported for a VLAN.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704</a></li> <li>• <a href="#">Configuring Firewall Filters (CLI Procedure) on page 4804</a></li> <li>• <a href="#">Configuring Firewall Filters (J-Web Procedure) on page 4813</a></li> <li>• <a href="#">Firewall Filters for EX Series Switches Overview on page 4696</a></li> </ul>

## loss-priority high then discard (Three-Color Policer)

<b>Syntax</b>	loss-priority high then discard;
<b>Hierarchy Level</b>	[edit <b>firewall three-color-policer</b> <i>policer-name</i> <b>action</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2 for EX Series switches.
<b>Description</b>	<p>For packets with high loss priority, discard the packets. The loss priority setting is implicit and cannot be configured. Include this statement if you do not want the local switch to forward packets that have high packet loss priority.</p> <p>For single-rate three-color policers, Junos OS assigns high loss priority to packets that exceed the committed information rate and the excess burst size.</p> <p>For two-rate three-color policers, Junos OS assigns high loss priority to packets that exceed the peak information rate and the peak burst size.</p>
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Tricolor Marking Policers on page 4827</a></li> </ul>

## peak-burst-size

---

<b>Syntax</b>	<code>peak-burst-size bytes;</code>
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2 for EX Series switches.
<b>Description</b>	Configure the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red).
<div> <b>NOTE:</b> When you include the <code>peak-burst-size</code> statement in the configuration, you must also include the <code>committed-burst-size</code> and <code>peak-information-rate</code> statements at the same hierarchy level.</div>	
<b>Options</b>	<b>bytes</b> —Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000). <b>Range:</b> 1500 through 100,000,000,000 bytes
<b>Required Privilege Level</b>	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Tricolor Marking Policers on page 4827</a></li></ul>

## policer

<b>Syntax</b>	<pre> policer <i>policer-name</i> {   counter {     counter-id <i>counter-index</i>;   }   filter-specific;   if-exceeding {     bandwidth-limit <i>bps</i>;     bandwidth-percent <i>percent</i>     burst-size-limit <i>bytes</i>;   }   then {     <i>policer-action</i>;   } } </pre>
<b>Hierarchy Level</b>	<p>[edit <a href="#">firewall</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> firewall]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Logical systems support introduced in Junos OS Release 9.3.</p>
<b>Description</b>	<p>Configure policer rate limits and actions. To activate a policer, you must include the <b>policer</b> action modifier in the <b>then</b> statement in a firewall filter term. Except for EX8200 switches, each policer that you configure includes an implicit counter. To obtain term-specific packet counts, configure a policer for each term in the filter that requires policing. For EX8200 switches, configure a policer and associate it with a global management counter using the <i>counter</i> option.</p>
<b>Options</b>	<p><b><i>policer-name</i></b>—Name that identifies the policer. The name can include letters, numbers, hyphens (-), and can contain up to 64 characters.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li> <li>• <a href="#">Example: Configuring CoS on EX Series Switches on page 2075</a></li> <li>• <a href="#">Configuring Policers to Control Traffic Rates (CLI Procedure) on page 4818</a></li> <li>• <a href="#">Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure)</a></li> <li>• <a href="#">Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)</a></li> <li>• <a href="#">Understanding the Use of Policers in Firewall Filters on page 4767</a></li> <li>• <a href="#">Basic Single-Rate Two-Color Policers</a></li> </ul>

## policy-options

---

### Syntax

**Hierarchy Level**    [\[edit\]](#)

**Release Information**    Statement introduced in Junos OS Release 12.1 for the EX Series.

**Description**    Configure options such as application maps for DCBX application protocol exchange and policy statements.

**Required Privilege Level**    storage—To view this statement in the configuration.  
storage-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring DCBX to Support an iSCSI Application*

## policy-statement

<b>Syntax</b>	<pre> policy-statement <i>policy-name</i> {   term <i>term-name</i> {     from {       family <i>family-name</i>;       match-conditions;       policy <i>subroutine-policy-name</i>;       prefix-list <i>prefix-list-name</i>;       prefix-list-filter <i>prefix-list-name</i> match-type &lt;actions&gt;;       route-filter <i>destination-prefix</i> match-type &lt;actions&gt;;       source-address-filter <i>source-prefix</i> match-type &lt;actions&gt;;     }     to {       match-conditions;       policy <i>subroutine-policy-name</i>;     }     then <i>actions</i>;   } } </pre>
<b>Hierarchy Level</b>	<p>[edit dynamic policy-options],</p> <p>[edit logical-systems <i>logical-system-name</i> policy-options],</p> <p>[edit policy-options]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>inet-mdt</b> option introduced in Junos OS Release 10.0R2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>route-target</b> option introduced in Junos OS Release 12.2.</p>
<b>Description</b>	<p>Define a routing policy, including subroutine policies.</p> <p>A <i>term</i> is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.</p> <p>Each term contains a set of match conditions and a set of actions:</p> <ul style="list-style-type: none"> <li>Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.</li> <li>Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.</li> </ul> <p>Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of</p>

**accept** or **reject** is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the **[edit policy-options]** hierarchy level by **policy-statement *policy-name*** in alphabetical order, enter the **show policy-options** configuration command.

**Options** *actions*—(Optional) One or more actions to take if the conditions match. The actions are described in *Configuring Flow Control Actions*.

**family** *family-name*—(Optional) Specify an address family protocol. Specify **inet** for IPv4. Specify **inet6** for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify **iso**. For IPv4 multicast VPN traffic, specify **inet-mvpn**. For IPv6 multicast VPN traffic, specify **inet6-mvpn**. For multicast-distribution-tree (MDT) IPv4 traffic, specify **inet-mdt**. For BGP route target VPN traffic, specify **route-target**.



**NOTE:** When *family* is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the *family* statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

**from**—(Optional) Match a route based on its source address.

**match-conditions**—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in *Routing Policy Match Conditions*.

**policy** *subroutine-policy-name*—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policy names cannot take the form **\_\_.\*-internal\_\_**, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

**policy-name**—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").

**prefix-list** *prefix-list-name*—Name of a list of IPv4 or IPv6 prefixes.

**prefix-list-filter** *prefix-list-name*—Name of a prefix list to evaluate using qualifiers; *match-type* is the type of match (see *Configuring Prefix List Filters*), and *actions* is the action to take if the prefixes match.

**route-filter** *destination-prefix match-type <actions>*—(Optional) List of routes on which to perform an immediate match; *destination-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see *Configuring Route Lists*), and *actions* is the action to take if the *destination-prefix* matches.

**source-address-filter** *source-prefix match-type <actions>*—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. *source-prefix* is

the IPv4 or IPv6 route prefix to match, **match-type** is the type of match (see *Configuring Route Lists*), and **actions** is the action to take if the **source-prefix** matches.

**term term-name**—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the **term** statement when defining match conditions and actions.

**to**—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

**then**—(Optional) Actions to take on matching routes. The actions are described in *Configuring Flow Control Actions* and *Configuring Actions That Manipulate Route Characteristics*.

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">dynamic-db on page 4848</a></li></ul>
------------------------------	---



## prefix-list

<b>Syntax</b>	<pre>prefix-list <i>name</i> {     <i>ip-addresses</i>;     apply-path <i>path</i>; }</pre>
<b>Hierarchy Level</b>	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Support for the <b>vpls</b> protocol family introduced in Junos OS Release 10.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>Define a list of IPv4 or IPv6 address prefixes for use in a routing policy statement or firewall filter statement.</p> <p>You can configure up to 85,325 prefixes in each prefix list. To configure more than 85,325 prefixes, configure multiple prefix lists and apply them to multiple firewall filter terms.</p>
<b>Options</b>	<p><b><i>name</i></b>—Name that identifies the list of IPv4 or IPv6 address prefixes.</p> <p><b><i>ip-addresses</i></b>—List of IPv4 or IPv6 address prefixes, one IP address per line in the configuration.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding Prefix Lists for Use in Routing Policy Match Conditions</i></li> <li>• <i>Firewall Filter Match Conditions Based on Address Fields</i></li> <li>• <i>Example: Configuring Routing Policy Prefix Lists</i></li> <li>• <i>Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List</i></li> </ul>

## routing-instance

---

<b>Syntax</b>	<code>routing-instance <i>routing-instance-name</i>;</code>
<b>Hierarchy Level</b>	[edit <code>firewall</code> family inet <code>filter filter-name</code> <code>term term-name</code> <code>then</code> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4 for EX Series switches.
<b>Description</b>	Specify a specific virtual routing instance to which the switch sends matched packets.
<b>Options</b>	<i>routing-instance-name</i> —Name of a virtual routing instance.
<b>Required Privilege Level</b>	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 4795</a></li><li>• <a href="#">Configuring Virtual Routing Instances (CLI Procedure) on page 2347</a></li><li>• <a href="#">Understanding Filter-Based Forwarding for EX Series Switches on page 4770</a></li></ul>

## single-rate

<b>Syntax</b>	<pre>single-rate {   (color-aware   color-blind);   committed-burst-size bytes;   committed-information-rate bps;   excess-burst-size bytes; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2 for EX Series switches.
<b>Description</b>	<p>Configure a single-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).</p> <p>Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but do not exceed the EBS are assigned medium-high loss priority (yellow). Packets that exceed the EBS are assigned high loss priority (red).</p> <p>Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Tricolor Marking Policers on page 4827</a></li> </ul>

## term

---

<b>Syntax</b>	<pre>term <i>term-name</i> {     from {         <i>match-conditions</i>;     }     then {         <i>action</i>;         <i>action-modifiers</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit <b>firewall family</b> <i>family-name</i> <b>filter</b> <i>filter-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Define a firewall filter term.
<b>Options</b>	<p><b><i>term-name</i></b>—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704</a></li><li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li><li>• <a href="#">Configuring Firewall Filters (CLI Procedure) on page 4804</a></li><li>• <a href="#">Configuring Firewall Filters (J-Web Procedure) on page 4813</a></li><li>• <a href="#">Firewall Filters for EX Series Switches Overview on page 4696</a></li></ul>

## then (Firewall Filters)

---

<b>Syntax</b>	<pre>then {     action;     action-modifiers; }</pre>
<b>Hierarchy Level</b>	[edit <b>firewall family</b> <i>family-name</i> <b>filter</b> <i>filter-name</i> <b>term</b> <i>term-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure a filter action.
<b>Options</b>	<p><b>action</b>—Action to accept, discard, or forward packets that match all match conditions specified in a filter term.</p> <p><b>action-modifiers</b>—Additional actions to analyze, classify, count, or police packets that match all conditions specified in a filter term.</p>
<b>Required Privilege Level</b>	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches on page 4704</a></li> <li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li> <li>• <a href="#">Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on EX Series Switches on page 4795</a></li> <li>• <a href="#">Configuring Firewall Filters (CLI Procedure) on page 4804</a></li> <li>• <a href="#">Configuring Firewall Filters (J-Web Procedure) on page 4813</a></li> <li>• <a href="#">Understanding Firewall Filter Match Conditions on page 4762</a></li> </ul>

## then (Policer Action)

---

<b>Syntax</b>	<pre>then {     <i>policer-action</i>; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">firewall policer <i>policer-name</i></a> ] [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure a policer action.
<b>Options</b>	<p><i>policer-action</i>—Actions to take are:</p> <ul style="list-style-type: none"><li>• <b>discard</b>—Discard traffic that exceeds the rate limits defined by the policer.</li><li>• <b>forwarding-class <i>class-name</i></b>—For routers only, classify traffic that exceeds the rate limits defined by the policer.</li><li>• <b>loss-priority</b>—Set the loss priority for traffic that exceeds the rate limits defined by the policer.</li></ul>
<b>Required Privilege Level</b>	firewall—To view this statement in the configuration. firewall -control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li><li>• <a href="#">Configuring Policers to Control Traffic Rates (CLI Procedure) on page 4818</a></li><li>• <a href="#">Configuring Firewall Filters (CLI Procedure) on page 4804</a></li><li>• <a href="#">Configuring Firewall Filters (J-Web Procedure) on page 4813</a></li><li>• <a href="#">Understanding the Use of Policers in Firewall Filters on page 4767</a></li><li>• <a href="#">Example: Configuring CoS for a PBB Network on MX Series Routers</a></li><li>• <a href="#">Basic Single-Rate Two-Color Policers</a></li></ul>

## three-color-policer (Configuring)

```
Syntax  three-color-policer policer-name {
        action {
            loss-priority high then discard;
        }
        single-rate {
            (color-aware | color-blind);
            committed-burst-size bytes;
            committed-information-rate bps;
            excess-burst-size bytes;
        }
        two-rate {
            (color-aware | color-blind);
            committed-burst-size bytes;
            committed-information-rate bps;
            peak-burst-size bytes;
            peak-information-rate bps;
        }
    }
```

**Hierarchy Level** [edit [firewall](#)]

**Release Information** Statement introduced in Junos OS Release 11.2 for EX Series switches.

**Description** Configure a three-color policer.

**Options** *policer-name*—Name of the three-color policer. Reference this name when you apply the policer to an interface.

The remaining statements are explained separately.

**Required Privilege Level** firewall—To view this statement in the configuration.  
firewall-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Tricolor Marking Policers on page 4827](#)

## two-rate

---

<b>Syntax</b>	<pre>two-rate {   (color-aware   color-blind);   committed-burst-size bytes;   committed-information-rate bps;   peak-burst-size bytes;   peak-information-rate bps; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> ],
<b>Release Information</b>	Statement introduced in Junos OS Release 11.2 for EX Series switches.
<b>Description</b>	<p>Configure a two-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), peak information rate (PIR), and peak burst size (PBS).</p> <p>Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the PIR and the PBS are assigned high loss priority (red).</p> <p>Green packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Tricolor Marking Policers on page 4827</a></li></ul>



# Administration

- [Routine Monitoring on page 4873](#)
- [Operational Commands on page 4876](#)

## Routine Monitoring

---

- [Verifying That Firewall Filters Are Operational on page 4873](#)
- [Verifying That Policers Are Operational on page 4874](#)
- [Monitoring Firewall Filter Traffic on page 4875](#)

### Verifying That Firewall Filters Are Operational

**Purpose** After you configure and apply firewall filters to ports, VLANs, or Layer 3 interfaces, you can perform the following task to verify that the firewall filters configured on EX Series switches are working properly.

**Action** Use the operational mode command to verify that the firewall filters on the switch are working properly:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                               Bytes          Packets
counter-employee-web                0              0
Filter: ingress-port-voip-class-limit-tcp-icmp
Counters:
Name                               Bytes          Packets
icmp-counter                        0              0
Policers:
Name                               Packets
icmp-connection-policer            0
tcp-connection-policer             0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

**Meaning** The **show firewall** command displays the names of all firewall filters, policers, and counters that are configured on the switch. For each counter that is specified in a filter configuration, the output field shows the byte count and packet count for the term in which the counter is specified. For each policer that is specified in a filter configuration, the output field shows the packet count for packets that exceed the specified rate limits.

- Related Documentation**
- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
  - [Configuring Firewall Filters \(J-Web Procedure\) on page 4813](#)
  - [Configuring Policers to Control Traffic Rates \(CLI Procedure\) on page 4818](#)
  - [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
  - [Monitoring Firewall Filter Traffic on page 4875](#)

## Verifying That Policers Are Operational

**Purpose** After you configure policers and include them in firewall filter configurations, you can perform the following tasks to verify that the policers configured on EX Series switches are working properly.

**Action** Use the operational mode command to verify that the policers on the switch are working properly:

```
user@switch> show policer
Filter: egress-vlan-watch-employee
Filter: ingress-port-filter
Filter: ingress-port-voip-class-limit-tcp-icmp
Policies:
Name                                     Packets
icmp-connection-policer                  0
tcp-connection-policer                  0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

**Meaning** The **show policer** command displays the names of all firewall filters and policers that are configured on the switch. For each policer that is specified in a filter configuration, the output field shows the current packet count for all packets that exceed the specified rate limits.

- Related Documentation**
- [Configuring Policers to Control Traffic Rates \(CLI Procedure\) on page 4818](#)
  - [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
  - [Configuring Firewall Filters \(J-Web Procedure\) on page 4813](#)
  - [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
  - [Monitoring Firewall Filter Traffic on page 4875](#)

## Monitoring Firewall Filter Traffic

You can monitor firewall filter traffic on EX Series switches.

- [Monitoring Traffic for All Firewall Filters and Policers That Are Configured on the Switch on page 4875](#)
- [Monitoring Traffic for a Specific Firewall Filter on page 4875](#)
- [Monitoring Traffic for a Specific Policar on page 4876](#)

### Monitoring Traffic for All Firewall Filters and Policers That Are Configured on the Switch

**Purpose** Perform the following task to monitor the number of packets and bytes that matched the firewall filters and monitor the number of packets that exceeded policer rate limits:

**Action** Use the operational mode command:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                               Bytes          Packets
counter-employee-web               3348            27
Filter: ingress-port-voip-class-limit-tcp-icmp
Counters:
Name                               Bytes          Packets
icmp-counter                       4100            49
Policers:
Name                               Packets
icmp-connection-policer           0
tcp-connection-policer            0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

**Meaning** The **show firewall** command displays the names of all firewall filters, policers, and counters that are configured on the switch. The output fields show byte and packet counts for counters and packet count for policers.

### Monitoring Traffic for a Specific Firewall Filter

**Purpose** Perform the following task to monitor the number of packets and bytes that matched a firewall filter and monitor the number of packets that exceeded the policer rate limits.

**Action** Use the operational mode command:

```
user@switch> show firewall filter ingress-vlan-rogue-block
Filter: ingress-vlan-rogue-block
Counters:
Name                               Bytes          Packets
rogue-counter                       2308            20
```

**Meaning** The **show firewall filter *filter-name*** command displays the name of the firewall filter, the packet and byte count for all counters configured with the filter, and the packet count for all policers configured with the filter.

### Monitoring Traffic for a Specific Policer

---

**Purpose** Perform the following task to monitor the number of packets that exceeded policer rate limits:

**Action** Use the operational mode command:

```
user@switch> show policer tcp-connection-policer
Filter: ingress-port-voip-class-limit-tcp-icmp
Policers:
Name                                     Packets
tcp-connection-policer                  0
```

**Meaning** The **show policer *policer-name*** command displays the name of the firewall filter that specifies the policer-action and displays the number of packets that exceeded rate limits for the specified filter.


- Related Documentation**
- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
  - [Configuring Firewall Filters \(J-Web Procedure\) on page 4813](#)
  - [Configuring Policers to Control Traffic Rates \(CLI Procedure\) on page 4818](#)
  - [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
  - [Verifying That Firewall Filters Are Operational on page 4873](#)

## Operational Commands

---

- [clear firewall](#)
- [show firewall](#)
- [show firewall log](#)
- [show policer](#)
- [show policy](#)
- [show policy conditions](#)
- [test policy](#)

## clear firewall

<b>List of Syntax</b>	<a href="#">Syntax on page 4877</a> <a href="#">Syntax (EX Series Switches) on page 4877</a>
<b>Syntax</b>	clear firewall (all   counter <i>counter-name</i>   filter <i>filter-name</i>   log (all   <i>logical-system-name</i> )   logical-system <i>logical-system-name</i> )
<b>Syntax (EX Series Switches)</b>	clear firewall (all   counter <i>counter-name</i>   filter <i>filter-name</i>   log (all   <i>logical-system-name</i> )   policer counter (all   counter-id <i>counter-index</i> ))
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>logical-system</b> option introduced in Junos OS Release 9.3.</p> <p><b>log</b> option introduced before Junos OS Release 11.4.</p>
<b>Description</b>	<p>Clear statistics about configured firewall filters.</p> <p>When you clear the counters of a filter, this impacts not only the counters shown by the CLI, but also the ones tracked by SNMP2.</p> <p>Subscriber management uses firewall filters to capture and report the volume-based service accounting counters that are used for subscriber billing. The <b>clear firewall</b> command also clears the service accounting counters that are reported to the RADIUS accounting server. For this reason, you must be cautious in specifying which firewall statistics you want to clear.</p>
<div>  <p><b>NOTE:</b> The <b>clear firewall</b> command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover (GRES).</p> </div>	
<p>If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the <b>prefix-action</b> action on matched packets, wait at least 5 seconds before you enter the <b>show firewall prefix-action-stats</b> command. A 5-second pause between issuing the <b>clear firewall</b> and <b>show firewall prefix-action-stats</b> commands avoids a possible timeout of the <b>show firewall prefix-action-stats</b> command.</p>	
<b>Options</b>	<p><b>all</b>—Clear the packet and byte counts for all filters. On EX Series switches, this option also clears the packet counts for all policer counters.</p> <p><b>counter <i>counter-name</i></b>—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.</p> <p><b>filter <i>filter-name</i></b>—Clear the packet and byte counts for the specified firewall filter.</p> <p><b>log (all   <i>logical-system-name</i>)</b>—Clear log entries for IPv4 firewall filters that have <b>then log</b> as an action. Use <b>log all</b> to clear all log entries or <b>log <i>logical-system-name</i></b> to clear log entries for the specified logical system.</p>

**logical-system** *logical-system-name*—Clear the packet and byte counts for the specified logical system.

**policer counter** (**all** | **counter-id** *counter-index*)—(EX8200 switches only) Clear all policer counters using the **policer counter all** command, or clear a specific policer counter using the **policer counter counter-id** *counter-index* command. The value of *counter-index* can be 0, 1, or 2.

**Required Privilege Level**

clear

**Related Documentation**

- [show firewall on page 4879](#)

**List of Sample Output**

[clear firewall all on page 4878](#)

[clear firewall \(counter counter-name\) on page 4878](#)

[clear firewall \(filter filter-name\) on page 4878](#)

[clear firewall \(policer counter all\) \(EX8200 Switch\) on page 4878](#)

[clear firewall \(policer counter counter-id counter-index\) \(EX8200 Switch\) on page 4878](#)

## Sample Output

clear firewall all

```
user@host> clear firewall all
```

clear firewall (counter counter-name)

```
user@host> clear firewall counter port-filter-counter
```

clear firewall (filter filter-name)

```
user@host> clear firewall filter ingress-port-filter
```

clear firewall (policer counter all) (EX8200 Switch)

```
user@switch> clear firewall policer counter all
```

clear firewall (policer counter counter-id counter-index) (EX8200 Switch)

```
user@switch> clear firewall policer counter counter-id 0
```

## show firewall

<b>List of Syntax</b>	<a href="#">Syntax on page 4879</a> <a href="#">Syntax (EX Series Switches) on page 4879</a>
<b>Syntax</b>	<pre>show firewall &lt;counter <i>counter-name</i>&gt; &lt;detail&gt; &lt;filter <i>filter-name</i>&gt; &lt;log&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;terse&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show firewall &lt;counter <i>counter-name</i>&gt; &lt;detail&gt; &lt;filter <i>filter-name</i>&gt; &lt;log &lt;(detail   interface <i>interface-name</i>)&gt;&gt; &lt;policer counters &lt;(detail   counter-id <i>counter-index</i> &lt;detail&gt;)&gt;&gt; &lt;terse&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>logical-system</b> introduced in Junos OS Release 9.3.</p> <p>Option <b>terse</b> introduced in Junos OS Release 9.4.</p> <p>Option <b>policer counters</b> introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>Option <b>detail</b> introduced in Junos OS Release 12.3 for EX Series switches.</p> <p>Option <b>detail</b> introduced in Junos OS Release 14.1 for MX Series routers.</p>
<b>Description</b>	Display enhanced statistics and counters for all configured firewall filters.
<b>Options</b>	<p><b>none</b>—(Optional) Display statistics and counters for all configured firewall filters and counters. For EX Series switches, this command also displays statistics about all configured policers.</p> <p><b>counter <i>counter-name</i></b>—(Optional) Name of a filter counter.</p> <p><b>detail</b>—(EX Series switches and MX Series routers only) (Optional) Display firewall filter statistics and enhanced policer statistics and counters.</p> <p><b>filter <i>filter-name</i></b>—(Optional) Name of a configured filter.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>log</b>—(Optional) Display log entries for firewall filters.</p> <p><b>log &lt;(detail   interface <i>interface-name</i>)&gt;</b>—(EX Series switches only) (Optional) Display detailed log entries of firewall activity or log information about a specific interface.</p> <p><b>policer counters &lt;(detail   counter-id <i>counter-index</i> &lt;detail&gt;)&gt;</b>—(EX8200 switches only) (Optional) Display policer counter statistics in brief or in detail.</p>

**terse**—(Optional) Display firewall filter names only.

**Required Privilege Level** view

- Related Documentation**
- [clear firewall on page 4877](#)
  - [show firewall log on page 4886](#)
  - *show policer*
  - *Enhanced Policer Statistics Overview*
  - *enhanced-policer*

- List of Sample Output**
- [show firewall filter \(MX Series Router and EX Series Switch\) on page 4883](#)
  - [show firewall filter \(non MX Series Router and EX Series Switch\) on page 4883](#)
  - [show firewall filter \(Dynamic Input Filter\) on page 4883](#)
  - [show firewall \(Logical Systems\) on page 4883](#)
  - [show firewall \(counter counter-name\) on page 4884](#)
  - [show firewall log on page 4884](#)
  - [show firewall policer counters \(EX8200 Switch\) on page 4884](#)
  - [show firewall policer counters \(detail\) \(EX8200 Switch\) on page 4884](#)
  - [show firewall policer counters \(counter-id counter-index\) \(EX8200 Switch\) on page 4885](#)
  - [show firewall policer counters \(counter-id counter-index detail\) \(EX8200 Switch\) on page 4885](#)
  - [show firewall detail on page 4885](#)

**Output Fields** Table 539 on page 4880 lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

**Table 539: show firewall Output Fields**

Field Name	Field Description
<b>Filter</b>	<p>Name of a filter that has been configured with the <b>filter</b> statement at the <b>[edit firewall]</b> hierarchy level.</p> <p>Except on EX Series switches:</p> <ul style="list-style-type: none"> <li>• When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either <b>-i</b> for an input filter or <b>-o</b> for an output filter.</li> <li>• When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either <b>-in</b> for an input filter or <b>-out</b> for an output filter. When a logical system-specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, __ls1/filter1).</li> <li>• When a service filter is displayed that uses a service set, the separator between the service-set name and the service-filter name is a semicolon (;).</li> </ul> <p><b>NOTE:</b> For <b>bridge family filter</b>, the <b>ip-protocol</b> match criteria is supported only for IPv4 and not for IPv6. This is applicable for line cards that support the Junos Trio chipset, such as the MX 3D MPC line cards.</p>



Table 539: show firewall Output Fields (*continued*)

Field Name	Field Description
<b>Counters</b>	<p>Display filter counter information:</p> <ul style="list-style-type: none"> <li>• <b>Name</b>—Name of a filter counter that has been configured with the <b>counter</b> firewall filter action.</li> <li>• <b>Bytes</b>—Number of bytes that match the filter term under which the <b>counter</b> action is specified.</li> <li>• <b>Packets</b>—Number of packets that matched the filter term under which the <b>counter</b> action is specified.</li> </ul> <p><b>NOTE:</b> On M and T series routers, firewall filters cannot count <b>ip-options</b> packets on a per option type and per interface basis. A limited work around is to use the <b>show pfe statistics ip options</b> command to see <b>ip-options</b> statistics on a per Packet Forwarding Engine (PFE) basis. See <i>show pfe statistics ip</i> for sample output.</p>
<b>Policers</b>	<p>Display policer information:</p> <ul style="list-style-type: none"> <li>• <b>Name</b>—Name of policer.</li> <li>• <b>Bytes</b>—(For two-color policers on MX Series routers and EX Series switches, and for hierarchical policers on MS-DPC, MIC, and MPC interfaces on MX Series routers) Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer. For other combinations of policer type, device, and line card type, this field is blank.</li> <li>• <b>Packets</b>—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.</li> </ul>
<b>Policer Counter Index</b>	(EX8200 switch only) Global management counter ID. The counter ID value ( <i>counter-index</i> ) can be 0, 1, or 2.
<b>Green</b>	(EX8200 switch only) Number of packets within the limits. The number of packets is smaller than the committed information rate (CIR).
<b>Yellow</b>	(EX8200 switch only) Number of packets partially within the limits. The number of packets is greater than the CIR, but the burst size is within the excess burst size (EBS) limit.
<b>Discard</b>	(EX8200 switch only) Number of discarded packets.
<b>Bytes</b>	(EX8200 switch only) Number of green, yellow, red, or discarded packets in bytes.
<b>Packets</b>	(EX8200 switch only) Number of green, yellow, red, or discarded packets.
<b>Filter name</b>	(EX8200 switch only) Name of the filter with a term associated to a policer.
<b>Term name</b>	(EX8200 switch only) Name of the term associated with a policer.
<b>Policer name</b>	(EX8200 switch only) Name of the policer that is associated with a global management counter.

Table 539: show firewall Output Fields (*continued*)

Field Name	Field Description
PI-t1	<ul style="list-style-type: none"><li>• OOS packet statistics for packets that are marked out-of-specification (out-of-spec) by the policer. Changes to all packets that have out-of-spec actions, such as discard, color marking, or forwarding-class, are included in this counter.</li><li>• Offered packet statistics for traffic subjected to policing.</li><li>• Transmitted packet statistics for traffic that is not discarded by the policer. When the policer action is discard, the statistics are the same as the in-spec statistics; when the policer action is non-discard (loss-priority or forwarding-class), the statistics are included in this counter.</li></ul>

---

## Sample Output

### show firewall filter (MX Series Router and EX Series Switch)

```

user@host> show firewall filter test
Filter: test
Counters:
Name                               Bytes          Packets
Counter-1                         0              0
Counter-2                         0              0
Policers:
Name                               Bytes          Packets
Policer-1                        2770           70

```

### show firewall filter (non MX Series Router and EX Series Switch)

```

user@host> show firewall filter test
Filter: test
Counters:
Name                               Bytes          Packets
Counter-1                         0              0
Counter-2                         0              0
Policers:
Name                               Bytes          Packets
Policer-1                        70

```

### show firewall filter (Dynamic Input Filter)

```

user@host> show firewall filter dfwd-ge-5/0/0.1-in
Filter: dfwd-ge-5/0/0.1-in
Counters:
Name                               Bytes          Packets
c1-ge-5/0/0.1-in                  0              0

```

### show firewall (Logical Systems)

```

user@host> show firewall

Filter: __lr1/test
Counters:
Name                               Bytes          Packets
icmp                               420            5
Filter: __default_bpdu_filter__
Filter: __lr1/inet_filter1
Counters:
Name                               Bytes          Packets
inet_tcp_count                     0              0
inet_udp_count                     0              0
Filter: __lr1/inet_filter2
Counters:
Name                               Bytes          Packets
inet_icmp_count                    0              0
inet_pim_count                     0              0
Filter: __lr2/inet_filter1
Counters:
Name                               Bytes          Packets
inet_tcp_count                     0              0
inet_udp_count                     0              0

```

### show firewall (counter counter-name)

```
user@host> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                                     Bytes      Packets
icmp-counter                             0           0
```

### show firewall log

```
user@host> show firewall log
Log :

Time      Filter  Action Interface  Protocol  Src Addr
      Dest Addr
08:00:53  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:52  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:51  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:50  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:49  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:48  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:47  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
```

### show firewall policer counters (EX8200 Switch)

```
user@switch> show firewall policer counters
Policer Counter Index 0:

          Bytes      Packets
Green:         73      15914
Yellow:         9      1962
Discard:       119     25942

Policer Counter Index 1:

          Bytes      Packets
Green:         0         0
Yellow:         0         0
Discard:         0         0

Policer Counter Index 2:

          Bytes      Packets
Green:         0         0
Yellow:         0         0
Discard:         0         0
```

### show firewall policer counters (detail) (EX8200 Switch)

```
user@switch> show firewall policer counters detail
Policer Counter Index 0:

          Bytes      Packets
Green:         73      15914
Yellow:         9      1962
Discard:       119     25942
```

Filter name	Term name	Policer name
myfilter	polcr-term-1	myfilter-polcr-1
inet-filter-ae	ae-snmp	policer-1
inet-filter-ae	ae-ssh	policer-2

## Policer Counter Index 1:

	Bytes	Packets
Green:	0	0
Yellow:	0	0
Discard:	0	0

Filter name	Term name	Policer name
-------------	-----------	--------------

## Policer Counter Index 2:

	Bytes	Packets
Green:	0	0
Yellow:	0	0
Discard:	0	0

Filter name	Term name	Policer name
-------------	-----------	--------------

**show firewall policer counters (counter-id counter-index) (EX8200 Switch)**

```
user@switch> show firewall policer counters counter-id 0
```

## Policer Counter Index 0:

	Bytes	Packets
Green:	73	15914
Yellow:	9	1962
Discard:	119	25942

**show firewall policer counters (counter-id counter-index detail) (EX8200 Switch)**

```
user@switch> show firewall policer counters counter-id 0 detail
```

## Policer Counter Index 0:

	Bytes	Packets
Green:	73	15914
Yellow:	9	1962
Discard:	119	25942

Filter name	Term name	Policer name
myfilter	polcr-term-1	myfilter-polcr-1
inet-filter-ae	ae-snmp	policer-1
inet-filter-ae	ae-ssh	policer-2

**show firewall detail**

```
user@host> show firewall detail
```

```
Filter: __default_bpdu_filter__
```

```
Filter: foo
```

```
Counters:
```

Name	Bytes	Packets
c1	17652140	160474

```
Policers:
```

Name	Bytes	Packets
P1-t1		
OOS	0	18286
Offered	0	18446744073709376546
Transmitted	0	18446744073709358260

## show firewall log

<b>List of Syntax</b>	<a href="#">Syntax on page 4886</a> <a href="#">Syntax (EX Series Switches) on page 4886</a>
<b>Syntax</b>	<pre>show firewall log &lt;detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (<i>logical-system-name</i>   all)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show firewall log &lt;detail&gt; &lt;interface <i>interface-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>logical-system</b> option introduced in Junos OS Release 9.3.</p>
<b>Description</b>	Display log information about firewall filters.
<b>Options</b>	<p><b>none</b>—Display log information about firewall filters.</p> <p><b>detail</b>—(Optional) Display detailed information.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display log information about a specific interface.</p> <p><b>logical-system (<i>logical-system-name</i>   all)</b>—(Optional) Perform this operation on all logical systems or on a particular system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show firewall log on page 4887</a> <a href="#">show firewall log detail on page 4887</a>
<b>Output Fields</b>	<p><a href="#">Table 540 on page 4886</a> lists the output fields for the <b>show firewall log</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 540: show firewall log Output Fields**

Field Name	Field Description
<b>Time of Log</b>	Time that the event occurred.
<b>Filter</b>	<ul style="list-style-type: none"> <li>Displays the name of a configured firewall filter or service filter only if the packet hit the filter's <b>log</b> action in a kernel filter (in the control plane). For any traffic that reaches the Routing Engine, the packets hit the <b>log</b> action in the kernel.</li> <li>For all other logged packets (packet hit the filter's <b>log</b> action in the Packet Forwarding Engine), this field displays <b>pfe</b> instead of a configured filter name.</li> </ul>

Table 540: show firewall log Output Fields (*continued*)

Field Name	Field Description
Filter Action	Filter action: <ul style="list-style-type: none"> <li>• <b>A</b>—Accept</li> <li>• <b>D</b>—Discard</li> <li>• <b>R</b>—Reject</li> </ul>
Name of Interface	<ul style="list-style-type: none"> <li>• Displays a physical interface name if the packet arrived at a port on a line card.</li> <li>• Displays <b>local</b> if the packet was generated by the device's internal Ethernet interface, <b>em1</b> or <b>fxp1</b>, which connects the Routing Engine with the router's packet-forwarding components.</li> </ul>
Name of protocol	Packet's protocol name: <b>egp</b> , <b>gre</b> , <b>icmp</b> , <b>ipip</b> , <b>ospf</b> , <b>pim</b> , <b>rsvp</b> , <b>tcp</b> , or <b>udp</b> .
Packet length	Length of the packet.
Source address	Packet's source address.
Destination address	Packet's destination address and port.

## Sample Output

### show firewall log

```
user@host>show firewall log
Time      Filter  Action Interface    Protocol  Src Addr    Dest Addr
13:10:12  pfe      D      rlsq0.902     ICMP     180.1.177.2 180.1.177.1
13:10:11  pfe      D      rlsq0.902     ICMP     180.1.177.2 180.1.177.1
```

### show firewall log detail

```
user@host> show firewall log detail
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
```

```
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,  
Destination address: 192.168.70.66:513  
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of  
interface: fxp0.0  
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,  
Destination address: 192.168.70.66:513  
....
```



## show policer

<b>Syntax</b>	<b>show policer</b> <b>&lt;policer-name&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Display statistics about configured policers.
<b>Options</b>	<b>none</b> —Display the count of policed packets for all configured policers in the system. <b>policer-name</b> —(Optional) Display the count of policed packets for the specified policer.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773</a></li> <li>• <a href="#">Verifying That Firewall Filters Are Operational on page 4873</a></li> <li>• <a href="#">Verifying That Policers Are Operational on page 4874</a></li> <li>• <a href="#">Firewall Filters for EX Series Switches Overview on page 4696</a></li> <li>• <a href="#">Understanding the Use of Policers in Firewall Filters on page 4767</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show policer on page 4889</a> <a href="#">show policer (policer-name) on page 4890</a>
<b>Output Fields</b>	<a href="#">Table 541 on page 4889</a> lists the output fields for the <b>show policer</b> command. Output fields are listed in the approximate order in which they appear.

**Table 541: show policer Output Fields**

Field Name	Field Description	Level of Output
<b>Filter</b>	Name of filter that is configured with the <b>filter</b> statement at the <b>[edit firewall]</b> hierarchy level.	All levels
<b>Policers</b>	Display policer information: <ul style="list-style-type: none"> <li>• <b>Filter</b>—Name of filter that specifies the policer action.</li> <li>• <b>Name</b>—Name of policer.</li> <li>• <b>Packets</b>—Number of packets that matched the filter term where the policer action is specified. This is the number of packets that exceed the rate limits that the policer specifies.</li> </ul>	All levels

## Sample Output

### show policer

```
user@host> show policer
```

```
Filter: egress-vlan-filter
Filter: ingress-port-filter
Policers:
Name                               Packets
icmp-connection-policer           0
tcp-connection-policer            0
Filter: ingress-vlan-rogu-block
```

#### show policer (policer-name)

```
user@host> show policer tcp-connection-policer
Filter: ingress-port-filter
Policers:
Name                               Packets
tcp-connection-policer            0
```

## show policy

List of Syntax	<a href="#">Syntax on page 4891</a> <a href="#">Syntax (EX Series Switches) on page 4891</a>
Syntax	<pre>show policy &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>policy-name</i>&gt;</pre>
Syntax (EX Series Switches)	<pre>show policy &lt;<i>policy-name</i>&gt;</pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display information about configured routing policies.
Options	<p><b>none</b>—List the names of all configured routing policies.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>policy-name</i></b>—(Optional) Show the contents of the specified policy.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <li>• <a href="#">show policy damping on page 3072</a></li> </ul>
List of Sample Output	<a href="#">show policy on page 4892</a> <a href="#">show policy <i>policy-name</i> on page 4892</a> <a href="#">show policy (Multicast Scoping) on page 4892</a>
Output Fields	<p><a href="#">Table 542 on page 4891</a> lists the output fields for the <b>show policy</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 542: show policy Output Fields**

Field Name	Field Description
<i>policy-name</i>	Name of the policy listed.
<i>term</i>	Policy term listed.
<i>from</i>	Match condition for the policy.
<i>then</i>	Action for the policy.

## Sample Output

### show policy

```
user@host> show policy
Configured policies:
__vrf-export-red-internal__
__vrf-import-red-internal__
red-export
all_routes
```

### show policy policy-name

```
user@host> show policy test-statics
Policy test-statics:
  from
    3.0.0.0/8  accept
    3.1.0.0/16  accept
  then reject
```

### show policy (Multicast Scoping)

```
user@host> show policy test-statics
Policy test-statics:
  from
    multicast-scoping == 8
```

## show policy conditions

<b>Syntax</b>	<pre>show policy conditions &lt;condition-name&gt; &lt;detail&gt; &lt;dynamic&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show policy conditions &lt;condition-name&gt; &lt;detail&gt; &lt;dynamic&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Display all the configured conditions as well as the routing tables with which the configuration manager is interacting. If the <b>detail</b> keyword is included, the output also displays dependent routes for each condition.</p>
<b>Options</b>	<p><b>none</b>—Display all configured conditions and associated routing tables.</p> <p><b>condition-name</b>—(Optional) Display information about the specified condition only.</p> <p><b>detail</b>—(Optional) Display the specified level of output.</p> <p><b>dynamic</b>—(Optional) Display information about the conditions in the dynamic database.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show policy conditions detail on page 4894</a>
<b>Output Fields</b>	<p><a href="#">Table 543 on page 4893</a> lists the output fields for the <b>show policy conditions</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 543: show policy conditions Output Fields**

Field Name	Field Description	Level of Output
<b>Condition</b>	Name of configured condition.	All levels
<b>event</b>	Condition type. If the <b>if-route-exists</b> option is configured, the event type is: <b>Existence of a route in a specific routing table.</b>	All levels
<b>Dependent routes</b>	List of routes dependent on the condition, along with the latest generation number.	<b>detail</b>
<b>Condition tables</b>	List of routing tables associated with the condition, along with the latest generation number and number of dependencies.	All levels

Table 543: show policy conditions Output Fields (*continued*)

Field Name	Field Description	Level of Output
If-route-exists conditions	List of conditions configured to look for a route in the specified table.	All levels

## Sample Output


### show policy conditions detail

```
user@host> show policy conditions detail
Configured conditions:
Condition primary (static), event: Existence of a route in a specific routing
table
Dependent routes:
  8.41.0.0/24, generation 18

Condition standby (static), event: Existence of a route in a specific routing
table
Dependent routes:
  8.41.0.0/24, generation 18

Condition tables:
Table mpls.0, generation 0, dependencies 0, If-route-exists conditions: primary
(static) standby (static)
Table l3vpn.inet.0, generation 633, dependencies 2
```

## test policy

<b>Syntax</b>	<code>test policy <i>policy-name</i> <i>prefix</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Test a policy configuration to determine which prefixes match routes in the routing table.
<div>  <p><b>NOTE:</b> If you are using the <code>test policy</code> command on a logical system, you must first set the CLI to the logical system context. For example, if you want to test a routing policy that is configured on logical system R2, first run the <code>set cli logical-system R2</code> command.</p> </div>	
<b>Options</b>	<p><i>policy-name</i>—Name of a policy.</p> <p><i>prefix</i>—Destination prefix to match.</p>
<b>Additional Information</b>	All prefixes in the default unicast routing table (inet.0) that match prefixes that are the same as or longer than the specific prefix are processed by the <b>from</b> clause in the specified policy. All prefixes accepted by the policy are displayed. The <b>test policy</b> command evaluates a policy differently from the BGP import process. When testing a policy that contains an <b>interface</b> match condition in the <b>from</b> clause, the <b>test policy</b> command uses the match condition. In contrast, BGP does not use the <b>interface</b> match condition when evaluating the policy against routes learned from internal BGP (IBGP) or external BGP (EGBP) multihop peers.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding Routing Policy Tests</i></li> <li>• <i>Example: Testing a Routing Policy with Complex Regular Expressions</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">test policy on page 4895</a>
<b>Output Fields</b>	For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.

## Sample Output

### test policy

```
user@host> test policy test-statics 3.0.0.1/8
inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
Prefixes passing policy:
```

```
3.0.0.0/8      *[BGP/170] 16:22:46, localpref 100, from 10.255.255.41
               AS Path: 50888 I
               > to 10.11.4.32 via en0.2, label-switched-path 12
3.3.3.1/32     *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
3.3.3.2/32     *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
3.3.3.3/32     *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
3.3.3.4/32     *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
               > to 10.0.4.7 via fxp0.0
Policy test-statics: 5 prefixes accepted, 0 prefixes rejected
```



# Troubleshooting Procedures

- [Troubleshooting Firewall Filters on page 4897](#)

## Troubleshooting Firewall Filters

---

Troubleshooting issues with firewall filters on EX Series switches:

1. [A Firewall Filter Configuration Returns a “No Space Available in TCAM” Message on page 4897](#)

### A Firewall Filter Configuration Returns a “No Space Available in TCAM” Message

**Problem**    **Description:** When a firewall filter configuration exceeds the amount of available ternary content addressable memory (TCAM) space, the switch returns the following system log (**syslogd**) message:

```
No space available in tcam.  
Rules for filter filter-name will not be installed.
```

The switch returns this error message during the commit operation in the following instances:

- If the firewall filter that you have applied to a port, VLAN, or Layer 3 interface requires more than the amount of available TCAM space.
- If you delete and add large firewall filters in the same commit operation. In this case, the large firewall filter might not be deleted from the TCAM space, because of which there will be no TCAM space freed up for the new firewall filter to be added to it. In addition to the **syslogd** message, the following error message is displayed in the CLI:

```
fpc<device-id> dfw_grph_merge_dfw_bind: rules for filter filter-name will not  
be installed
```

However, in both these instances, the commit operation for the firewall filter configuration is completed in the CLI.

**Solution**    When a firewall filter configuration exceeds the amount of available TCAM table space, you must configure a new firewall filter with fewer filter terms or, if you had deleted and created a firewall filter with a large number of terms (on the order of 1000 or more), you must delete and add the large firewall filters in separate commit operations.

The first procedure (set of steps) in this Solution section tells you how to delete a firewall filter and its bind point and associate a new firewall filter with that existing bind point.

The second procedure in this Solution section tells you how to create a new firewall filter with fewer terms (without deleting the bind point) and bind the new firewall filter with the existing bind point, when you want to create a firewall filter with fewer terms. Do not use the second procedure if you need to replace one large firewall filter with another large firewall filter—you must delete the original large firewall filter and commit that delete operation, and then add the new large firewall filter.

To delete the firewall filter and its bind point and apply a new firewall filter to the same bind point:

1. Delete the firewall filter configuration and its bind points to ports, VLANs, or Layer 3 interfaces—for example:

```
[edit]
user@switch# delete firewall family ethernet-switching filter mini-filter-ingress-vlan
user@switch# delete vlans voice-vlan description "filter to block rogue devices on voice-vlan"
user@switch# delete vlans voice-vlan filter input mini-filter-ingress-vlan
```

2. Commit the operation:

```
[edit]
user@switch# commit
```



**NOTE:** Use separate commit operations for deleting and adding large firewall filters.

3. Configure a firewall filter with fewer terms (if the error message appeared when you tried to create a new filter) or configure a large filter (if the error message appeared when you tried to delete and add large firewall filters)—for example:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-filter-ingress-vlan ...
```



**NOTE:** See [“Firewall Filters for EX Series Switches Overview”](#) on page 4696 to ascertain the maximum number of terms allowed for various firewall filters on EX Series switches.

4. Apply (bind) the new firewall filter to a port, VLAN, or Layer 3 interface—for example:

```
[edit]
user@switch# set vlans voice-vlan description "filter to block rogue devices on voice-vlan"
user@switch# set vlans voice-vlan filter input new-filter-ingress-vlan
```

5. Commit the operation:

```
[edit]
user@switch# commit
```

To create a new firewall filter and attach it to the existing bind point:

1. Configure a firewall filter with fewer terms than the original filter:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-filter-ingress-vlan...
```

2. Apply the firewall filter to the port, VLAN, or Layer 3 interfaces to overwrite the bind points of the original filter—for example:

```
[edit]
user@switch# set vlans voice-vlan description "smaller filter to block rogue devices on voice-vlan"
user@switch# set vlans voice-vlan filter input new-filter-ingress-vlan
```

As a bind point can be attached to only one firewall filter, this configuration detaches the bind point from the previous firewall filter that contained many terms and attaches the bind point to the new firewall filter.

3. Commit the operation:

```
[edit]
user@switch# commit
```

**Related  
Documentation**

- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches on page 4773](#)
- [Verifying That Firewall Filters Are Operational on page 4873](#)
- [Configuring Firewall Filters \(CLI Procedure\) on page 4804](#)
- [Configuring Firewall Filters \(J-Web Procedure\) on page 4813](#)



## PART 25

# Spanning-Tree Protocols

- [Overview on page 4903](#)
- [Configuration on page 4919](#)
- [Administration on page 5027](#)



## CHAPTER 81

# Overview

- [Spanning Trees Overview on page 4903](#)

### Spanning Trees Overview

---

- [Understanding MSTP for EX Series Switches on page 4904](#)
- [Understanding RSTP for EX Series Switches on page 4906](#)
- [Understanding STP for EX Series Switches on page 4910](#)
- [RSTP or VSTP Forced to Run as IEEE 802.1D STP on page 4912](#)
- [Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches on page 4914](#)
- [Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches on page 4916](#)
- [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches on page 4917](#)

## Understanding MSTP for EX Series Switches

Ethernet networks are susceptible to broadcast storms if loops are introduced. However, an Ethernet network should always include loops because they provide redundant paths in case of a link failure. Spanning-tree protocols address both of these issues because they provide link redundancy while simultaneously preventing undesirable loops.

Spanning-tree protocols intelligently avoid loops in a network by creating a tree topology (spanning-tree) of the entire bridged network with only one available path between the tree root and a leaf. All other paths are forced into a standby state. The tree *root* is a switch within the network elected by the STA (spanning-tree algorithm) to use when computing the best path between bridges throughout the network and the root bridge. Frames travel through the network to their destination—a *leaf* such as Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches an end-user PC—along branches. A tree *branch* is a network segment, or link, between bridges. Switches that forward frames through an STP spanning-tree are called *designated bridges*.

Juniper Networks EX Series Ethernet Switches provide Layer 2 loop prevention through Spanning-Tree Protocol (STP), Rapid Spanning-Tree Protocol (RSTP), Multiple Spanning-Tree Protocol (MSTP), and VLAN Spanning-Tree Protocol (VSTP). This topic explains MSTP.

This topic describes:

- [MSTP Maps Multiple VLANs on page 4904](#)
- [Configuring MSTP Regions on page 4905](#)
- [Selecting a Spanning-Tree Protocol on page 4905](#)

---

### MSTP Maps Multiple VLANs

MSTP is an extension of RSTP that maps multiple independent spanning-tree instances onto one physical topology. Each spanning-tree instance (STI) includes one or more VLANs. Unlike in STP and RSTP configurations, a port may belong to multiple VLANs and be dynamically blocked in one spanning-tree instance but forwarding in another. This behavior significantly improves network resource utilization by load-balancing across the network and maintaining switch CPU loads at moderate levels. MSTP also leverages the fast re-convergence time of RSTP when a network, switch, or port failure occurs within a spanning-tree instance.

MSTP creates a Common and Internal Spanning-Tree (CIST) to interconnect and manage all MSTP regions and even individual devices that run RSTP or STP, which are recognized as distinct spanning-tree regions by MSTP. The CIST views each MSTP region as a virtual bridge, regardless of the actual number of devices participating in the MSTP region, and enables MSTIs to link to other regions. The CIST is a single topology that connects all switches (STP, RSTP, and MSTP devices) through an active topology, ensuring connectivity between LANs and devices within a bridged network. This functionality provided by MSTP enables you to better utilize network resources while remaining backward-compatible with older network devices.



### Configuring MSTP Regions

When enabling MSTP, you define one or more MSTP regions. An MSTP region defines a logical domain where MSTIs can be administered independently of MSTIs in other regions, setting the boundary for Bridge Protocol Data Units (BPDUs) sent by one MSTI. An MSTP region is a group of switches that is defined by three parameters:

- Region name—User-defined alphanumeric name for the region.
- Revision level—User-defined value that identifies the region.
- Mapping table—Numerical digest of VLAN-to-instance mappings.

An MSTP region can support up to 64 MST instances, and each MSTI can support from 1 to 4094 VLANs. When you define a region, MSTP automatically creates an internal spanning-tree instance (IST instance 0) that provides the root switch for the region and includes all currently configured VLANs that are not specifically assigned to a user-defined Multiple Spanning-Tree Instance (MSTI). An MSTI includes all static VLANs that you specifically add to it. The switch places any dynamically created VLANs in the IST instance by default, unless you explicitly map them to another MSTI. Once you assign a VLAN to a user-defined MSTI, the switch removes the VLAN from the IST instance.

### Selecting a Spanning-Tree Protocol

The default factory configuration for EX Series switches is RSTP, a faster version of STP. To determine which spanning-tree protocol is best for your situation, see [Table 544 on page 4905](#) below.

**Table 544: Selecting a Spanning-Tree Protocol**

Protocol	Advantages	Disadvantages
RSTP	<ul style="list-style-type: none"> <li>• Rapid Spanning-Tree Protocol is the default switch configuration and is recommended for most network configurations because it converges more quickly than STP after a failure.</li> <li>• Voice and video work better with RSTP than they do with STP.</li> <li>• RSTP is backward compatible with STP so switches do not all have to run RSTP.</li> </ul>	<ul style="list-style-type: none"> <li>• RSTP does not work with 802.1D 1998 bridges.</li> <li>• RSTP is not recommended for multiple VLAN networks because it is not VLAN-aware—as a result, all VLANs within a LAN share the same spanning-tree. This limits the number of forwarding paths for data traffic.</li> </ul>
STP	<ul style="list-style-type: none"> <li>• Spanning-Tree Protocol works with 802.1D 1998 bridges.</li> <li>• RSTP is backward compatible with STP so switches do not all have to run STP.</li> </ul>	<ul style="list-style-type: none"> <li>• STP is slower than RSTP.</li> <li>• STP is not recommended for multiple VLAN networks because it is not VLAN-aware—as a result, all VLANs within a LAN share the same spanning-tree. This limits the number of forwarding paths for data traffic.</li> </ul>
MSTP	<ul style="list-style-type: none"> <li>• Multiple Spanning-Tree Protocol works with most VLANs.</li> <li>• RSTP and STP are recognized as distinct spanning-tree regions by MSTP.</li> </ul>	<ul style="list-style-type: none"> <li>• Some protocols require compatibility that is not provided by MSTP. In this case, use VSTP.</li> <li>• MSTP uses more CPU than RSTP and does not converge as fast as RSTP.</li> </ul>

Table 544: Selecting a Spanning-Tree Protocol (*continued*)

Protocol	Advantages	Disadvantages
VSTP	<ul style="list-style-type: none"> <li>VLAN Spanning-Tree Protocol works with VLANs that require device compatibility.</li> <li>VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on a switch.</li> </ul>	<ul style="list-style-type: none"> <li>With VSTP there can be only STP instance per VLAN, whereas MSTP lets you combine multiple VLANs in one instance.</li> <li>VSTP supports a limited number of ports compared to RSTP.</li> <li>VSTP uses more CPU than RSTP and does not converge as fast as RSTP.</li> <li>Having a large number of VSTP and RSTP instances can cause continuous changes in the topology. Ensure to check the scale limits before configuring large number of VSTP instances.</li> </ul>

- Related Documentation**
- [Understanding STP for EX Series Switches on page 4910](#)
  - [Understanding RSTP for EX Series Switches on page 4906](#)
  - [Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912](#)
  - [Understanding Layer 2 Protocol Tunneling on EX Series Switches](#)
  - [Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936](#)

## Understanding RSTP for EX Series Switches

Ethernet networks are susceptible to broadcast storms if loops are introduced. However, an Ethernet network should always include loops because they provide redundant paths in case of a link failure. Spanning tree protocols address both of these issues because they provide link redundancy while simultaneously preventing undesirable loops. Rapid Spanning-Tree Protocol (RSTP) is the default spanning-tree protocol for preventing loops on Ethernet networks.

This topic describes:

- [Spanning-Tree Protocols Help Prevent Broadcast Storms on page 4906](#)
- [RSTP is an Enhancement of the Original STP on page 4907](#)
- [Port Roles Determine Participation in The Spanning-Tree on page 4907](#)
- [Port States Determine How a Port Processes a Frame on page 4907](#)
- [Edge Ports Connect to Devices That Cannot Be Part of a Spanning-Tree on page 4908](#)
- [BPDUs Maintain the Spanning-Tree on page 4908](#)
- [When an RSTP Root Bridge Fails on page 4909](#)
- [Switches Must Relearn MAC Addresses After a Link Failure on page 4909](#)
- [Selecting a Spanning-Tree Protocol on page 4909](#)

### Spanning-Tree Protocols Help Prevent Broadcast Storms

Spanning tree protocols intelligently avoid loops in a network by creating a tree topology (spanning-tree) of the entire bridged network with only one available path between the

tree root and a leaf. All other paths are forced into a standby state. The tree *root* is a switch within the network elected by the STA (spanning-tree algorithm) to use when computing the best path between bridges throughout the network and the root bridge. Frames travel through the network to their destination—a *leaf* such as an end-user PC—along branches. A tree *branch* is a network segment, or link, between bridges. Switches that forward frames through an STP spanning-tree are called *designated bridges*.

Juniper Networks EX Series Ethernet Switches provide Layer 2 loop prevention through Spanning-Tree Protocol (STP), Rapid Spanning-Tree Protocol (RSTP), Multiple Spanning-Tree Protocol (MSTP), and VLAN Spanning-Tree Protocol (VSTP). This topic explains the spanning-tree default RSTP.

### RSTP is an Enhancement of the Original STP

RSTP evolved from the original STP IEEE 802.1D protocol to provide faster spanning-tree re-convergence after a switch port, switch, or LAN failure. Where STP took up to 50 seconds to respond to topology changes, RSTP responds to changes within the timeframe of three hello BPDUs (bridge protocol data units), or 6 seconds. This is the primary reason that RSTP is the default configuration on EX Series switches. In addition, note that EX Series switches configured to use STP actually run RSTP force version 0, which is compatible with STP.

### Port Roles Determine Participation in The Spanning-Tree

Each port has both a state and a role. A port's *role* determines how it participates in the spanning-tree. The five port roles used in RSTP are:

- Root port—The port closest to the root bridge (has the lowest path cost from a bridge). This is the only port that receives frames from and forwards frames to the root bridge.
- Designated port—The port that forwards traffic away from the root bridge toward a leaf. A designated bridge has one designated port for every link connection it serves. A root bridge forwards frames from all of its ports, which serve as designated ports.
- Alternate port—A port that provides an alternate path toward the root bridge if the root port fails and is placed in the discarding state. This port is not part of the active spanning-tree, but if the root port fails, the alternate port immediately takes over.
- Backup port—A port that provides a backup path toward the leaves of the spanning-tree if a designated port fails and is placed in the discarding state. A backup port can only exist where two or more bridge ports connect to the same LAN for which the bridge serves as the designated bridge. A backup port for a designated port immediately takes over if the port fails.
- Disabled port—The port is not part of the active spanning-tree.

### Port States Determine How a Port Processes a Frame

Each port has both a state and a role. A port's *state* determines how it processes a frame. RSTP places each port of a designated bridge in one of three states:

- Discarding—The port discards all BPDUs. A port in this state discards all frames it receives and does not learn MAC addresses.

- Learning—The port prepares to forward traffic by examining received frames for location information in order to build its MAC address table.
- Forwarding—The port filters and forwards frames. A port in the forwarding state is part of the active spanning-tree.

### Edge Ports Connect to Devices That Cannot Be Part of a Spanning-Tree

---

RSTP also defines the concept of an *edge port*, which is a designated port that connects to non-STP-capable devices, such as PCs, servers, routers, or hubs that are not connected to other switches. Because edge ports connect directly to end stations, they cannot create network loops and can transition to the forwarding state immediately. You can manually configure edge ports, and a switch can also detect edge ports by noting the absence of communication from the end stations.

The edge ports themselves do send BPDUs to the spanning-tree. If you have a good understanding of the implications on your network and want to modify RSTP on the edge port interface, see *Configuring RSTP (CLI Procedure)*.

### BPDUs Maintain the Spanning-Tree

---

Spanning-tree protocols use frames called bridge protocol data units (BPDUs) to create and maintain the spanning-tree. A BPDU frame is a message sent from one switch to another to communicate information about itself, such as its bridge ID, root path costs, and port MAC addresses. The initial exchange of BPDUs between switches determines the root bridge. Simultaneously, BPDUs are used to communicate the cost of each link between branch devices, which is based upon port speed or user configuration. RSTP uses this path cost to determine the ideal route for data frames to travel from one leaf to another leaf and then blocks all other routes. If an edge port receives a BPDU, it automatically transitions to a regular RSTP port.

When the network is in a steady state, the spanning-tree converges when the spanning-tree algorithm (STA) identifies both the root and designated bridges and all ports are in either a forwarding or blocking state. To maintain the tree, the root bridge continues to send BPDUs at a “hello time” interval (default 2 seconds). These BPDUs continue to communicate the current tree topology. When a port receives a hello BPDU, it compares the information to that already stored for the receiving port. One of three actions takes place when a switch receives a BPDU:

- If the BPDU data matches the existing entry in the MAC address table, the port resets a timer called “max age” to zero and then forwards a new BPDU with the current active topology information to the next port in the spanning-tree.
- If the topology in the BPDU has been changed, the information is updated in the MAC address table, “max age” is again set to zero, and a new BPDU is forwarded with the current active topology information to the next port in the spanning-tree.
- When an RSTP port does not receive a BPDU for three hello times, it reacts one of two ways. If the port is the root port, a complete rework of the spanning-tree occurs—see *When an RSTP Root Bridge Fails*. If the bridge is any non-root bridge, RSTP detects that the connected device cannot send BPDUs and converts that port to an edge port.

### When an RSTP Root Bridge Fails

When a link to the root port goes down, a flag called a topology change notification (TCN) is added to the BPDU. When this BPDU reaches the next port in the VLAN, the MAC address table is flushed and the BPDU is sent to the next bridge. Eventually, all ports in the VLAN have flushed their MAC address tables. Then, RSTP configures a new root port.

After a root port or a designated port fails, the alternate or backup port takes over after an exchange of BPDUs called the proposal-agreement handshake. RSTP propagates this handshake over *point-to-point links*, which are dedicated links between two network nodes, or switches, that connect one port to another. If a local port becomes a new root or designated port, it negotiates a rapid transition with the receiving port on the nearest neighboring switch by using the proposal-agreement handshake to ensure a loop-free topology.

### Switches Must Relearn MAC Addresses After a Link Failure

Because a link failure causes all associated ports to flush their MAC address table, the network may be slower as it floods to relearn the MAC addresses. There is a way to speed up this relearning process. During TCN propagation, the Layer 2 forwarding table of switches is flushed, resulting in a flood of data packets. The ARP feature causes the switch to proactively send ARP requests for IP addresses in the ARP cache (present because of Layer 3 VLAN interface). With ARP on STP enabled, as the reply comes through, the switches build up the Layer 2 forwarding table, thus limiting the flooding later. Enabling ARP on STP is most useful to prevent excessive flooding in large Layer 2 networks using RVIs.

### Selecting a Spanning-Tree Protocol

The default factory configuration for EX Series switches is RSTP, a faster version of the original STP. To determine which spanning-tree protocol is best for your situation, see [Table 545 on page 4909](#) below.

**Table 545: Selecting a Spanning-Tree Protocol**

Protocol	Advantages	Disadvantages
RSTP	<ul style="list-style-type: none"> <li>Rapid Spanning-Tree Protocol is the default switch configuration and is recommended for most network configurations because it converges more quickly than STP after a failure.</li> <li>Voice and video work better with RSTP than they do with STP.</li> <li>RSTP is backward compatible with STP so switches do not all have to run RSTP.</li> <li>RSTP supports more ports than MSTP or VSTP</li> </ul>	<ul style="list-style-type: none"> <li>RSTP does not work with 802.1D 1998 bridges.</li> <li>RSTP is not recommended for multiple VLAN networks because it is not VLAN-aware—as a result, all VLANs within a LAN share the same spanning-tree. This limits the number of forwarding paths for data traffic.</li> </ul>
STP	<ul style="list-style-type: none"> <li>Spanning-Tree Protocol works with 802.1D 1998 bridges.</li> <li>RSTP is backward compatible with STP so switches do not all have to run STP.</li> </ul>	<ul style="list-style-type: none"> <li>STP is slower than RSTP.</li> <li>STP is not recommended for multiple VLAN networks because it is not VLAN-aware—as a result, all VLANs within a LAN share the same spanning-tree. This limits the number of forwarding paths for data traffic.</li> </ul>

Table 545: Selecting a Spanning-Tree Protocol (*continued*)

Protocol	Advantages	Disadvantages
MSTP	<ul style="list-style-type: none"> <li>Multiple Spanning-Tree Protocol works with most VLANs.</li> <li>RSTP and STP are recognized as distinct spanning-tree regions by MSTP.</li> </ul>	<ul style="list-style-type: none"> <li>Some protocols require compatibility that is not provided by MSTP. In this case, use VSTP.</li> <li>MSTP supports a limited number of ports.</li> <li>MSTP uses more CPU than RSTP and does not converge as fast as RSTP.</li> </ul>
VSTP	<ul style="list-style-type: none"> <li>VLAN Spanning-Tree Protocol works with VLANs that require device compatibility.</li> <li>VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on a switch.</li> </ul>	<ul style="list-style-type: none"> <li>With VSTP there can be only STP instance per VLAN, whereas MSTP lets you combine multiple VLANs in one instance.</li> <li>VSTP supports a limited number of ports compared to RSTP.</li> <li>VSTP uses more CPU than RSTP and does not converge as fast as RSTP.</li> <li>Having a large number of VSTP and RSTP instances can cause continuous changes in the topology. As a workaround, reduce the number of VSTP instances to fewer than 190.</li> </ul>

**Related Documentation**

- [Understanding STP for EX Series Switches on page 4910](#)
- [Understanding MSTP for EX Series Switches on page 4904](#)
- [Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912](#)
- [Understanding Layer 2 Protocol Tunneling on EX Series Switches](#)
- [Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919](#)

## Understanding STP for EX Series Switches

Ethernet networks are susceptible to broadcast storms if loops are introduced. However, an Ethernet network should always include loops because they provide redundant paths in case of a link failure. Spanning-tree protocols address both of these issues because they provide link redundancy while simultaneously preventing undesirable loops.

Juniper Networks EX Series Ethernet Switches provide Layer 2 loop prevention through Spanning-Tree Protocol (STP), Rapid Spanning-Tree Protocol (RSTP), Multiple Spanning-Tree Protocol (MSTP), and VLAN Spanning-Tree Protocol (VSTP). Configure STP when you need to support older 802.1D 1998 bridges. However, note that EX Series switches configured to use STP actually run RSTP force version 0, which is compatible with STP. For an explanation of RSTP, see [“Understanding RSTP for EX Series Switches” on page 4906](#)

This topic describes:

- [Selecting a Spanning-Tree Protocol on page 4911](#)

### Selecting a Spanning-Tree Protocol

The default factory configuration for EX Series switches is RSTP, a faster version of STP. To determine which spanning-tree protocol is best for your situation, see [Table 546 on page 4911](#) below.

**Table 546: Selecting a Spanning-Tree Protocol**

Protocol	Advantages	Disadvantages
RSTP	<ul style="list-style-type: none"> <li>Rapid Spanning-Tree Protocol is the default switch configuration and is recommended for most network configurations because it converges more quickly than STP after a failure.</li> <li>Voice and video work better with RSTP than they do with STP.</li> <li>RSTP is backward compatible with STP so switches do not all have to run RSTP.</li> <li>RSTP supports more ports than MSTP or VSTP</li> </ul>	<ul style="list-style-type: none"> <li>RSTP does not work with 802.1D 1998 bridges.</li> <li>RSTP is not recommended for multiple VLAN networks because it is not VLAN-aware—as a result, all VLANs within a LAN share the same spanning-tree. This limits the number of forwarding paths for data traffic.</li> </ul>
STP	<ul style="list-style-type: none"> <li>Spanning-Tree Protocol works with 802.1D 1998 bridges.</li> <li>RSTP is backward compatible with STP so switches do not all have to run STP.</li> </ul>	<ul style="list-style-type: none"> <li>STP is slower than RSTP.</li> <li>STP is not recommended for multiple VLAN networks because it is not VLAN-aware—as a result, all VLANs within a LAN share the same spanning-tree. This limits the number of forwarding paths for data traffic.</li> </ul>
MSTP	<ul style="list-style-type: none"> <li>Multiple Spanning-Tree Protocol works with most VLANs.</li> <li>RSTP and STP are recognized as distinct spanning-tree regions by MSTP.</li> </ul>	<ul style="list-style-type: none"> <li>Some protocols require compatibility that is not provided by MSTP. In this case, use VSTP.</li> <li>MSTP supports a limited number of ports.</li> <li>MSTP uses more CPU than RSTP and does not converge as fast as RSTP.</li> </ul>
VSTP	<ul style="list-style-type: none"> <li>VLAN Spanning-Tree Protocol works with VLANs that require device compatibility.</li> <li>VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on a switch.</li> </ul>	<ul style="list-style-type: none"> <li>With VSTP there can be only STP instance per VLAN, whereas MSTP lets you combine multiple VLANs in one instance.</li> <li>VSTP supports a limited number of ports compared to RSTP.</li> <li>VSTP uses more CPU than RSTP and does not converge as fast as RSTP.</li> <li>Having a large number of VSTP and RSTP instances can cause continuous changes in the topology. As a workaround, reduce the number of VSTP instances to fewer than 190.</li> </ul>

- Related Documentation**
- [Configuring STP \(CLI Procedure\)](#)
  - [Understanding RSTP for EX Series Switches on page 4906](#)
  - [Understanding MSTP for EX Series Switches on page 4904](#)
  - [Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912](#)
  - [Understanding Layer 2 Protocol Tunneling on EX Series Switches](#)

## RSTP or VSTP Forced to Run as IEEE 802.1D STP

On MX Series routers and EX Series switches in a Layer 2 environment, you can force the configured Rapid Spanning Tree Protocol (RSTP) or VLAN Spanning Tree Protocol (VSTP) to run as the original IEEE 802.1D Spanning Tree Protocol (STP) version. Configure original IEEE\_802.1D STP for compatibility with older bridges that do not support RSTP or VSTP.

Keep the following limitations in mind when RSTP or VSTP are forced to run as the original STP version:

- If you configure an instance interface as an edge port, the configuration statement is ignored.
- If you configure point-to-point link mode for an instance interface, the configuration statement is ignored.

### Related Documentation

- *Spanning-Tree Protocols Supported*
- *Configuring Rapid Spanning-Tree Protocol*
- *Configuring VLAN Spanning-Tree Protocol*
- [Reverting to RSTP or VSTP from Forced IEEE 802.1D STP on page 4985](#)
- [Forcing RSTP or VSTP to Run as IEEE 802.1D STP \(CLI Procedure\) on page 4976](#)
- [force-version on page 4999](#)

## Understanding VSTP for EX Series Switches and QFX Series Switches

Juniper Networks EX Series Ethernet Switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP). The default factory configuration for EX Series switches uses RSTP. If you use VLANs, however, we recommend that you enable MSTP unless your network requires the device compatibility provided by VSTP. Switches configured to run VSTP automatically assign each VLAN to one spanning-tree instance that runs RSTP. While this approach is useful to optimize network usage in small networks with a limited number of VLANs, a VSTP configuration in a network with several hundred VLANs can overload switch CPUs. MSTP ensures that your network is not slowed down by the increased network traffic caused by hundreds of VLANs, each with its own spanning-tree instance.

When using VSTP, you can selectively configure up to 253 VLANs per switch—additional VLANs use RSTP. (VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on a switch.)



**NOTE:** When you configure VSTP, we recommend that you enable VSTP on all VLANs that can receive VSTP bridge protocol data units (BPDUs).

---





**NOTE:** When you configure VSTP with the `set protocol vstp vlan all` command, VLAN ID 1 is not set; it is excluded so that the configuration is compatible with Cisco PVST+. If you want VLAN ID 1 to be included in the VSTP configuration on your switch, you must set it separately with the `set protocol vstp vlan 1` command.



**NOTE:** Option `vlan all` is not supported in Junos OS Release 13.2X50.

### Selecting a Spanning-Tree Protocol

The default factory configuration for EX Series switches is RSTP, a faster version of STP. To determine which spanning-tree protocol is best for your situation, see [Table 547 on page 4913](#).

**Table 547: Selecting a Spanning-Tree Protocol**

Protocol	Advantages	Disadvantages
RSTP	<ul style="list-style-type: none"> <li>Rapid Spanning Tree Protocol is the default switch configuration and is recommended for most network configurations because it converges more quickly than STP after a failure.</li> <li>Voice and video work better with RSTP than they do with STP.</li> <li>RSTP is backward compatible with STP so switches do not all have to run RSTP.</li> <li>RSTP supports more ports than MSTP or VSTP.</li> </ul>	<ul style="list-style-type: none"> <li>RSTP does not work with 802.1D 1998 bridges.</li> <li>RSTP is not recommended for multiple VLAN networks because it is not VLAN-aware—as a result, all VLANs within a LAN share the same spanning tree. This limits the number of forwarding paths for data traffic.</li> </ul>
STP	<ul style="list-style-type: none"> <li>Spanning Tree Protocol works with 802.1D 1998 bridges.</li> <li>RSTP is backward compatible with STP so switches do not all have to run STP.</li> </ul>	<ul style="list-style-type: none"> <li>STP is slower than RSTP.</li> <li>STP is not recommended for multiple VLAN networks because it is not VLAN-aware—as a result, all VLANs within a LAN share the same spanning-tree. This limits the number of forwarding paths for data traffic.</li> </ul>
MSTP	<ul style="list-style-type: none"> <li>Multiple Spanning Tree Protocol works with most VLANs.</li> <li>RSTP and STP are recognized as distinct Spanning Tree regions by MSTP.</li> </ul>	<ul style="list-style-type: none"> <li>Some protocols require compatibility that is not provided by MSTP. In this case, use VSTP.</li> <li>MSTP supports a limited number of ports.</li> <li>MSTP uses more CPU than RSTP and does not converge as fast as RSTP.</li> </ul>
VSTP	<ul style="list-style-type: none"> <li>VLAN Spanning Tree Protocol works with VLANs that require device compatibility.</li> <li>VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on a switch.</li> </ul>	<ul style="list-style-type: none"> <li>With VSTP there can be only STP instance per VLAN, whereas MSTP lets you combine multiple VLANs in one instance.</li> <li>VSTP supports a limited number of ports compared to RSTP.</li> <li>VSTP uses more CPU than RSTP and does not converge as fast as RSTP.</li> </ul>

- Related Documentation**
- [Understanding STP for EX Series Switches on page 4910](#)
  - [Understanding RSTP for EX Series Switches on page 4906](#)
  - [Understanding MSTP for EX Series Switches on page 4904](#)
  - [Understanding Layer 2 Protocol Tunneling on EX Series Switches](#)
  - [Configuring VSTP \(CLI Procedure\) on page 4982](#)

## Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches

Networks frequently use multiple protocols simultaneously to achieve different goals and in some cases those protocols might conflict with each other. One such case is when spanning-tree protocols are active on the network, where a special type of switching frame called a bridge protocol data unit (BPDU) can conflict with BPDUs generated on other devices such as PCs. The different kinds of BPDUs are not compatible but they can still be recognized by other devices that use BPDUs and cause network outages. You need to protect any device that recognizes BPDUs from picking up incompatible BPDUs.

- [Different Kinds of BPDUs on page 4914](#)
- [Protecting Switches From Incompatible BPDUs on page 4914](#)

### Different Kinds of BPDUs

---

Spanning-tree protocols such as Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP) generate their own BPDUs. These peer STP applications use their BPDUs to communicate, and ultimately, the exchange of BPDUs determines which interfaces block traffic and which interfaces become root ports and forward traffic.

User bridge applications running on a PC can also generate BPDUs. If these BPDUs are picked up by STP applications running on the switch, they can trigger STP miscalculations, and those miscalculations can lead to network outages. Similarly, BPDUs generated by STP protocols can cause problems if they are picked up by devices like PCs that are not using STP. Some mechanism for BPDU protection must be implemented in these cases.

### Protecting Switches From Incompatible BPDUs

---

To protect the state of spanning-tree protocols on switches from outside BPDUs, enable BPDU protection on the interfaces of a switch on which spanning-tree protocols are configured and are connected to user devices (such as PCs)—for example, on edge ports connected to PCs. Use the same strategy when a device on which STP is not configured is connected to a switch through a trunk interface that could be forwarding BPDUs generated by spanning-tree protocols. In this case, you would protect the device from BPDUs generated by the STP on the switch.

To configure BPDU protection on a switch on which spanning-tree protocols are configured, include the **bpdu-block-on-edge** statement at the `[edit protocols (stp| mstp|rstp) ]` hierarchy level. To prevent such a switch from forwarding BPDUs generated by spanning-tree protocols to devices, include the **bpdu-block** statement at the `[edit ethernet-switching-options]` hierarchy level.



**NOTE:** You can configure the `drop` statement under the `bpdudrop` statement only on interfaces on which no spanning-tree protocol is configured.

When an interface configured with BPDU protection encounters an incompatible BPDU, it drops that BPDU and then, either shuts down or continues to receive packets other than spanning-tree protocol BPDUs depending on the configuration defined in the `bpdudrop` statement. If the interface continues to be open after dropping all incompatible BPDUs, all packets except incompatible BPDUs continue to ingress and egress through the interface.

If the interface shuts down after dropping all BPDUs, there are two ways to re-enable the interface:

- Include the `disable-timeout (Spanning Trees)` statement in the BPDU configuration to enable the interface(s) to automatically return to service when the specified timer expires.
- Issue the operational mode command `clear ethernet-switching bpdudrop-error` on the switch.



**NOTE:** You can also configure BPDU drop protection on a specified interface or any interface of an access switch if the interfaces or interfaces do *not* have a spanning tree protocol configured. See `drop` for additional information about that option.

#### Related Documentation

- [Configuring BPDU Protection on an Interface \(CLI Procedure\)](#)
- [Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4956](#)
- [Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4961](#)
- [Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches on page 4916](#)
- [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches on page 4917](#)
- [Understanding STP for EX Series Switches on page 4910](#)

## Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches

Juniper Networks EX Series Ethernet Switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). Loop protection increases the efficiency of STP, RSTP, and MSTP by preventing ports from moving into a forwarding state that would result in a loop opening up in the network.

A loop-free network in spanning-tree topologies is supported through the exchange of a special type of frame called bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic (preventing loops) and which interfaces become root ports and forward traffic.

However, a blocking interface can transition to the forwarding state in error if the interface stops receiving BPDUs from its designated port on the segment. Such a transition error can occur when there is a hardware error on the switch or software configuration error between the switch and its neighbor.

When loop protection is enabled, the spanning-tree topology detects root ports and blocked ports and makes sure both keep receiving BPDUs. If a loop-protection-enabled interface stops receiving BPDUs from its designated port, it reacts as it would react to a problem with the physical connection on this interface. It doesn't transition the interface to a forwarding state, but instead transitions it to a loop-inconsistent state. The interface recovers and then it transitions back to the spanning-tree blocking state as soon as it receives a BPDU.

We recommend that you enable loop protection on all switch interfaces that have a chance of becoming root or designated ports. Loop protection is most effective when enabled in the entire switched network. When you enable loop protection, you must configure at least one action (**log**, **block**, or both).

Note that an interface can be configured for either loop protection or root protection, but not for both.

### Related Documentation

- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches on page 4967](#)
- [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches on page 4917](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches on page 4914](#)
- [Understanding MSTP for EX Series Switches on page 4904](#)
- [Understanding RSTP for EX Series Switches on page 4906](#)
- [Understanding STP for EX Series Switches on page 4910](#)
- [Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912](#)

## Understanding Root Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches

Juniper Networks EX Series Ethernet Switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). A loop-free network is supported through the exchange of a special type of frame called bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic and which interfaces become root ports and forward traffic.

However, a root port elected through this process has the possibility of being wrongly elected. A user bridge application running on a PC can generate BPDUs, too, and interfere with root port election. Root protection allows network administrators to manually enforce the root bridge placement in the network.

Enable root protection on interfaces that should not receive superior BPDUs from the root bridge and should not be elected as the root port. These interfaces become designated ports and are typically located on an administrative boundary. If the bridge receives superior STP BPDUs on a port that has root protection enabled, that port transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge that should not be the root bridge from being elected the root bridge. After the bridge stops receiving superior STP BPDUs on the interface with root protection, the interface returns to a listening state, followed by a learning state, and ultimately back to a forwarding state. Recovery back to the forwarding state is automatic.

When root protection is enabled on an interface, it is enabled for all the STP instances on that interface. The interface is blocked only for instances for which it receives superior BPDUs. Otherwise, it participates in the spanning-tree topology.

An interface can be configured for either root protection or loop protection, but not for both.

### Related Documentation

- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on EX Series Switches on page 4971](#)
- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches on page 4967](#)
- [Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4956](#)
- [Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4961](#)
- [Understanding MSTP for EX Series Switches on page 4904](#)
- [Understanding RSTP for EX Series Switches on page 4906](#)
- [Understanding STP for EX Series Switches on page 4910](#)
- [Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912](#)



## CHAPTER 82

# Configuration

- [Configuration Examples on page 4919](#)
- [Configuration Tasks on page 4976](#)
- [Operational Tasks on page 4984](#)
- [Configuration Statements on page 4985](#)

### Configuration Examples

---

- [Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919](#)
- [Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936](#)
- [Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4956](#)
- [Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4961](#)
- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches on page 4967](#)
- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on EX Series Switches on page 4971](#)

### Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches

EX Series switches use Rapid Spanning Tree Protocol (RSTP) by default to provide a loop-free topology.

When switches that support redundant Routing Engines use RSTP, it is important to keep RSTP synchronized on both Routing Engines so that no loss of service occurs after a Routing Engine switchover. Nonstop bridging protocol keeps Routing Engines synchronized.

This example describes how to configure RSTP and NSB on four EX Series switches:

- [Requirements on page 4920](#)
- [Overview and Topology on page 4920](#)
- [Configuring RSTP and Nonstop Bridging on Switch 1 on page 4922](#)
- [Configuring RSTP and Nonstop Bridging on Switch 2 on page 4925](#)

- [Configuring RSTP and Nonstop Bridging on Switch 3 on page 4928](#)
- [Configuring RSTP and Nonstop Bridging on Switch 4 on page 4931](#)
- [Verification on page 4934](#)

## Requirements

---

This example uses the following hardware and software components:

- Junos OS Release 11.3 or later for EX Series switches
- Four EX Series switches

Before you configure the switches for RSTP, be sure you have:

- Installed the four switches. See *Connecting and Configuring an EX Series Switch (J-Web Procedure)*.
- Performed the initial software configuration on all switches. See *Installing and Connecting an EX3200 Switch*.

## Overview and Topology

---

RSTP works by identifying certain links as point to point links and blocking other possible paths. When one of the point-to-point links fails, a designated alternate link transitions to the forwarding state and take over. Configuring nonstop bridging (NSB) on a switch with redundant Routing Engines keeps RSTP synchronized on both Routing Engines. This way, RSTP remains active immediately after a switchover because it is already synchronized to the backup Routing Engine. RSTP does not have to reconverge after a Routing Engine switchover when NSB is enabled because the neighbor devices do not detect an RSTP change on the switch. In this example, four EX Series switches are connected in the topology displayed in [Figure 73 on page 4921](#) to create a loop-free topology with NSB applied to switches with dual Routing Engines.



Figure 73: Network Topology for RSTP




Table 548 on page 4921 shows the components of the topology for this example.



**NOTE:** You can configure RSTP on logical or physical interfaces. This example shows RSTP configured on logical interfaces.

Table 548: Components of the Topology for Configuring RSTP

Property	Settings
Switch 1	<p>The following interfaces on Switch 1 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>ge-0/0/9</b> is connected to Switch 2</li> <li>• <b>ge-0/0/13</b> is connected to Switch 4</li> <li>• <b>ge-0/0/11</b> is connected to Switch 3</li> </ul>
Switch 2	<p>The following interfaces on Switch 2 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>ge-0/0/14</b> is connected to Switch 1</li> <li>• <b>ge-0/0/18</b> is connected to Switch 3</li> </ul>
Switch 3	<p>The following interfaces on Switch 3 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>ge-0/0/26</b> is connected to Switch 1</li> <li>• <b>ge-0/0/28</b> is connected to Switch 2</li> <li>• <b>ge-0/0/24</b> is connected to Switch 4</li> </ul>

Table 548: Components of the Topology for Configuring RSTP (*continued*)

Property	Settings
Switch 4	<p>The following interfaces on Switch 4 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>ge-0/0/19</b> is connected to Switch 1</li> <li>• <b>ge-0/0/23</b> is connected to Switch 3</li> </ul>
VLAN names and tag IDs	<p><b>voice-vlan</b>, tag 10</p> <p><b>employee-vlan</b>, tag 20</p> <p><b>guest-vlan</b>, tag 30</p> <p><b>camera-vlan</b>, tag 40</p>

This configuration example creates a loop-free topology between four EX Series switches using RSTP.

An RSTP topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.
- The *backup port* is a backup port for the designated port. When a designated port goes down, the backup port becomes the active designated port and starts forwarding data.



**NOTE:** You also can create a loop-free topology between the aggregation layer and the distribution layer using redundant trunk links. For more information about configuring redundant trunk links, see *Example: Configuring Redundant Trunk Links for Faster Recovery*.

### Configuring RSTP and Nonstop Bridging on Switch 1

#### CLI Quick Configuration

To quickly configure RSTP and nonstop bridging on Switch 1, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/13 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/9 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
```

```

set protocols rstp bridge-priority 16k
set protocols rstp interface ge-0/0/13.0 cost 1000
set protocols rstp interface ge-0/0/13.0 mode point-to-point
set protocols rstp interface ge-0/0/9.0 cost 1000
set protocols rstp interface ge-0/0/9.0 mode point-to-point
set protocols rstp interface ge-0/0/11.0 cost 1000
set protocols rstp interface ge-0/0/11.0 mode point-to-point

```

If Switch 1 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 1, copy the following commands and paste them into the switch terminal window:

```

set chassis redundancy graceful switchover
set system commit synchronize
set ethernet-switching-options nonstop-bridging

```

### Step-by-Step Procedure

To configure RSTP and nonstop bridging on Switch 1:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```

[edit vlans]
user@switch1# set voice-vlan description "Voice VLAN"
user@switch1# set voice-vlan vlan-id 10
user@switch1# set employee-vlan description "Employee VLAN"
user@switch1# set employee-vlan vlan-id 20
user@switch1# set guest-vlan description "Guest VLAN"
user@switch1# set guest-vlan vlan-id 30
user@switch1# set camera-vlan description "Camera VLAN"
user@switch1# set camera-vlan vlan-id 40

```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```

[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]

```

3. Configure the port mode for the interfaces:

```

[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching port-mode trunk
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching port-mode trunk
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching port-mode trunk

```

4. Configure RSTP on the switch:

```

[edit protocols]
user@switch1# rstp bridge-priority 16k
user@switch1# rstp interface ge-0/0/13.0 cost 1000
user@switch1# rstp interface ge-0/0/13.0 mode point-to-point
user@switch1# rstp interface ge-0/0/9.0 cost 1000
user@switch1# rstp interface ge-0/0/9.0 mode point-to-point
user@switch1# rstp interface ge-0/0/11.0 cost 1000
user@switch1# rstp interface ge-0/0/11.0 mode point-to-point

```

### Step-by-Step Procedure

If Switch 1 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 1:

1. Enable graceful Routing Engine switchover (GRES):

- ```
[edit chassis redundancy]
user@switch1# set graceful-switchover
```
2. Configure the switch to always synchronize configuration changes between the Routing Engines:
- ```
[edit system]
user@switch1# set commit synchronize
```
- If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.
3. Enable nonstop bridging:
- ```
[edit ethernet-switching-options]
user@switch1# set nonstop-bridging
```



**NOTE:** This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

**Results** Check the results of the configuration:

```
user@switch1> show configuration
interfaces {
  ge-0/0/13 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  ge-0/0/9 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  ge-0/0/11 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
}
```

```

}
protocols {
  rstp {
    bridge-priority 16k;
    interface ge-0/0/13.0 {
      cost 1000;
      mode point-to-point;
    }
    interface ge-0/0/9.0 {
      cost 1000;
      mode point-to-point;
    }
    interface ge-0/0/11.0 {
      cost 1000;
      mode point-to-point;
    }
  }
}
}
vlands {
  voice-vlan {
    vlan-id 10;
  }
  employee-vlan {
    vlan-id 20;
  }
  guest-vlan {
    vlan-id 30;
  }
  camera-vlan {
    vlan-id 40;
  }
}
system {
  commit synchronize;
}
chassis {
  redundancy {
    graceful-switchover;
  }
}
ethernet-switching-options {
  nonstop-bridging;
}

```

### Configuring RSTP and Nonstop Bridging on Switch 2

#### CLI Quick Configuration

To quickly configure RSTP and nonstop bridging on Switch 2, copy the following commands and paste them into the switch terminal window:

```

[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30

```

```

set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/18 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 32k
set protocols rstp interface ge-0/0/14.0 cost 1000
set protocols rstp interface ge-0/0/14.0 mode point-to-point
set protocols rstp interface ge-0/0/18.0 cost 1000
set protocols rstp interface ge-0/0/18.0 mode point-to-point

```

If Switch 2 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 2, copy the following commands and paste them into the switch terminal window:

```

set chassis redundancy graceful switchover
set system commit synchronize
set ethernet-switching-options nonstop-bridging

```

#### Step-by-Step Procedure

To configure RSTP and nonstop bridging on Switch 2:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```

[edit vlans]
user@switch2# set voice-vlan description "Voice VLAN"
user@switch2# set voice-vlan vlan-id 10
user@switch2# set employee-vlan description "Employee VLAN"
user@switch2# set employee-vlan vlan-id 20
user@switch2# set guest-vlan description "Guest VLAN"
user@switch2# set guest-vlan vlan-id 30
user@switch2# set camera-vlan description "Camera VLAN"
user@switch2# set camera-vlan vlan-id 40

```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```

[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch2# set ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]

```

3. Configure the port mode for the interfaces:

```

[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
user@switch2# set ge-0/0/18 unit 0 family ethernet-switching port-mode trunk

```

4. Configure RSTP on the switch:

```

[edit protocols]
user@switch2# rstp bridge-priority 32k
user@switch2# rstp interface ge-0/0/14.0 cost 1000
user@switch2# rstp interface ge-0/0/14.0 mode point-to-point
user@switch2# rstp interface ge-0/0/18.0 cost 1000
user@switch2# rstp interface ge-0/0/18.0 mode point-to-point

```

#### Step-by-Step Procedure

If Switch 2 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 2:

1. Enable graceful Routing Engine switchover (GRES):

```

[edit chassis redundancy]

```

- ```
user@switch2# set graceful-switchover
```
2. Configure the switch to always synchronize configuration changes between the Routing Engines:
 

```
[edit system]
user@switch2# set commit synchronize
```

If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.
  3. Enable nonstop bridging:
 

```
[edit ethernet-switching-options]
user@switch2# set nonstop-bridging
```



**NOTE:** This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

**Results** Check the results of the configuration:

```
user@switch2> show configuration
interfaces {
  ge-0/0/14 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  ge-0/0/18 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
}
protocols {
  rstp {
    bridge-priority 32k;
    interface ge-0/0/14.0 {
      cost 1000;
      mode point-to-point;
    }
    interface ge-0/0/18.0 {
      cost 1000;
      mode point-to-point;
    }
  }
}
```

```
    }  
  }  
}  
vpls {  
  voice-vlan {  
    vlan-id 10;  
  }  
  employee-vlan {  
    vlan-id 20;  
  }  
  guest-vlan {  
    vlan-id 30;  
  }  
  camera-vlan {  
    vlan-id 40;  
  }  
}  
system {  
  commit synchronize;  
}  
chassis {  
  redundancy {  
    graceful-switchover;  
  }  
  ethernet-switching-options {  
    nonstop-bridging;  
  }  
}
```

---

### Configuring RSTP and Nonstop Bridging on Switch 3

#### CLI Quick Configuration

To quickly configure RSTP and nonstop bridging on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]  
set vlans voice-vlan description "Voice VLAN"  
set vlans voice-vlan vlan-id 10  
set vlans employee-vlan description "Employee VLAN"  
set vlans employee-vlan vlan-id 20  
set vlans guest-vlan description "Guest VLAN"  
set vlans guest-vlan vlan-id 30  
set vlans camera-vlan description "Camera VLAN"  
set vlans camera-vlan vlan-id 40  
set interfaces ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]  
set interfaces ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]  
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]  
set interfaces ge-0/0/26 unit 0 family ethernet-switching port-mode trunk  
set interfaces ge-0/0/28 unit 0 family ethernet-switching port-mode trunk  
set interfaces ge-0/0/24 unit 0 family ethernet-switching port-mode trunk  
set protocols rstp bridge-priority 8k  
set protocols rstp interface ge-0/0/26.0 cost 1000  
set protocols rstp interface ge-0/0/26.0 mode point-to-point  
set protocols rstp interface ge-0/0/28.0 cost 1000  
set protocols rstp interface ge-0/0/28.0 mode point-to-point  
set protocols rstp interface ge-0/0/24.0 cost 1000  
set protocols rstp interface ge-0/0/24.0 mode point-to-point
```



If Switch 3 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 3, copy the following commands and paste them into the switch terminal window:

```
set chassis redundancy graceful switchover
set system commit synchronize
set ethernet-switching-options nonstop-bridging
```

#### Step-by-Step Procedure

To configure RSTP and nonstop bridging on Switch 3:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch3# set voice-vlan description "Voice VLAN"
user@switch3# set voice-vlan vlan-id 10
user@switch3# set employee-vlan description "Employee VLAN"
user@switch3# set employee-vlan vlan-id 20
user@switch3# set guest-vlan description "Guest VLAN"
user@switch3# set guest-vlan vlan-id 30
user@switch3# set camera-vlan description "Camera VLAN"
user@switch3# set camera-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching port-mode trunk
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching port-mode trunk
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching port-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch3# rstp bridge-priority 8k
user@switch3# rstp interface ge-0/0/26.0 cost 1000
user@switch3# rstp interface ge-0/0/26.0 mode point-to-point
user@switch3# rstp interface ge-0/0/28.0 cost 1000
user@switch3# rstp interface ge-0/0/28.0 mode point-to-point
user@switch3# rstp interface ge-0/0/24.0 cost 1000
user@switch3# rstp interface ge-0/0/24.0 mode point-to-point
```

#### Step-by-Step Procedure

If Switch 3 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 3:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch3# set graceful-switchover
```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch3# set commit synchronize
```

If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:  

```
[edit ethernet-switching-options]  
user@switch3# set nonstop-bridging
```



**NOTE:** This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

**Results** Check the results of the configuration:

```
user@switch3> show configuration  
interfaces {  
  ge-0/0/26 {  
    unit 0 {  
      family ethernet-switching {  
        port-mode trunk;  
        vlan {  
          members [10 20 30 40];  
        }  
      }  
    }  
  }  
  ge-0/0/28 {  
    unit 0 {  
      family ethernet-switching {  
        port-mode trunk;  
        vlan {  
          members [10 20 30 40];  
        }  
      }  
    }  
  }  
  ge-0/0/24 {  
    unit 0 {  
      family ethernet-switching {  
        port-mode trunk;  
        vlan {  
          members [10 20 30 40];  
        }  
      }  
    }  
  }  
}  
protocols {  
  rstp {  
    bridge-priority 8k;  
    interface ge-0/0/26.0 {  
      cost 1000;  
    }  
  }  
}
```

```

        mode point-to-point;
    }
    interface ge-0/0/28.0 {
        cost 1000;
        mode point-to-point;
    }
    interface ge-0/0/24.0 {
        cost 1000;
        mode point-to-point;
    }
    }
    bridge-priority 8k;
    }
}
}
vllans {
    voice-vlan {
        vlan-id 10;
    }
    employee-vlan {
        vlan-id 20;
    }
    guest-vlan {
        vlan-id 30;
    }
    camera-vlan {
        vlan-id 40;
    }
}
system {
    commit synchronize;
}
chassis {
    redundancy {
        graceful-switchover;
    }
}
ethernet-switching-options {
    nonstop-bridging;
}

```

### Configuring RSTP and Nonstop Bridging on Switch 4

#### CLI Quick Configuration

To quickly configure RSTP and nonstop bridging on Switch 4, copy the following commands and paste them into the switch terminal window:

```

[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]

```

```
set interfaces ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/23 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/19 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 16k
set protocols rstp interface ge-0/0/23.0 cost 1000
set protocols rstp interface ge-0/0/23.0 mode point-to-point
set protocols rstp interface ge-0/0/19.0 cost 1000
set protocols rstp interface ge-0/0/19.0 mode point-to-point
```

If Switch 4 includes dual Routing Engines, configure NSB. To quickly configure nonstop bridging on Switch 4, copy the following commands and paste them into the switch terminal window:

```
set chassis redundancy graceful switchover
set system commit synchronize
set ethernet-switching-options nonstop-bridging
```

#### Step-by-Step Procedure

To configure RSTP and nonstop bridging on Switch 4:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch4# set voice-vlan description "Voice VLAN"
user@switch4# set voice-vlan vlan-id 10
user@switch4# set employee-vlan description "Employee VLAN"
user@switch4# set employee-vlan vlan-id 20
user@switch4# set guest-vlan description "Guest VLAN"
user@switch4# set guest-vlan vlan-id 30
user@switch4# set camera-vlan description "Camera VLAN"
user@switch4# set camera-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching port-mode trunk
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching port-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch4# rstp bridge-priority 16k
user@switch4# rstp interface all cost 1000
user@switch4# rstp interface ge-0/0/23.0 cost 1000
user@switch4# rstp interface ge-0/0/23.0 mode point-to-point
user@switch4# rstp interface ge-0/0/19.0 cost 1000
user@switch4# rstp interface ge-0/0/19.0 mode point-to-point
```

#### Step-by-Step Procedure

If Switch 4 includes dual Routing Engines, configure nonstop bridging. To configure NSB on Switch 4:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch4# set graceful-switchover
```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch4# set commit synchronize
```

If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:

```
[edit ethernet-switching-options]
user@switch4# set nonstop-bridging
```



**NOTE:** This process enables NSB for all NSB-supported Layer 2 protocols on the switch, including RSTP.

**Results** Check the results of the configuration:

```
user@switch4> show configuration
interfaces {
  ge-0/0/23 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  ge-0/0/19 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
}
protocols {
  rstp {
    bridge-priority 16k;
    interface ge-0/0/23.0 {
      cost 1000;
      mode point-to-point;
    }
    interface ge-0/0/19.0 {
      cost 1000;
      mode point-to-point;
    }
  }
}
```

```
    }  
  }  
}  
vpls {  
  voice-vlan {  
    vlan-id 10;  
  }  
  employee-vlan {  
    vlan-id 20;  
  }  
  guest-vlan {  
    vlan-id 30;  
  }  
  camera-vlan {  
    vlan-id 40;  
  }  
}  
system {  
  commit synchronize;  
}  
chassis {  
  redundancy {  
    graceful-switchover;  
  }  
}  
ethernet-switching-options {  
  nonstop-bridging;  
}
```

---

## Verification

To confirm that the configuration is working properly, perform these tasks on both Routing Engines:

- [Verifying RSTP Configuration on Switch 1 on page 4934](#)
- [Verifying RSTP Configuration on Switch 2 on page 4935](#)
- [Verifying RSTP Configuration on Switch 3 on page 4935](#)
- [Verifying RSTP Configuration on Switch 4 on page 4936](#)

### ***Verifying RSTP Configuration on Switch 1***

**Purpose** Verify the RSTP configuration on Switch 1.

**Action** Use the operational mode command:

```
user@switch1> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ge-0/0/13.0	128:526	128:526	16384.0019e25040e0	1000	BLK	ALT
ge-0/0/9.0	128:522	128:522	32768.0019e2503d20	1000	BLK	ALT
ge-0/0/11.0	128:524	128:524	8192.0019e25051e0	1000	FWD	ROOT

**Meaning** Refer to the topology in [Figure 73 on page 4921](#). The operational mode command **show spanning-tree interface** shows that **ge-0/0/13.0** is in a forwarding state. The other interfaces on Switch 1 are blocking.

#### *Verifying RSTP Configuration on Switch 2*

**Purpose** Use this procedure to verify the RSTP configuration on both Switch 2 Routing Engines.

**Action** Use the operational mode command:

```
user@switch2> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ge-0/0/14.0	128:527	128:527	32768.0019e2503d20	1000	FWD	DESC
ge-0/0/18.0	128:529	128:529	8192.0019e25051e0	1000	FWD	ROOT

**Meaning** Refer to the topology in [Figure 73 on page 4921](#). The operational mode command **show spanning-tree interface** shows that **ge-0/0/18.0** is in a forwarding state and is the root port.

#### *Verifying RSTP Configuration on Switch 3*

**Purpose** Use this procedure to verify the RSTP configuration on both Switch 3 Routing Engines.

**Action** Use the operational mode commands:

```
user@switch3> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ge-0/0/26.0	128:539	128:539	8192.0019e25051e0	1000	FWD	DESG
ge-0/0/28.0	128:541	128:541	8192.0019e25051e0	1000	FWD	DESG
ge-0/0/24.0	128:537	128:537	8192.0019e25051e0	1000	FWD	DESG

**Meaning** Refer to the topology in [Figure 73 on page 4921](#). The operational mode command **show spanning-tree interface** shows that no interface is the root interface.

#### *Verifying RSTP Configuration on Switch 4*

**Purpose** Use this procedure to verify the RSTP configuration on both Switch 4 Routing Engines.

**Action** Use the operational mode commands:

```
user@switch4> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ge-0/0/23.0	128:536	128:536	8192.0019e25051e0	1000	FWD	ROOT
ge-0/0/19.0	128:532	128:532	16384.0019e25040e0	1000	FWD	DESG

**Meaning** Refer to the topology in [Figure 73 on page 4921](#). The operational mode command **show spanning-tree interface** shows that interface **ge-0/0/23.0** is the root interface and forwarding.

- Related Documentation**
- [Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936](#)
  - [Understanding RSTP for EX Series Switches on page 4906](#)

### Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches

Multiple Spanning Tree Protocol (MSTP) is used to create a loop-free topology in networks using multiple spanning-tree regions in which each region contains multiple spanning-tree instances (MSTIs). MSTIs provide different paths for different VLANs. This functionality facilitates better load sharing across redundant links.

Up to 64 MSTIs can be created for an EX Series switch, and each MSTI can support up to 4094 VLANs.



This example describes how to configure MSTP on four EX Series switches:

- [Requirements on page 4937](#)
- [Overview and Topology on page 4937](#)
- [Configuring MSTP on Switch 1 on page 4940](#)
- [Configuring MSTP on Switch 2 on page 4943](#)
- [Configuring MSTP on Switch 3 on page 4945](#)
- [Configuring MSTP on Switch 4 on page 4948](#)
- [Verification on page 4950](#)

### Requirements

---

This example uses the following hardware and software components:

- Junos OS Release 9.0 or later for EX Series switches
- Four EX Series switches

Before you configure the switches for MSTP, be sure you have:


- Installed and connected the four switches. See the hardware documentation for your switch.
- Performed the initial software configuration on all switches. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* or *Connecting and Configuring an EX Series Switch (J-Web Procedure)*.

### Overview and Topology

---

When the number of VLANs grows in a network, MSTP provides an efficient way of creating a loop-free topology by using MSTIs. Each MSTI in the spanning-tree domain maintains its own tree. Each tree can be mapped to different links, utilizing bandwidth that would be unavailable to a single tree. MSTIs reduce the demand on system resources.

Figure 74: Network Topology for MSTP



The interfaces shown in [Figure 74 on page 4938](#) will be configured for MSTP.



**NOTE:** You can configure MSTP on logical or physical interfaces. This example shows configuring MSTP on logical interfaces.

Table 549: Components of the Topology for Configuring MSTP on EX Series Switches

Property	Settings
Switch 1	<p>The following interfaces on Switch 1 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <code>ge-0/0/9</code> is connected to Switch 2</li> <li>• <code>ge-0/0/13</code> is connected to Switch 4</li> <li>• <code>ge-0/0/11</code> is connected to Switch 3</li> </ul>
Switch 2	<p>The following interfaces on Switch 2 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <code>ge-0/0/14</code> is connected to Switch 1</li> <li>• <code>ge-0/0/18</code> is connected to Switch 3</li> </ul>

**Table 549: Components of the Topology for Configuring MSTP on EX Series Switches (*continued*)**

Property	Settings
Switch 3	<p>The following interfaces on Switch 3 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>ge-0/0/26</b> is connected to Switch 1</li> <li>• <b>ge-0/0/28</b> is connected to Switch 2</li> <li>• <b>ge-0/0/24</b> is connected to Switch 4</li> </ul>
Switch 4	<p>The following interfaces on Switch 4 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>ge-0/0/19</b> is connected to Switch 1</li> <li>• <b>ge-0/0/23</b> is connected to Switch 3</li> </ul>
VLAN names and tag IDs	<p><b>voice-vlan</b>, tag 10  <b>employee-vlan</b>, tag 20  <b>guest-vlan</b>, tag 30  <b>camera-vlan</b>, tag 40</p>
MSTIs	<p>1 2</p>
MSTI region	<b>region1</b>

The topology in [Figure 74 on page 4938](#) shows a common and internal spanning tree (CIST). The CIST is a single spanning tree connecting all devices in the network. The switch with the lowest bridge priority is elected as the root bridge of the CIST. You can control the election of the root bridge by configuring the bridge priority. Switch 3 is the root bridge of the CIST.

The ports in an MSTP topology have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.
- The *backup port* becomes the active designated port and starts forwarding data when the designated port goes down.

In this example, one MSTP region contains Switch 1, Switch 2, Switch 3, and Switch 4. Within the region, four VLANs are created:

- **voice-vlan** supports voice traffic and has the VLAN tag identifier of **10**.
- **employee-vlan** supports data traffic and has the VLAN tag identifier of **20**.
- **guest-vlan** supports guest VLAN traffic (for supplicants that fail authentication) and has the VLAN tag identifier of **30**.
- **camera-vlan** supports video traffic and has the VLAN tag identifier of **40**.

The VLANs are associated with specific interfaces on each of the four switches. Two MSTIs, 1 and 2, are then associated with the VLAN tag identifiers, and some MSTP parameters, such as cost, are configured on each switch.

### Configuring MSTP on Switch 1

---

#### CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 1, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/13 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/9 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 16k
set protocols mstp interface ge-0/0/13.0 cost 1000
set protocols mstp interface ge-0/0/13.0 mode point-to-point
set protocols mstp interface ge-0/0/9.0 cost 1000
set protocols mstp interface ge-0/0/9.0 mode point-to-point
set protocols mstp interface ge-0/0/11.0 cost 1000
set protocols mstp interface ge-0/0/11.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 16k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 1 interface ge-0/0/11.0 cost 1000
set protocols mstp msti 2 bridge-priority 8k
set protocols mstp msti 2 vlan [30 40]
```

#### Step-by-Step Procedure

To configure interfaces and MSTP on Switch 1:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch1# set voice-vlan description "Voice VLAN"
user@switch1# set voice-vlan vlan-id 10
user@switch1# set employee-vlan description "Employee VLAN"
user@switch1# set employee-vlan vlan-id 20
user@switch1# set guest-vlan description "Guest VLAN"
user@switch1# set guest-vlan vlan-id 30
user@switch1# set camera-vlan description "Camera VLAN"
user@switch1# set guest-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

- ```

user@switch1# set ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]

```
3. Configure the port mode for the interfaces:
 

```

[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching port-mode trunk
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching port-mode trunk
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching port-mode trunk

```
  4. Configure MSTP on the switch, including the two MSTIs:
 

```

[edit protocols]
user@switch1# mstp configuration-name region1
user@switch1# mstp bridge-priority 16k
user@switch1# mstp interface ge-0/0/13.0 cost 1000
user@switch1# mstp interface ge-0/0/13.0 mode point-to-point
user@switch1# mstp interface ge-0/0/9.0 cost 1000
user@switch1# mstp interface ge-0/0/9.0 mode point-to-point
user@switch1# mstp interface ge-0/0/11.0 cost 1000
user@switch1# mstp interface ge-0/0/11.0 mode point-to-point
user@switch1# mstp msti 1 bridge-priority 16k
user@switch1# mstp msti 1 vlan [10 20]
user@switch1# mstp msti 1 interface ge-0/0/11.0 cost 1000
user@switch1# mstp msti 2 bridge-priority 8k
user@switch1# mstp msti 2 vlan [30 40]

```

**Results** Check the results of the configuration:

```

user@switch1> show configuration
interfaces {
  ge-0/0/13 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
  ge-0/0/9 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
  ge-0/0/11 {
    unit 0 {
      family ethernet-switching {

```

Copyright © 2015, Juniper Networks, Inc.

## Configuring MSTP on Switch 2

**CLI Quick Configuration** To quickly configure interfaces and MSTP on Switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/18 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 32k
set protocols mstp interface ge-0/0/14.0 cost 1000
set protocols mstp interface ge-0/0/14.0 mode point-to-point
set protocols mstp interface ge-0/0/18.0 cost 1000
set protocols mstp interface ge-0/0/18.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 32k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 4k
set protocols mstp msti 2 vlan [30 40]
```

**Step-by-Step Procedure** To configure interfaces and MSTP on Switch 2:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch2# set voice-vlan description "Voice VLAN"
user@switch2# set voice-vlan vlan-id 10
user@switch2# set employee-vlan description "Employee VLAN"
user@switch2# set employee-vlan vlan-id 20
user@switch2# set guest-vlan description "Guest VLAN"
user@switch2# set guest-vlan vlan-id 30
user@switch2# set camera-vlan vlan-description "Camera VLAN"
user@switch2# set guest-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch2# set ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
user@switch2# set ge-0/0/18 unit 0 family ethernet-switching port-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch2# mstp configuration-name region1
user@switch2# mstp bridge-priority 32k
```

```
user@switch2# mstp interface ge-0/0/14.0 cost 1000
user@switch2# mstp interface ge-0/0/14.0 mode point-to-point
user@switch2# mstp interface ge-0/0/18.0 cost 1000
user@switch2# mstp interface ge-0/0/18.0 mode point-to-point
user@switch2# mstp msti 1 bridge-priority 32k
user@switch2# mstp msti 1 vlan [10 20]
user@switch2# mstp msti 2 bridge-priority 4k
user@switch2# mstp msti 2 vlan [30 40]
```

**Results** Check the results of the configuration:

```
user@switch2> show configuration
interfaces {
  ge-0/0/14 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
  ge-0/0/18 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
}
protocols {
  mstp {
    configuration-name region1;
    bridge-priority 32k;
    interface ge-0/0/14.0 {
      cost 1000;
      mode point-to-point;
    }
    interface ge-0/0/18.0 {
      cost 1000;
      mode point-to-point;
    }
    msti 1 {
      bridge-priority 32k;
      vlan [10 20];
    }
  }
}
```



```

        msti 2 {
            bridge-priority 4k;
            vlan [30 40];
        }
    }
}
vlangs {
    voice-vlan {
        vlan-id 10;
    }
    employee-vlan {
        vlan-id 20;
    }
    guest-vlan {
        vlan-id 30;
    }
    camera-vlan {
        vlan-id 40;
    }
}
}

```

### Configuring MSTP on Switch 3

#### CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 3, copy the following commands and paste them into the switch terminal window:

```

[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/26 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/28 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/24 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 8k
set protocols mstp interface ge-0/0/26.0 cost 1000
set protocols mstp interface ge-0/0/26.0 mode point-to-point
set protocols mstp interface ge-0/0/28.0 cost 1000
set protocols mstp interface ge-0/0/28.0 mode point-to-point
set protocols mstp interface ge-0/0/24.0 cost 1000
set protocols mstp interface ge-0/0/24.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 4k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 16k
set protocols mstp msti 2 vlan [30 40]

```

**Step-by-Step Procedure** To configure interfaces and MSTP on Switch 3:

1. Configure the VLANs voice-vlan, employee-vlan, guest-vlan, and camera-vlan:

```
[edit vlans]
user@switch3# set voice-vlan description "Voice VLAN"
user@switch3# set voice-vlan vlan-id 10
user@switch3# set employee-vlan description "Employee VLAN"
user@switch3# set employee-vlan vlan-id 20
user@switch3# set guest-vlan description "Guest VLAN"
user@switch3# set guest-vlan vlan-id 30
user@switch3# set camera-vlan description "Camera VLAN"
user@switch3# set guest-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching port-mode trunk
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching port-mode trunk
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching port-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch3# mstp configuration-name region1
user@switch3# mstp bridge-priority 8k
user@switch3# mstp interface ge-0/0/26.0 cost 1000
user@switch3# mstp interface ge-0/0/26.0 mode point-to-point
user@switch3# mstp interface ge-0/0/28.0 cost 1000
user@switch3# mstp interface ge-0/0/28.0 mode point-to-point
user@switch3# mstp interface ge-0/0/24.0 cost 1000
user@switch3# mstp interface ge-0/0/24.0 mode point-to-point
user@switch3# mstp msti 1 bridge-priority 4k
user@switch3# mstp msti 1 vlan [10 20]
user@switch3# mstp msti 2 bridge-priority 16k
user@switch3# mstp msti 2 vlan [30 40]
```

**Results** Check the results of the configuration:

```
user@switch3> show configuration
interfaces {
  ge-0/0/26 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
}
```



```
}
vlands {
  voice-vlan {
    vlan-id 10;
  }
  employee-vlan {
    vlan-id 20;
  }
  guest-vlan {
    vlan-id 30;
  }
  camera-vlan {
    vlan-id 40;
  }
}
```

---

### Configuring MSTP on Switch 4

#### CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 4, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description "Voice VLAN"
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description "Employee VLAN"
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description "Guest VLAN"
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description "Camera VLAN"
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/23 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/19 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 16k
set protocols mstp interface ge-0/0/23.0 cost 1000
set protocols mstp interface ge-0/0/23.0 mode point-to-point
set protocols mstp interface ge-0/0/19.0 cost 1000
set protocols mstp interface ge-0/0/19.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 16k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 32k
set protocols mstp msti 2 vlan [30 40]
```

#### Step-by-Step Procedure

To configure interfaces and MSTP on Switch 4:

1. Configure the VLANs `voice-vlan`, `employee-vlan`, `guest-vlan`, and `camera-vlan`:

```
[edit vlans]
user@switch4# set voice-vlan description "Voice VLAN"
user@switch4# set voice-vlan vlan-id 10
user@switch4# set employee-vlan description "Employee VLAN"
user@switch4# set employee-vlan vlan-id 20
user@switch4# set guest-vlan description "Guest VLAN"
user@switch4# set guest-vlan vlan-id 30
user@switch4# set camera-vlan description "Camera VLAN"
user@switch4# set camera-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching port-mode trunk
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching port-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch4# mstp configuration-name region1
user@switch4# mstp bridge-priority 16k
user@switch4# mstp interface ge-0/0/23.0 cost 1000
user@switch4# mstp interface ge-0/0/23.0 mode point-to-point
user@switch4# mstp interface ge-0/0/19.0 cost 1000
user@switch4# mstp interface ge-0/0/19.0 mode point-to-point
user@switch4# mstp msti 1 bridge-priority 16k
user@switch4# mstp msti 1 vlan [10 20]
user@switch4# mstp msti 2 bridge-priority 32k
user@switch4# mstp msti 2 vlan [30 40]
```

**Results** Check the results of the configuration:

```
user@switch4> show configuration
interfaces {
  ge-0/0/23 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
  ge-0/0/19 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
}
```

```
protocols {
  mstp {
    configuration-name region1;
    bridge-priority 16k;
    interface ge-0/0/23.0 {
      cost 1000;
      mode point-to-point;
    }
    interface ge-0/0/19.0 {
      cost 1000;
      mode point-to-point;
    }
    msti 1 {
      bridge-priority 16k;
      vlan [10 20];
    }
    msti 2 {
      bridge-priority 32k;
      vlan [30 40];
    }
  }
}
vpls {
  voice-vlan {
    vlan-id 10;
  }
  employee-vlan {
    vlan-id 20;
  }
  guest-vlan {
    vlan-id 30;
  }
  camera-vlan {
    vlan-id 40;
  }
}
```

---

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying MSTP Configuration on Switch 1 on page 4950](#)
- [Verifying MSTP Configuration on Switch 2 on page 4952](#)
- [Verifying MSTP Configuration on Switch 3 on page 4953](#)
- [Verifying MSTP Configuration on Switch 4 on page 4955](#)

### *Verifying MSTP Configuration on Switch 1*

**Purpose** Verify the MSTP configuration on Switch 1.

**Action** Issue the operational mode commands **show spanning-tree interface** and **show spanning-tree bridge**:

```
user@switch1> show spanning-tree interface
```

## Spanning tree interface parameters for instance 0

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/13.0 | 128:527 | 128:525               | 16384.0019e25040e0      | 1000         | BLK   | ALT  |
| ge-0/0/9.0  | 128:529 | 128:513               | 32768.0019e2503d20      | 1000         | BLK   | ALT  |
| ge-0/0/11.0 | 128:531 | 128:513               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |

## Spanning tree interface parameters for instance 1

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/13.0 | 128:527 | 128:525               | 16384.0019e25040e0      | 1000         | BLK   | ALT  |
| ge-0/0/9.0  | 128:529 | 128:513               | 32768.0019e2503d20      | 1000         | BLK   | ALT  |
| ge-0/0/11.0 | 128:531 | 128:513               | 4096.0019e25051e0       | 1000         | FWD   | ROOT |

## Spanning tree interface parameters for instance 2

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/13.0 | 128:527 | 128:527               | 8192.0019e25044e0       | 1000         | FWD   | DESG |
| ge-0/0/9.0  | 128:529 | 128:513               | 4096.0019e2503d20       | 1000         | FWD   | ROOT |
| ge-0/0/11.0 | 128:531 | 128:531               | 8192.0019e25044e0       | 1000         | FWD   | DESG |

```
user@switch1> show spanning-tree bridge
```

## STP bridge parameters

```
Context ID                : 0
Enabled protocol          : MSTP
```

## STP bridge parameters for CIST

```
Root ID                   : 8192.00:19:e2:50:51:e0
Root cost                  : 1000
Root port                  : ge-0/0/13.0
CIST regional root         : 8192.00:19:e2:50:51:e0
CIST internal root cost    : 2000
Hello time                  : 2 seconds
Maximum age                 : 20 seconds
Forward delay               : 15 seconds
Hop count                   : 18
Message age                 : 0
Number of topology changes : 3
Time since last topology change : 921 seconds
Local parameters
  Bridge ID                 : 16384.00:19:e2:50:44:e0
  Extended system ID        : 0
  Internal instance ID      : 0
```

## STP bridge parameters for MSTI 1

```
MSTI regional root        : 4096.00:19:e2:50:51:e0
Root cost                  : 1000
Root port                  : ge-0/0/13.0
Hello time                  : 2 seconds
Maximum age                 : 20 seconds
Forward delay               : 15 seconds
Hop count                   : 18
Local parameters
  Bridge ID                 : 16384.00:19:e2:50:44:e0
  Extended system ID        : 0
  Internal instance ID      : 1
```

## STP bridge parameters for MSTI 2

```

MSTI regional root      : 4096.00:19:e2:50:3d:20
Root cost               : 1000
Root port              : ge-0/0/9.0
Hello time              : 2 seconds
Maximum age             : 20 seconds
Forward delay           : 15 seconds
Hop count               : 19
Local parameters
  Bridge ID             : 8192.00:19:e2:50:44:e0
  Extended system ID    : 0
  Internal instance ID   : 2

```

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

#### *Verifying MSTP Configuration on Switch 2*

**Purpose** Verify the MSTP configuration on Switch 2.

**Action** Issue the operational mode commands **show spanning-tree interface** and **show spanning-tree bridge**:

```
user@switch2> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/14.0 | 128:513 | 128:513               | 32768.0019e2503d20      | 1000         | FWD   | DESC |
| ge-0/0/18.0 | 128:519 | 128:515               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |

Spanning tree interface parameters for instance 1

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/14.0 | 128:513 | 128:513               | 32769.0019e2503d20      | 1000         | FWD   | DESC |
| ge-0/0/18.0 | 128:519 | 128:515               | 4097.0019e25051e0       | 1000         | FWD   | ROOT |

Spanning tree interface parameters for instance 2

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/14.0 | 128:513 | 128:513               | 4098.0019e2503d20       | 1000         | FWD   | DESC |
| ge-0/0/18.0 | 128:519 | 128:519               | 4098.0019e2503d20       | 1000         | FWD   | DESC |

```
user@switch2> show spanning-tree bridge
```

STP bridge parameters

```

Context ID              : 0
Enabled protocol        : MSTP

```

STP bridge parameters for CIST



```

Root ID                               : 8192.00:19:e2:50:51:e0
Root cost                             : 1000
Root port                             : ge-0/0/18.0
CIST regional root                     : 8192.00:19:e2:50:51:e0
CIST internal root cost                : 1000
Hello time                            : 2 seconds
Maximum age                           : 20 seconds
Forward delay                         : 15 seconds
Hop count                             : 19
Message age                           : 0
Number of topology changes             : 1
Time since last topology change       : 782 seconds
Local parameters
  Bridge ID                           : 32768.00:19:e2:50:3d:20
  Extended system ID                  : 0
  Internal instance ID                : 0

STP bridge parameters for MSTI 1
MSTI regional root                    : 4096.00:19:e2:50:51:e0
Root cost                             : 1000
Root port                             : ge-0/0/18.0
Hello time                            : 2 seconds
Maximum age                           : 20 seconds
Forward delay                         : 15 seconds
Hop count                             : 19
Local parameters
  Bridge ID                           : 32768.00:19:e2:50:3d:20
  Extended system ID                  : 0
  Internal instance ID                : 1

STP bridge parameters for MSTI 2
MSTI regional root                    : 4096.00:19:e2:50:3d:20
Hello time                            : 2 seconds
Maximum age                           : 20 seconds
Forward delay                         : 15 seconds
Local parameters
  Bridge ID                           : 4096.00:19:e2:50:3d:20
  Extended system ID                  : 0
  Internal instance ID                : 2

```

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles. The spanning-tree interface parameters for instance 2 show that both ports are designated ports, which means Switch 2 is the root bridge for this instance.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

### *Verifying MSTP Configuration on Switch 3*

**Purpose** Verify the MSTP configuration on Switch 3.

**Action** Issue the operational mode commands **show spanning-tree interface** and **show spanning-tree bridge**:

```
user@switch3> show spanning-tree interface
```

## Spanning tree interface parameters for instance 0

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/26.0 | 128:513 | 128:513               | 8192.0019e25051e0       | 1000         | FWD   | DESC |
| ge-0/0/28.0 | 128:515 | 128:515               | 8192.0019e25051e0       | 1000         | FWD   | DESC |
| ge-0/0/24.0 | 128:517 | 128:517               | 8192.0019e25051e0       | 1000         | FWD   | DESC |

## Spanning tree interface parameters for instance 1

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/26.0 | 128:513 | 128:513               | 4096.0019e25051e0       | 1000         | FWD   | DESC |
| ge-0/0/28.0 | 128:515 | 128:515               | 4096.0019e25051e0       | 1000         | FWD   | DESC |
| ge-0/0/24.0 | 128:517 | 128:517               | 4096.0019e25051e0       | 1000         | FWD   | DESC |

## Spanning tree interface parameters for instance 2

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/26.0 | 128:513 | 128:531               | 8192.0019e25044e0       | 1000         | BLK   | ALT  |
| ge-0/0/28.0 | 128:515 | 128:519               | 4096.0019e2503d20       | 1000         | FWD   | ROOT |
| ge-0/0/24.0 | 128:517 | 128:517               | 16384.0019e25051e0      | 1000         | FWD   | DESC |

```
user@switch3> show spanning-tree bridge
```

## STP bridge parameters

```
Context ID           : 0
Enabled protocol     : MSTP
```

## STP bridge parameters for CIST

```
Root ID              : 8192.00:19:e2:50:51:e0
CIST regional root   : 8192.00:19:e2:50:51:e0
CIST internal root cost : 0
Hello time           : 2 seconds
Maximum age          : 20 seconds
Forward delay        : 15 seconds
Number of topology changes : 3
Time since last topology change : 843 seconds
Local parameters
  Bridge ID          : 8192.00:19:e2:50:51:e0
  Extended system ID : 0
  Internal instance ID : 0
```

## STP bridge parameters for MSTI 1

```
MSTI regional root   : 4096.00:19:e2:50:51:e0
Hello time           : 2 seconds
Maximum age          : 20 seconds
Forward delay        : 15 seconds
Local parameters
  Bridge ID          : 4096.00:19:e2:50:51:e0
  Extended system ID : 0
  Internal instance ID : 1
```

## STP bridge parameters for MSTI 2

```
MSTI regional root   : 4096.00:19:e2:50:3d:20
Root cost             : 1000
Root port             : ge-0/0/28.0
Hello time           : 2 seconds
Maximum age          : 20 seconds
```

```

Forward delay           : 15 seconds
Hop count               : 19
Local parameters
  Bridge ID             : 16384.00:19:e2:50:51:e0
  Extended system ID    : 0
  Internal instance ID   : 2

```

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles. Switch 3 is the root bridge for instance 0, which is the CIST, as well as for instance 1. In both instances, all ports on Switch 3 are designated ports.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

#### *Verifying MSTP Configuration on Switch 4*

**Purpose** Verify the MSTP configuration on Switch 4.

**Action** Issue the operational mode commands **show spanning-tree interface** and **show spanning-tree bridge**:

```
user@switch4> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/23.0 | 128:523 | 128:517               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |
| ge-0/0/19.0 | 128:525 | 128:525               | 16384.0019e25040e0      | 1000         | FWD   | DESG |

```
Spanning tree interface parameters for instance 1
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/23.0 | 128:523 | 128:517               | 4096.0019e25051e0       | 1000         | FWD   | ROOT |
| ge-0/0/19.0 | 128:525 | 128:525               | 16384.0019e25040e0      | 1000         | FWD   | DESG |

```
Spanning tree interface parameters for instance 2
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/23.0 | 128:523 | 128:517               | 16384.0019e25051e0      | 1000         | BLK   | ALT  |
| ge-0/0/19.0 | 128:525 | 128:527               | 8192.0019e25044e0       | 1000         | FWD   | ROOT |

```
user@switch4> show spanning-tree bridge
```

```
STP bridge parameters
```

```
Context ID           : 0
Enabled protocol     : MSTP
```

```
STP bridge parameters for CIST
```

```

Root ID              : 8192.00:19:e2:50:51:e0
Root cost            : 0
Root port            : ge-0/0/23.0
CIST regional root   : 8192.00:19:e2:50:51:e0

```

```
CIST internal root cost      : 1000
Hello time                  : 2 seconds
Maximum age                 : 20 seconds
Forward delay               : 15 seconds
Hop count                   : 19
Message age                 : 0
Number of topology changes  : 4
Time since last topology change : 887 seconds
Local parameters
  Bridge ID                  : 16384.00:19:e2:50:40:e0
  Extended system ID         : 0
  Internal instance ID       : 0

STP bridge parameters for MSTI 1
MSTI regional root         : 4096.00:19:e2:50:51:e0
Root cost                   : 1000
Root port                   : ge-0/0/23.0
Hello time                  : 2 seconds
Maximum age                 : 20 seconds
Forward delay               : 15 seconds
Hop count                   : 19
Local parameters
  Bridge ID                  : 16384.00:19:e2:50:40:e0
  Extended system ID         : 0
  Internal instance ID       : 1

STP bridge parameters for MSTI 2
MSTI regional root         : 4096.00:19:e2:50:3d:20
Root cost                   : 2000
Root port                   : ge-0/0/19.0
Hello time                  : 2 seconds
Maximum age                 : 20 seconds
Forward delay               : 15 seconds
Hop count                   : 18
Local parameters
  Bridge ID                  : 32768.00:19:e2:50:40:e0
  Extended system ID         : 0
  Internal instance ID       : 2
```

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

- Related Documentation**
- [Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919](#)
  - [Understanding MSTP for EX Series Switches on page 4904](#)

## Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches

EX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). All spanning-tree protocols use a special type of frame called a bridge protocol

data unit (BPDU) to communicate. Other devices—PC bridging applications, for example, also use BPDUs and generate their own BPDUs. These different BPDUs are not compatible. When BPDUs generated by spanning-tree protocols are transmitted to a device that uses another type of BPDU, they can cause problems on the device. Similarly, if switches within a spanning-tree topology receive BPDUs from other devices, network outages can occur because of STP miscalculations.

This example configures BPDU protection on an EX Series switch that uses RSTP. The upstream configuration is done on the edge interfaces, where outside BPDUs are often received from other devices:

- [Requirements on page 4957](#)
- [Overview and Topology on page 4957](#)
- [Configuration on page 4958](#)
- [Verification on page 4959](#)

### Requirements

This example uses the following hardware and software components:

- Two EX Series switches in an RSTP topology
- Junos OS Release 9.1 or later for EX Series switches

Before you configure the interfaces on Switch 2 for BPDU protection, be sure you have:

- RSTP enabled on the switches.



**NOTE:** By default, RSTP is enabled on all EX Series switches.

### Overview and Topology

The switches, being in an RSTP topology, support a loop-free network through the exchange of BPDUs. Receipt of outside BPDUs in an STP, RSTP, or MSTP topology, however, can lead to network outages by triggering an STP misconfiguration. To prevent such outages, enable BPDU protection on STP interfaces that could receive outside BPDUs. If an outside BPDU is received on a BPDU-protected interface, the interface shuts down to prevent the outside BPDU from accessing the STP interface.

[Figure 75 on page 4958](#) shows the topology for this example. In this example, Switch 1 and Switch 2 are configured for RSTP and create a loop-free topology. The interfaces on Switch 2 are edge access ports—edge access ports frequently receive outside BPDUs generated by PC applications.

This example configures interface **ge-0/0/5.0** and interface **ge-0/0/6.0** as edge ports on Switch 2, and then configures BPDU protection on those ports. With BPDU protection enabled, these interfaces shut down when they encounter an outside BPDU sent by the PCs connected to Switch 2.

Figure 75: BPDU Protection Topology




Table 550 on page 4958 shows the components that will be configured for BPDU protection.

Table 550: Components of the Topology for Configuring BPDU Protection on EX Series Switches

| Property                      | Settings                                                                                                                                           |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch 1 (Distribution Layer) | Switch 1 is connected to Switch 2 on a trunk interface.                                                                                            |
| Switch 2 (Access Layer)       | Switch 2 has these access ports that require BPDU protection: <ul style="list-style-type: none"> <li>• ge-0/0/5.0</li> <li>• ge-0/0/6.0</li> </ul> |

This configuration example uses RSTP topology. You also can configure BPDU protection for STP or MSTP topologies at the `[edit protocols (mstp | stp)]` hierarchy level.

### Configuration

To configure BPDU protection on two access interfaces:

#### CLI Quick Configuration

Quickly configure RSTP on the two Switch 2 interfaces, and then configure BPDU protection on all edge ports on Switch 2 by copying the following commands and pasting them into the switch terminal window:

```
[edit]
set protocols rstp interface ge-0/0/5.0 edge
set protocols rstp interface ge-0/0/6.0 edge
set protocols rstp bpdv-block-on-edge
```

**Step-by-Step Procedure** To configure RSTP on the two Switch 2 interfaces, and then configure BPDU protection:

1. Configure RSTP on interface **ge-0/0/5.0** and interface **ge-0/0/6.0**, and configure them as edge ports:

```
[edit protocols rstp]
user@switch# set interface ge-0/0/5.0 edge
user@switch# set interface ge-0/0/6.0 edge
```

2. Configure BPDU protection on all edge ports on this switch:

```
[edit protocols rstp]
user@switch# set bpdu-block-on-edge
```

**Results** Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface ge-0/0/5.0 {
  edge;
}
interface ge-0/0/6.0 {
  edge;
}
bpdu-block-on-edge;
```

### Verification

To confirm that the configuration is working properly:

- [Displaying the Interface State Before BPDU Protection Is Triggered on page 4959](#)
- [Verifying That BPDU Protection Is Working Correctly on page 4960](#)

#### *Displaying the Interface State Before BPDU Protection Is Triggered*

**Purpose** Before BPDUs can be received from PCs connected to interface **ge-0/0/5.0** and interface **ge-0/0/6.0**, confirm the interface state.

**Action** Use the operational mode command:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/5.0 | 128:518 | 128:518               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/6.0 | 128:519 | 128:519               | 32768.0019e2503f00      | 20000        | FWD   | DESG |

[output truncated]

**Meaning** The output from the operational mode command **show spanning-tree interface** shows that **ge-0/0/5.0** and interface **ge-0/0/6.0** are ports in a forwarding state.

**Verifying That BPDU Protection Is Working Correctly**

**Purpose** In this example, the PCs connected to Switch 2 start sending BPDUs to interface **ge-0/0/5.0** and interface **ge-0/0/6.0**. Verify that BPDU protection is working on the interfaces.

**Action** Use the operational mode command:

```
user@switch> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface                  | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|----------------------------|---------|--------------------|----------------------|-----------|-------|------|
| ge-0/0/0.0                 | 128:513 | 128:513            | 32768.0019e2503f00   | 20000     | BLK   | DIS  |
| ge-0/0/1.0                 | 128:514 | 128:514            | 32768.0019e2503f00   | 20000     | BLK   | DIS  |
| ge-0/0/2.0                 | 128:515 | 128:515            | 32768.0019e2503f00   | 20000     | BLK   | DIS  |
| ge-0/0/3.0                 | 128:516 | 128:516            | 32768.0019e2503f00   | 20000     | FWD   | DESG |
| ge-0/0/4.0                 | 128:517 | 128:517            | 32768.0019e2503f00   | 20000     | FWD   | DESG |
| ge-0/0/5.0<br>(Bpdu-Incon) | 128:518 | 128:518            | 32768.0019e2503f00   | 20000     | BLK   | DIS  |
| ge-0/0/6.0<br>(Bpdu-Incon) | 128:519 | 128:519            | 32768.0019e2503f00   | 20000     | BLK   | DIS  |
| ge-0/0/7.0                 | 128:520 | 128:1              | 16384.00aabbcc0348   | 20000     | FWD   | ROOT |
| ge-0/0/8.0                 | 128:521 | 128:521            | 32768.0019e2503f00   | 20000     | FWD   | DESG |
| [output truncated]         |         |                    |                      |           |       |      |

**Meaning** When BPDUs are sent from the PCs to interface **ge-0/0/5.0** and interface **ge-0/0/6.0** on Switch 2, the output from the operational mode command **show spanning-tree interface** shows that the interfaces have transitioned to a BPDU inconsistent state. The BPDU inconsistent state causes the interfaces to shut down.

Disabling the BPDU protection configuration on an interface does not automatically re-enable the interface. However, if the **disable-timeout (Spanning Trees)** statement has been included in the BPDU configuration, the interface does return to service after the timer expires. Otherwise, you must use the operational mode command **clear ethernet-switching bpd-error** to unblock and re-enable the interface.

If the PCs connected to Switch 2 send BPDUs to the interfaces again, BPDU protection is triggered once more and the interfaces transition back to the BPDU inconsistent state, causing them to shut down. In such cases, you need to find and repair the misconfiguration on the PCs that is sending BPDUs to Switch 2.

- Related Documentation**
- [Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919](#)
  - [Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4961](#)
  - [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches on page 4967](#)



- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on EX Series Switches on page 4971](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches on page 4914](#)

## Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches

Spanning-tree protocols support loop-free network communication through the exchange of a special type of frame called a bridge protocol data unit (BPDU). However, when BPDUs generated by spanning-tree protocols are communicated to devices on which spanning-tree protocols are not configured, these devices recognize the BPDUs, which can lead to network outages. You can, however, enable BPDU protection on switch interfaces to prevent BPDUs generated by spanning-tree protocols from passing through those interfaces. When BPDU protection is enabled, an interface shuts down or drops BPDU packets when any incompatible BPDU is encountered, thereby preventing the BPDUs generated by spanning-tree protocols from reaching the switch. When an interface is configured to drop BPDU packets, all traffic except the incompatible BPDUs can pass through the interface.



**NOTE:** The BPDU drop feature can be specified only on interfaces on which no spanning-tree protocol is configured.

This example configures BPDU protection on STP switch downstream interfaces that connect to two PCs:

- [Requirements on page 4961](#)
- [Overview and Topology on page 4962](#)
- [Configuration on page 4963](#)
- [Verification on page 4965](#)

### Requirements

This example uses the following hardware and software components:

- One EX Series switch in an RSTP topology
- One EX Series switch that is not in any spanning-tree topology
- Junos OS Release 9.1 or later for EX Series switches

Before you configure the interfaces on Switch 2 for BPDU protection, be sure you have:

- Ensured that RSTP is operating on Switch 1.
- Disabled or enabled RSTP on Switch 2 (depending on the configuration that you plan to implement.)

If you want to enable the BPDU shutdown feature, then it is optional to disable spanning-tree protocols on the interface.



**NOTE:** By default, RSTP is enabled on all EX Series switches.

---

## Overview and Topology

EX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). All spanning-tree protocols use a special type of frame called a BPDU to communicate. Other devices also use BPDUs—PC bridging applications, for example, generate their own BPDUs. These different BPDUs are not compatible. When BPDUs generated by spanning-tree protocols are transmitted to a device that uses another type of BPDU, they can cause problems on the device. Similarly, if switches within a spanning-tree topology receive BPDUs from other devices, network outages can occur because of the miscalculations caused by the outside BPDUs. Therefore, you must configure BPDU protection on interfaces in a spanning-tree topology to avoid network outages.

This example explains how to block outside BPDUs from reaching a switch interface connected to devices that are not part of the STP topology. This example addresses two scenarios. In the first scenario, an interface is shutdown when it encounters an outside BPDU. In the second scenario, an interface drops only BPDU packets while retaining the status of the interface as up and allowing all other traffic to pass through the interface.

[Figure 75 on page 4958](#) shows the topology for this example. Switch 1 and Switch 2 are connected through a trunk interface. Switch 1 is configured for RSTP while Switch 2 has a spanning-tree protocol configured on it for the first scenario, and does not have a spanning-tree protocol configured on it for the second scenario.

In the first scenario, this example configures downstream BPDU protection on Switch 2 interfaces **ge-0/0/5.0** and **ge-0/0/6.0** when the default spanning-tree protocol (RSTP) is not disabled on these interfaces. When BPDU protection is enabled with the **shutdown** statement, the switch interfaces will shut down if BPDUs generated by the laptops attempt to access Switch 2.

In the second scenario, this example configures downstream BPDU protection on Switch 2 interfaces **ge-0/0/5.0** and **ge-0/0/6.0** when the default spanning-tree protocol (RSTP) is disabled on these interfaces. When BPDU protection is enabled with the **drop** statement, the switch interfaces drop only the BPDUs while allowing remaining traffic to pass through and retaining their status as up if BPDUs generated by the laptops attempt to access Switch 2.



**CAUTION:** When configuring BPDU protection on an interface without spanning trees connected to a switch with spanning trees, be careful that you do not configure BPDU protection on all interfaces. Doing so could prevent

BPDUs being received on switch interfaces (such as a trunk interface) that you intended to have receive BPDUs from a switch with spanning trees.

Figure 76: BPDU Protection Topology




Table 551 on page 4963 shows the components that will be configured for BPDU protection.

Table 551: Components of the Topology for Configuring BPDU Protection on EX Series Switches

| Property                      | Settings                                                                                                                                                                      |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch 1 (Distribution Layer) | Switch 1 is connected to Switch 2 through a trunk interface. Switch 1 is configured for RSTP.                                                                                 |
| Switch 2 (Access Layer)       | Switch 2 has two downstream access ports connected to laptops: <ul style="list-style-type: none"> <li>• <code>ge-0/0/5.0</code></li> <li>• <code>ge-0/0/6.0</code></li> </ul> |

### Configuration

To configure BPDU protection on the interfaces:

#### CLI Quick Configuration

This is the first scenario that explains configuration for the **shutdown** statement. To quickly configure BPDU protection on Switch 2 for the **shutdown** statement, copy the following commands and paste them into the switch terminal window:

[edit]

```
user@switch# set ethernet-switching-options bpdu-block interface ge-0/0/5.0 shutdown
[edit]
user@switch# set ethernet-switching-options bpdu-block interface ge-0/0/6.0 shutdown
```

**Step-by-Step  
Procedure**

To configure BPDU protection for the **shutdown** statement:

1. Configure the BPDU **shutdown** statement on the downstream interface **ge-0/0/5.0** on Switch 2:

```
[edit ethernet-switching-options]
user@switch# set bpdu-block interface ge-0/0/5.0 shutdown
```

2. Configure the BPDU **shutdown** statement on the downstream interface **ge-0/0/6.0** on Switch 2:

```
[edit ethernet-switching-options]
user@switch# set bpdu-block interface ge-0/0/6.0 shutdown
```

**Results**

Check the results of the configuration:

```
user@switch> show ethernet-switching-options
bpdu-block {
  interface ge-0/0/5.0 {
    shutdown;
  }
  interface ge-0/0/6.0 {
    shutdown;
  }
}
```

**CLI Quick  
Configuration**

This is the second scenario that explains configuration for the **drop** statement. To quickly configure BPDU protection on Switch 2 for the **drop** statement, copy the following commands and paste them into the switch terminal window:

```
[edit]
user@switch# set protocols rstp interface ge-0/0/5.0 disable
user@switch# set protocols rstp interface ge-0/0/6.0 disable
user@switch# set ethernet-switching-options bpdu-block interface ge-0/0/5.0 drop
user@switch# set ethernet-switching-options bpdu-block interface ge-0/0/6.0 drop
```



**NOTE:** You can also disable RSTP globally using the `delete protocols rstp`, the `set protocols rstp disable`, or the `set protocols rstp interface all disable` command.

**Step-by-Step  
Procedure**

To configure BPDU protection for the **drop** statement:

1. Disable RSTP on both the interfaces **ge-0/0/5.0** and **ge-0/0/6.0** interfaces:

```
[edit]
user@switch# set protocols rstp interface ge-0/0/5.0 disable
user@switch# set protocols rstp interface ge-0/0/6.0 disable
```

2. Configure the BPDU **drop** statement on the downstream interface **ge-0/0/5.0** on Switch 2:

```
[edit ethernet-switching-options]
user@switch# set bpdu-block interface ge-0/0/5.0 drop
```

- Configure the BPDU **drop** statement on the downstream interface **ge-0/0/6.0** on Switch 2:

```
[edit ethernet-switching-options]
user@switch# set bpd-block interface ge-0/0/6.0 drop
```

**Results** Check the results of the configuration:

```
user@switch> show protocols rstp
interface ge-0/0/5.0 {
  disable;
}
interface ge-0/0/6.0 {
  disable;
}
user@switch> show ethernet-switching-options
bpd-block {
  interface ge-0/0/5.0 {
    drop;
  }
  interface ge-0/0/6.0 {
    drop;
  }
}
```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Displaying the Interface State Before BPDU Protection Is Triggered on page 4965](#)
- [Verifying That BPDU Shutdown Protection Is Working Correctly on page 4965](#)
- [Verifying That BPDU Drop Protection Is Working Correctly on page 4966](#)

#### *Displaying the Interface State Before BPDU Protection Is Triggered*

**Purpose** Before any BPDUs can be received on Switch 2 on either interface **ge-0/0/5.0** or interface **ge-0/0/6.0**, confirm the state of those interfaces.

**Action** Use the operational mode command **show ethernet-switching interfaces**:

```
user@switch> show ethernet-switching interfaces
```

| Interface  | State | VLAN members | Tag | Tagging  | Blocking  |
|------------|-------|--------------|-----|----------|-----------|
| ge-0/0/5.0 | up    | default      |     | untagged | unblocked |
| ge-0/0/6.0 | up    | default      |     | untagged | unblocked |

**Meaning** The output from the operational mode command **show ethernet-switching interfaces** shows that **ge-0/0/5.0** and interface **ge-0/0/6.0** are **up** and **unblocked**.

#### *Verifying That BPDU Shutdown Protection Is Working Correctly*

**Purpose** Verify that BPDU protection is working correctly in the network by checking to see whether BPDUs have been blocked appropriately.

**Action** Issue `show ethernet-switching interfaces` to see what happened when the BPDUs reached the two interfaces configured for BPDU protection on Switch 2:

```
user@switch> show ethernet-switching interfaces
Interface  State  VLAN members  Tag  Tagging  Blocking
ge-0/0/5.0  down  default              untagged  Disabled by bpdu-control
ge-0/0/6.0  down  default              untagged  Disabled by bpdu-control
```

**Meaning** When the BPDUs sent from laptops reached interfaces `ge-0/0/5.0` and `ge-0/0/6.0` on Switch 2, the interfaces transitioned to a BPDU inconsistent state, shutting down the two interfaces to prevent BPDUs from reaching the laptops.

You need to re-enable the blocked interfaces. There are two ways to do this. If you included the statement `disable-timeout (Spanning Trees)` in the BPDU configuration, the interface returns to service after the timer expires. Otherwise, use the operational mode command `clear ethernet-switching bpdu-error` to unblock and re-enable `ge-0/0/5.0` and `ge-0/0/6.0`. This command will only re-enable an interface but the BPDU configuration for the interface will continue to exist unless you remove the BPDU configuration explicitly.

If BPDUs reach the downstream interfaces on Switch 2 again, BPDU protection is triggered again and the interfaces shut down. In such cases, you must find and repair the misconfiguration that is sending BPDUs to interfaces `ge-0/0/5.0` and `ge-0/0/6.0`.

#### *Verifying That BPDU Drop Protection Is Working Correctly*

**Purpose** Verify that BPDU drop protection is working correctly in the network by checking to see whether BPDUs have been blocked appropriately.

**Action** Issue `show ethernet-switching interfaces` to see what happened when the BPDUs reached the two interfaces configured for BPDU protection on Switch 2:

```
user@switch> show ethernet-switching interfaces
Interface  State  VLAN members  Tag  Tagging  Blocking
ge-0/0/5.0  up     default              untagged  unblocked-xSTP bpdu
                                     filter enabled
ge-0/0/6.0  up     default              untagged  unblocked-xSTP bpdu
                                     filter enabled
```

**Meaning** When the BPDUs sent from laptops reached interfaces `ge-0/0/5.0` and `ge-0/0/6.0` on Switch 2, the interfaces dropped those BPDUs to prevent them from reaching Switch 2, and the state of both the interfaces is `up`.

**Related Documentation**

- [Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919](#)
- [Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4956](#)
- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches on page 4967](#)

- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on EX Series Switches on page 4971](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches on page 4914](#)

## Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches

EX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Loop protection increases the efficiency of STP, RSTP, and MSTP by preventing interfaces from moving into a forwarding state that would result in a loop opening up in the network.

This example describes how to configure loop protection for an interface on an EX Series switch in an RSTP topology:

- [Requirements on page 4967](#)
- [Overview and Topology on page 4967](#)
- [Configuration on page 4969](#)
- [Verification on page 4969](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.1 or later for EX Series switches
- Three EX Series switches in an RSTP topology

Before you configure the interface for loop protection, be sure you have:

- RSTP operating on the switches.



**NOTE:** By default, RSTP is enabled on all EX Series switches.

### Overview and Topology

A loop-free network in spanning-tree topologies is supported through the exchange of a special type of frame called bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic (preventing loops) and which interfaces become root ports and forward traffic.

A blocking interface can transition to the forwarding state in error if the interface stops receiving BPDUs from its designated port on the segment. Such a transition error can occur when there is a hardware error on the switch or software configuration error between the switch and its neighbor. When this happens, a loop opens up in the spanning tree.

Loops in a Layer 2 topology cause broadcast, unicast, and multicast frames to continuously circle the looped network. As a switch processes a flood of frames in a looped network, its resources become depleted and the ultimate result is a network outage.




**CAUTION:** An interface can be configured for either loop protection or root protection, but not for both.

Three EX Series switches are displayed in [Figure 77 on page 4968](#). In this example, they are configured for RSTP and create a loop-free topology. Interface **ge-0/0/6** is blocking traffic between Switch 3 and Switch 1; thus, traffic is forwarded through interface **ge-0/0/7** on Switch 2. BPDUs are being sent from the root bridge on Switch 1 to both of these interfaces.

This example shows how to configure loop protection on interface **ge-0/0/6** to prevent it from transitioning from a blocking state to a forwarding state and creating a loop in the spanning-tree topology.

**Figure 77: Network Topology for Loop Protection**



[Table 552 on page 4968](#) shows the components that will be configured for loop protection.

**Table 552: Components of the Topology for Configuring Loop Protection on EX Series Switches**

| Property | Settings                                                              |
|----------|-----------------------------------------------------------------------|
| Switch 1 | Switch 1 is the root bridge.                                          |
| Switch 2 | Switch 2 has the root port <b>ge-0/0/7</b> .                          |
| Switch 3 | Switch 3 is connected to Switch 1 through interface <b>ge-0/0/6</b> . |



A spanning-tree topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.

This configuration example uses an RSTP topology. However, you also can configure loop protection for STP or MSTP topologies at the `[edit protocols (mstp | stp)]` hierarchy level.

### Configuration

To configure loop protection on an interface:

#### CLI Quick Configuration

To quickly configure loop protection on interface **ge-0/0/6**:

```
[edit]
set protocols rstp interface ge-0/0/6 bpdutimeout-action block
```

#### Step-by-Step Procedure

To configure loop protection:

1. Configure interface **ge-0/0/6** on Switch 3:

```
[edit protocols rstp]
user@switch# set interface ge-0/0/6 bpdutimeout-action (Spanning Trees) block
```

#### Results

Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface ge-0/0/6.0 {
  bpdutimeout-action {
    block;
  }
}
```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Displaying the Interface State Before Loop Protection Is Triggered on page 4969](#)
- [Verifying That Loop Protection Is Working on an Interface on page 4970](#)

#### *Displaying the Interface State Before Loop Protection Is Triggered*

#### Purpose

Before loop protection is triggered on interface **ge-0/0/6**, confirm that the interface is blocking.

**Action** Use the operational mode command:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/5.0 | 128:518 | 128:518               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/6.0 | 128:519 | 128:2                 | 16384.00aabbcc0348      | 20000        | BLK   | ALT  |

[output truncated]

**Meaning** The output from the operational mode command **show spanning-tree interface** shows that **ge-0/0/6.0** is the alternate port and in a blocking state.

#### *Verifying That Loop Protection Is Working on an Interface*

**Purpose** Verify the loop protection configuration on interface **ge-0/0/6**. RSTP has been disabled on interface **ge-0/0/4** on Switch 1. This will stop BPDUs from being sent to interface **ge-0/0/6** and trigger loop protection on the interface.

**Action** Use the operational mode command:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/5.0 | 128:518 | 128:518               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/6.0 | 128:519 | 128:519               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |

(Loop-Incon)

[output truncated]

**Meaning** The operational mode command **show spanning-tree interface** shows that interface **ge-0/0/6.0** has detected that BPDUs are no longer being forwarded to it and has moved into a loop-inconsistent state. The loop-inconsistent state prevents the interface from transitioning to a forwarding state. The interface recovers and transitions back to its original state as soon as it receives BPDUs.

**Related Documentation**

- [Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919](#)

- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on EX Series Switches on page 4971](#)
- [Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4956](#)
- [Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4961](#)
- [Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches on page 4916](#)

## Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on EX Series Switches

EX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Root protection increases the efficiency of STP, RSTP, and MSTP by allowing network administrators to manually enforce the root bridge placement in the network.

This example describes how to configure root protection on an interface on an EX Series switch:

- [Requirements on page 4971](#)
- [Overview and Topology on page 4971](#)
- [Configuration on page 4974](#)
- [Verification on page 4974](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.1 or later for EX Series switches
- Four EX Series switches in an RSTP topology

Before you configure the interface for root protection, be sure you have:

- RSTP operating on the switches.



**NOTE:** By default, RSTP is enabled on all EX Series switches.

### Overview and Topology

Peer STP applications running on switch interfaces exchange a special type of frame called a bridge protocol data unit (BPDU). Switches communicate interface information using BPDUs to create a loop-free topology that ultimately determines the root bridge and which interfaces block or forward traffic in the spanning tree.

However, a root port elected through this process has the possibility of being wrongly elected. A user bridge application running on a PC can generate BPDUs, too, and interfere with root port election.

To prevent this from happening, enable root protection on interfaces that should not receive superior BPDUs from the root bridge and should not be elected as the root port. These interfaces are typically located on an administrative boundary and are designated ports.

When root protection is enabled on an interface:

- The interface is blocked from becoming the root port.
- Root protection is enabled for all STP instances on that interface.
- The interface is blocked only for instances for which it receives superior BPDUs. Otherwise, it participates in the spanning-tree topology.



**CAUTION:** An interface can be configured for either root protection or loop protection, but not for both.

Four EX Series switches are displayed in [Figure 78 on page 4973](#). In this example, they are configured for RSTP and create a loop-free topology. Interface **ge-0/0/7** on Switch 1 is a designated port on an administrative boundary. It connects to Switch 4. Switch 3 is the root bridge. Interface **ge-0/0/6** on Switch 1 is the root port.

This example shows how to configure root protection on interface **ge-0/0/7** to prevent it from transitioning to become the root port.

Figure 78: Network Topology for Root Protection

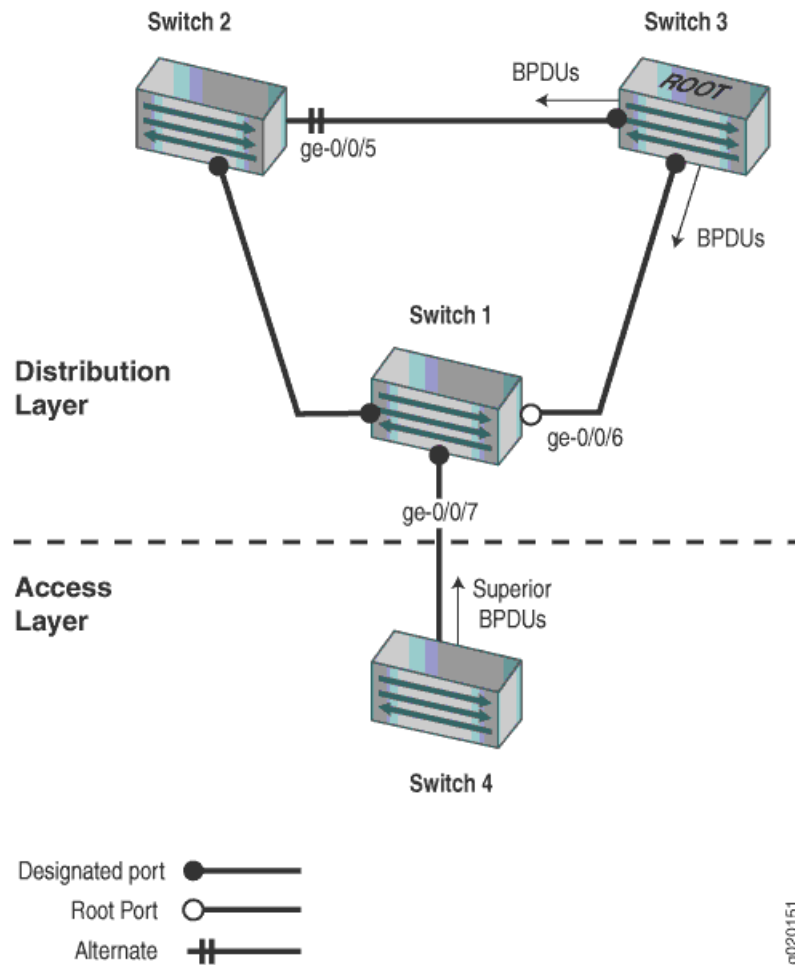


Table 553 on page 4973 shows the components that will be configured for root protection.

Table 553: Components of the Topology for Configuring Root Protection on EX Series Switches

| Property | Settings                                                                                                                                                                                                 |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switch 1 | Switch 1 is connected to Switch 4 through interface <b>ge-0/0/7</b> .                                                                                                                                    |
| Switch 2 | Switch 2 is connected to Switch 1 and Switch 3. Interface <b>ge-0/0/4</b> is the alternate port in the RSTP topology.                                                                                    |
| Switch 3 | Switch 3 is the root bridge and is connected to Switch 1 and Switch 2.                                                                                                                                   |
| Switch 4 | Switch 4 is connected to Switch 1. After root protection is configured on interface <b>ge-0/0/7</b> , Switch 4 will send superior BPDUs that will trigger root protection on interface <b>ge-0/0/7</b> . |

A spanning tree topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.

- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.

This configuration example uses an RSTP topology. However, you also can configure root protection for STP or MSTP topologies at the `[edit protocols (mstp | stp)]` hierarchy level.

---

### Configuration

To configure root protection on an interface:

#### CLI Quick Configuration

To quickly configure root protection on interface **ge-0/0/7**, copy the following command and paste it into the switch terminal window:

```
[edit]
set protocols rstp interface ge-0/0/7 no-root-port
```

#### Step-by-Step Procedure

To configure root protection:

1. Configure interface **ge-0/0/7**:  

```
[edit protocols rstp]
user@switch#
set interface ge-0/0/7 no-root-port (Spanning Trees)
```

#### Results

Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface ge-0/0/7.0 {
  no-root-port;
}
```

---

### Verification

To confirm that the configuration is working properly:

- [Displaying the Interface State Before Root Protection Is Triggered on page 4974](#)
- [Verifying That Root Protection Is Working on the Interface on page 4975](#)

#### *Displaying the Interface State Before Root Protection Is Triggered*

#### Purpose

Before root protection is triggered on interface **ge-0/0/7**, confirm the interface state.

**Action** Use the operational mode command:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/5.0 | 128:518 | 128:2                 | 16384.00aabbcc0348      | 20000        | BLK   | ALT  |
| ge-0/0/6.0 | 128:519 | 128:1                 | 16384.00aabbcc0348      | 20000        | FWD   | ROOT |
| ge-0/0/7.0 | 128:520 | 128:520               | 32768.0019e2503f00      | 20000        | FWD   | DESG |

[output truncated]

**Meaning** The output from the operational mode command **show spanning-tree interface** shows that **ge-0/0/7.0** is a designated port in a forwarding state.

#### *Verifying That Root Protection Is Working on the Interface*

**Purpose** A configuration change takes place on Switch 4. A smaller bridge priority on the Switch 4 causes it to send superior BPDUs to interface **ge-0/0/7**. Receipt of superior BPDUs on interface **ge-0/0/7** will trigger root protection. Verify that root protection is operating on interface **ge-0/0/7**.

**Action** Use the operational mode command:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/0.0 | 128:513 | 128:513               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/1.0 | 128:514 | 128:514               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/2.0 | 128:515 | 128:515               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |
| ge-0/0/3.0 | 128:516 | 128:516               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/4.0 | 128:517 | 128:517               | 32768.0019e2503f00      | 20000        | FWD   | DESG |
| ge-0/0/5.0 | 128:518 | 128:2                 | 16384.00aabbcc0348      | 20000        | BLK   | ALT  |
| ge-0/0/6.0 | 128:519 | 128:1                 | 16384.00aabbcc0348      | 20000        | FWD   | ROOT |
| ge-0/0/7.0 | 128:520 | 128:520               | 32768.0019e2503f00      | 20000        | BLK   | DIS  |

(Root-Incon)  
[output truncated]

**Meaning** The operational mode command **show spanning-tree interface** shows that interface **ge-0/0/7.0** has transitioned to a root inconsistent state. The root inconsistent state makes the interface block, discarding any received BPDUs, and prevents the interface from becoming a candidate for the root port. When the root bridge no longer receives superior STP BPDUs from the interface, the interface will recover and transition back to a forwarding state. Recovery is automatic.

**Related Documentation**

- [Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919](#)
- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches on page 4967](#)
- [Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4956](#)
- [Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4961](#)
- [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches on page 4917](#)

---

## Configuration Tasks

---

- [Forcing RSTP or VSTP to Run as IEEE 802.1D STP \(CLI Procedure\) on page 4976](#)
- [Configuring Spanning-Tree Protocols \(J-Web Procedure\) on page 4977](#)
- [Configuring VSTP \(CLI Procedure\) on page 4982](#)
- [Unblocking an Interface That Receives BPDUs in Error \(CLI Procedure\) on page 4984](#)

### Forcing RSTP or VSTP to Run as IEEE 802.1D STP (CLI Procedure)

---



**NOTE:** This procedure uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 3](#).

---

On EX Series switches running Rapid Spanning Tree Protocol (RSTP) (the default) or VLAN Spanning Tree Protocol (VSTP), you can force the original IEEE 802.1D Spanning Tree Protocol (STP) version to run in place of RSTP or VSTP. Configure the **force-version stp** statement for compatibility with older bridges that do not support RSTP or VSTP.

To force the spanning-tree protocol version to be the original IEEE 802.1D STP:

1. Enable IEEE 802.1D STP:

```
[edit protocols]  
user@switch# set (rstp | vstp) force-version stp
```

---



**NOTE:** After using the **force-version** statement to enable xSTP globally, apply the **force-version** statement for specific Layer 2 ports.

---

**Related Documentation**

- [RSTP or VSTP Forced to Run as IEEE 802.1D STP on page 4912](#)
- [Reverting to RSTP or VSTP from Forced IEEE 802.1D STP on page 4985](#)



## Configuring Spanning-Tree Protocols (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

For EX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP). You can configure STP, RSTP, and MSTP by using the J-Web interface. You can configure bridge protocol data unit (BPDU) protection on interfaces to prevent them from receiving BPDUs that could result in STP misconfigurations, which could lead to network outages.



**NOTE:** In EX4300 switches, you can configure STP only by enabling RSTP and forcing it to act as STP. You need to select the Force STP check box from the RSTP configuration page.

To configure STP, MSTP, or RSTP for an EX Series switch by using the J-Web interface:

1. Select **Configure > Switching > Spanning Tree**.

The Spanning Tree Configuration page displays the spanning-tree protocol configuration parameters and a list of interfaces configured for each spanning-tree protocol configuration.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:
  - **Add**—Creates a spanning-tree protocol configuration.
    - a. Select a protocol name.
    - b. Enter information as described in [Table 554 on page 4978](#).
    - c. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.
  - **Edit**—Modifies a selected spanning-tree protocol configuration.
    - a. Enter information as described in [Table 554 on page 4978](#).
    - b. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.
  - **Delete**—Deletes a selected spanning-tree protocol configuration.

Table 554: Spanning-Tree Protocol Configuration Parameters

| Field                                                          | Function                                                                                                                                                                                                                                                          | Your Action                                  |
|----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| General                                                        |                                                                                                                                                                                                                                                                   |                                              |
| Protocol Name                                                  | Specifies the spanning-tree protocol type: STP, MSTP, or RSTP.                                                                                                                                                                                                    | None.                                        |
| Disable                                                        | Disables spanning-tree protocols on the interface.                                                                                                                                                                                                                | To enable this option, select the check box. |
| BPDU Protect                                                   | Specifies BPDU protection on all edge interfaces on the switch.                                                                                                                                                                                                   | To enable this option, select the check box. |
| Bridge Priority                                                | Specifies the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment. | Select a value from the list.                |
| Forward Delay                                                  | Specifies the number of seconds an interface waits before changing from spanning-tree learning and listening states to the forwarding state.                                                                                                                      | Type a value.                                |
| Hello Time                                                     | Specifies the time interval in seconds at which the root bridge transmits configuration BPDUs.                                                                                                                                                                    | Type a value.                                |
| Max Age                                                        | Specifies the maximum-aging time in seconds for all MST instances. The maximum aging time is the number of seconds a switch waits without receiving spanning-tree configuration messages before attempting a reconfiguration.                                     | Type a value.                                |
| Max Hops                                                       | (MSTP only) Specifies the number of hops in a region before the BPDU is discarded.                                                                                                                                                                                | Type a value.                                |
| Configuration Name                                             | (MSTP only) Specifies the MSTP region name carried in the MSTP BPDUs.                                                                                                                                                                                             | Type a name.                                 |
| Revision Level                                                 | (MSTP only) Specifies the revision number of the MSTP configuration.                                                                                                                                                                                              | Type a value.                                |
| Force STP Version                                              | Enables or disables STP.                                                                                                                                                                                                                                          | To enable this option, select the check box. |
| <b>NOTE:</b> This option is supported only on EX4300 switches. |                                                                                                                                                                                                                                                                   |                                              |
| Ports                                                          |                                                                                                                                                                                                                                                                   |                                              |

Table 554: Spanning-Tree Protocol Configuration Parameters (*continued*)

| Field                                                         | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface Name                                                | Specifies an interface for the spanning-tree protocol.                                                                                                                                                                                                                                                                                                                                                                                                                              | <ol style="list-style-type: none"> <li>1. Click the <b>Ports</b> tab.</li> <li>2. Choose one of the following options: <ul style="list-style-type: none"> <li>• Click <b>Add</b> and select an interface from the list. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.</li> <li>• Select an interface in the <b>Port/State</b> table and click <b>Edit</b>.</li> <li>• To delete an interface from the configuration, select it in the <b>Port/State</b> table and click <b>Remove</b>.</li> </ul> </li> </ol> |
| Cost                                                          | Specifies the link cost to determine which bridge is the designated bridge and which interface is the designated interface.                                                                                                                                                                                                                                                                                                                                                         | Type a value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Priority                                                      | Specifies the interface priority to determine which interface is elected as the root port.                                                                                                                                                                                                                                                                                                                                                                                          | Select a value from the list.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Disable Port                                                  | Disables the spanning-tree protocol on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                               | To enable the option, select the check box.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>NOTE:</b> This option is not supported on EX4300 switches. |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Edge                                                          | Configures the interface as an edge interface. Edge interfaces immediately transition to a forwarding state.                                                                                                                                                                                                                                                                                                                                                                        | To enable the option, select the check box.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| No Root Port                                                  | Specifies an interface as a spanning-tree designated port. If the bridge receives superior STP BPDUs on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge that should not be the root bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked. | To enable the option, select the check box.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Interface Mode                                                | Specifies the link mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <ol style="list-style-type: none"> <li>1. To enable the option, select the check box.</li> <li>2. Select one of the following: <ul style="list-style-type: none"> <li>• <b>Point to Point</b>—For a full-duplex link, the default link mode is point-to-point.</li> <li>• <b>Shared</b>—For a half-duplex link, the default link mode is shared.</li> </ul> </li> </ol>                                                                                                                                                                                               |

Table 554: Spanning-Tree Protocol Configuration Parameters (*continued*)

| Field                      | Function                                                                                                                                                                                                                                                          | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BPDU Timeout Action        | Specifies the BPDU timeout action for the interface.                                                                                                                                                                                                              | <p>Select one of the following options:</p> <ul style="list-style-type: none"> <li>• <b>Log</b></li> <li>• <b>Block</b></li> </ul> <p><b>NOTE:</b> For EX4300 switches, you can select one of the following options:</p> <ul style="list-style-type: none"> <li>• <b>Alarm</b></li> <li>• <b>Block</b></li> </ul>                                                                                                                      |
| <b>MSTI</b><br>(MSTP only) |                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| MSTI Name                  | Specifies a name (an MSTI ID) for the MST instance.                                                                                                                                                                                                               | <ol style="list-style-type: none"> <li>1. Click the <b>MSTI</b> tab.</li> <li>2. Choose one of the following options: <ul style="list-style-type: none"> <li>• Click <b>Add</b>.</li> <li>• Select an MSTI ID and click <b>Edit</b>.</li> <li>• To delete an MSTI from the configuration, select the MSTI ID and click <b>Remove</b>.</li> </ul> </li> </ol> <p><b>NOTE:</b> For EX4300 switches, the MSTI ID can be 1 through 64.</p> |
| Bridge Priority            | Specifies the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment. | Select a value from the list.                                                                                                                                                                                                                                                                                                                                                                                                          |
| VLAN ID                    | Specifies the VLAN for the MST instance.                                                                                                                                                                                                                          | <p>In the VLAN box, choose one of the following options:</p> <ul style="list-style-type: none"> <li>• Click <b>Add</b>, select a VLAN from the list, and click <b>OK</b>.</li> <li>• To remove a VLAN association, select the VLAN ID, click <b>Remove</b>, and click <b>OK</b>.</li> </ul>                                                                                                                                            |

Table 554: Spanning-Tree Protocol Configuration Parameters (*continued*)

| Field      | Function                                     | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interfaces | Specifies an interface for the MST instance. | <ol style="list-style-type: none"> <li>1. In the Interfaces box, click <b>Add</b> and select an interface from the list, or select an interface from the list and click <b>Edit</b>.</li> <li>2. Specify the link cost to determine which bridge is the designated bridge and which interface is the designated interface.</li> <li>3. Specify the interface priority to determine which interface is elected as the root port.</li> <li>4. If you want to disable the interface, select the check box.</li> <li>5. Click OK.</li> </ol> <p>To delete an interface configuration, select the interface, click <b>Remove</b>, and click <b>OK</b>.</p> |

- Related Documentation**
- [Configuring STP \(CLI Procedure\)](#)
  - [Monitoring Spanning-Tree Protocols on page 5027](#)
  - [Unblocking an Interface That Receives BPDUs in Error \(CLI Procedure\) on page 4984](#)
  - [Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4956](#)
  - [Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936](#)
  - [Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919](#)

## Configuring VSTP (CLI Procedure)

The default spanning-tree protocol for EX Series switches is Rapid Spanning Tree Protocol (RSTP). VLAN Spanning Tree Protocol (VSTP) is an alternate protocol that allows EX Series switches to run one or more Spanning Tree Protocol (STP) or RSTP instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, VSTP improves intelligent tree spanning by defining best paths within the VLANs instead of within the entire network.



**NOTE:** EX Series switches can have a maximum of 253 VLANs on VSTP. Therefore, to have as many spanning-tree protocol VLANs as possible, use both VSTP and RSTP. RSTP will then be applied to VLANs that exceed the limit for VSTP. Because RSTP is enabled by default, you just need to additionally enable VSTP.

You can configure VSTP for an interface at the global level (for all configured VLANs) or for a specific VLAN.



**NOTE:**

- If you configure VSTP on an interface at both the global and the specific VLAN level, the interface configuration that is defined at the specific VLAN level overrides the interface configuration that is defined at the global level.
- If you specify VSTP to be configured on an interface that is not configured to belong to the VLAN (or VLANs), an error message is displayed.
- Option `vlan all` is not supported in Junos OS Release 13.2X50.

To configure VSTP:

1. Configure VSTP on a group of VLANs, on all VLANs, or on a specific VLAN, or a specific interface within a VLAN, or on a specific interface that belongs to all VLANs:

- To enable VSTP on multiple VLANs using a VLAN group:

```
[edit protocols]
```

```
user@switch# set vstp vlan-group group group-name vlan vlan-id-range
```



**NOTE:** EX Series switches can have a maximum of 253 VLANs on VSTP.



**NOTE:** The EX4300 switch does not support VLAN groups.

- To enable VSTP on all VLANs:



**NOTE:** When you configure VSTP with the `set protocol vstp vlan all` command, VLAN ID 1 is not set; it is excluded so that the configuration is compatible with Cisco PVST+. If you want VLAN ID 1 to be included in the VSTP configuration on your switch, you must set it separately with the `set protocol vstp vlan 1` command.

```
[edit protocols]
user@switch# set vlan vlan all
```



**NOTE:** EX Series switches can have a maximum of 253 VLANs on VSTP. RSTP will then be applied to VLANs that exceed the limit for VSTP. Be sure that RSTP is also enabled (the default setting) when you use the command `set vstp vlan all`. In addition to this being a recommended practice, your configuration will not commit if VSTP is enabled on a switch with more than 253 VLANs.

- To enable VSTP on a VLAN using a single VLAN ID:

```
[edit protocols]
user@switch# set vlan vlan vlan-id
```

- To enable VSTP on a VLAN using a single VLAN name:

```
[edit protocols]
user@switch# set vlan vlan vlan-name
```

- To enable VSTP on an interface at the global level:

```
[edit protocols]
user@switch# set vlan vlan all interface interface-name
```



**CAUTION:** Ensure that the interface is a member of all VLANs before you add the interface to the VSTP configuration. If the interface is not a member of all VLANs, this VSTP configuration will fail when you try to commit it.

- To enable VSTP on an interface for a specific VLAN:

```
[edit protocols]
user@switch# set vlan vlan vlan-id-or-vlan-name interface interface-name
```



**CAUTION:** Ensure that the interface is a member of the specified VLAN before you add the interface to the VSTP configuration. If the interface is not a member of the specified VLAN, this VSTP configuration will fail when you try to commit it.

2. (Optional) Enable the Address Resolution Protocol (ARP) for faster MAC address recovery only if a routed VLAN interface (RVI) is configured:

- To enable ARP on VSTP for all VLANs:

```
[edit protocols]
user@switch# set vlan vlan all arp-on-stp
```

- To enable ARP on VSTP on a VLAN using a single VLAN ID:

```
[edit protocols]
user@switch# set vlan vlan vlan-id arp-on-stp
```

- To enable ARP on VSTP on a VLAN using a single VLAN name:

```
[edit protocols]
user@switch# set vlan vlan vlan-name arp-on-stp
```

**Related  
Documentation**

- [show spanning-tree bridge on page 5035](#)
- [show spanning-tree interface on page 5045](#)
- [Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912](#)

## Unblocking an Interface That Receives BPDUs in Error (CLI Procedure)

EX Series switches use bridge protocol data unit (BPDU) protection on interfaces to prevent them from receiving BPDUs that could trigger a spanning-tree misconfiguration. If BPDUs are received on a BPDU-protected interface, the interface either shuts down or transitions to a blocking state and stops forwarding frames. In the latter scenario, after the misconfiguration that triggered the BPDUs being sent to an interface is fixed in the topology, the interface can be unblocked and returned to service.

To unblock an interface and return it to service using the CLI:

- Automatically unblock an interface by configuring a timer that expires:

```
[edit ethernet-switching-options]
user@switch# set bpdv-block disable-timeout 30
```

All interfaces on the switch will be re-enabled (unblocked) after the timer expires. However, once an interface on the switch receives a new spanning-tree protocol BPDU, the interface returns to the blocked state.

- Manually unblock an interface using the operational mode command:

```
user@switch> clear ethernet-switching bpdv-error interface ge-0/0/6.0
```

This command will only re-enable an interface but the BPDU configuration for the interface will continue to exist unless you remove the BPDU configuration explicitly.

**Related  
Documentation**

- [Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4956](#)
- [Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4961](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches on page 4914](#)

---

## Operational Tasks

- [Reverting to RSTP or VSTP from Forced IEEE 802.1D STP on page 4985](#)



## Reverting to RSTP or VSTP from Forced IEEE 802.1D STP

On MX Series routers and EX Series switches on which Rapid Spanning Tree Protocol (RSTP) or VLAN Spanning Tree Protocol (VSTP) has been forced to run as the original IEEE 802.1D Spanning Tree Protocol (STP) version, you can revert back to RSTP or VSTP.

To revert from the forced instance of the original IEEE 802.1D STP version to the originally configured RSTP or VSTP version:

1. Remove the **force-version** statement from the following RSTP or VSTP configuration:

```
user@host# delete protocols rstp force-version stp
user@host# delete protocols vstp force-version stp
```

Include this statement at the following hierarchy levels:

- [edit logical-systems *routing-instance-name* protocols rstp]
- [edit protocols rstp]
- [edit protocols vstp]
- [edit routing-instances *routing-instance-name* protocols rstp]
- [edit routing-instances *routing-instance-name* protocols vstp]

2. Revert the forced IEEE 802.1D STP to run as the configured RSTP or VSTP:

```
user@host# clear spanning-tree protocol-migration <interface interface-name>
<routing-instance routing-instance-name>
```

To revert the STP protocol globally, issue the statement without options (**clear spanning-tree protocol-migration**).

To revert the STP protocol for the specified interface only, specify the **interface *interface-name*** option.

To revert the STP protocol for a particular routing instance only, specify the **routing-instance *routing-instance-name*** option.

### Related Documentation

- [Spanning-Tree Protocols Supported](#)
- [RSTP or VSTP Forced to Run as IEEE 802.1D STP on page 4912](#)
- [Forcing RSTP or VSTP to Run as IEEE 802.1D STP \(CLI Procedure\) on page 4976](#)
- [Configuring Rapid Spanning-Tree Protocol](#)
- [Configuring VLAN Spanning-Tree Protocol](#)

## Configuration Statements

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 4986](#)
- [block \(Spanning Trees\) on page 4989](#)
- [bpdu-block on page 4990](#)

- [bpdu-block-on-edge](#) on page 4991
- [bpdu-timeout-action](#) (Spanning Trees) on page 4992
- [bridge-priority](#) (Spanning Trees) on page 4993
- [configuration-name](#) (Spanning Trees) on page 4994
- [cost](#) (Spanning Trees) on page 4995
- [disable](#) (Spanning Trees) on page 4996
- [disable-timeout](#) (Spanning Trees) on page 4997
- [edge](#) (Spanning Trees) on page 4998
- [force-version](#) (IEEE 802.1D STP) on page 4999
- [forward-delay](#) (Spanning Trees) on page 5000
- [hello-time](#) (Spanning Trees) on page 5001
- [interface](#) (BPDU Block) on page 5002
- [interface](#) (Spanning Trees) on page 5004
- [log](#) (Spanning Trees) on page 5005
- [max-age](#) (Spanning Trees) on page 5006
- [max-hops](#) (Spanning Trees) on page 5007
- [mode](#) (Spanning Trees) on page 5008
- [msti](#) (Spanning Trees) on page 5009
- [mstp](#) (Spanning Trees) on page 5010
- [no-root-port](#) (Spanning Trees) on page 5011
- [priority](#) (Spanning Trees) on page 5012
- [revision-level](#) (Spanning Trees) on page 5013
- [rstp](#) (Spanning Trees) on page 5014
- [stp](#) (Spanning Trees) on page 5016
- [traceoptions](#) (Spanning Trees) on page 5018
- [vlan](#) (Spanning Trees) on page 5021
- [vlan](#) (VSTP) on page 5023
- [vstp](#) on page 5025

## [\[edit protocols\]](#) Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the [\[edit protocols\]](#) hierarchy:

- [\[edit protocols bfd\]](#) Configuration Statement Hierarchy on EX Series Switches on page 391
- [\[edit protocols bgp\]](#) Configuration Statement Hierarchy on EX Series Switches on page 392
- [\[edit protocols connections\]](#) Configuration Statement Hierarchy on EX Series Switches on page 401

- [\[edit protocols dcbx\] Configuration Statement Hierarchy on EX Series Switches on page 403](#)
- [\[edit protocols dot1x\] Configuration Statement Hierarchy on EX Series Switches on page 404](#)
- [\[edit protocols igmp\] Configuration Statement Hierarchy on EX Series Switches on page 406](#)
- [\[edit protocols igmp-snooping\] Configuration Statement Hierarchy on EX Series Switches on page 407](#)
- [\[edit protocols isis\] Configuration Statement Hierarchy on EX Series Switches on page 408](#)
- [\[edit protocols lacp\] Configuration Statement Hierarchy on EX Series Switches on page 411](#)
- [\[edit protocols link-management\] Configuration Statement Hierarchy on EX Series Switches on page 412](#)
- [\[edit protocols lldp\] Configuration Statement Hierarchy on EX Series Switches on page 413](#)
- [\[edit protocols lldp-med\] Configuration Statement Hierarchy on EX Series Switches on page 415](#)
- [\[edit protocols mld\] Configuration Statement Hierarchy on EX Series Switches on page 416](#)
- [\[edit protocols mld-snooping\] Configuration Statement Hierarchy on EX Series Switches on page 417](#)
- [\[edit protocols mpls\] Configuration Statement Hierarchy on EX Series Switches on page 418](#)
- [\[edit protocols msdp\] Configuration Statement Hierarchy on EX Series Switches on page 429](#)
- [\[edit protocols mstp\] Configuration Statement Hierarchy on EX Series Switches on page 431](#)
- [\[edit protocols mvrp\] Configuration Statement Hierarchy on EX Series Switches on page 433](#)
- [\[edit protocols neighbor-discovery\] Configuration Statement Hierarchy on EX Series Switches on page 434](#)
- [\[edit protocols oam\] Configuration Statement Hierarchy on EX Series Switches on page 435](#)
- [\[edit protocols ospf\] Configuration Statement Hierarchy on EX Series Switches on page 438](#)
- [\[edit protocols ospf3\] Configuration Statement Hierarchy on EX Series Switches on page 441](#)
- [\[edit protocols pim\] Configuration Statement Hierarchy on EX Series Switches on page 444](#)

- [\[edit protocols rip\] Configuration Statement Hierarchy on EX Series Switches on page 447](#)
- [\[edit protocols ripng\] Configuration Statement Hierarchy on EX Series Switches on page 450](#)
- [\[edit protocols router-advertisement\] Configuration Statement Hierarchy on EX Series Switches on page 451](#)
- [\[edit protocols router-discovery\] Configuration Statement Hierarchy on EX Series Switches on page 452](#)
- [\[edit protocols rstp\] Configuration Statement Hierarchy on EX Series Switches on page 453](#)
- [\[edit protocols rsvp\] Configuration Statement Hierarchy on EX Series Switches on page 455](#)
- [\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches on page 459](#)
- [\[edit protocols stp\] Configuration Statement Hierarchy on EX Series Switches on page 460](#)
- [\[edit protocols uplink-failure-detection\] Configuration Statement Hierarchy on EX Series Switches on page 461](#)
- [\[edit protocols vrrp\] Configuration Statement Hierarchy on EX Series Switches on page 462](#)
- [\[edit protocols vstp\] Configuration Statement Hierarchy on EX Series Switches on page 463](#)

**Related  
Documentation**

- [EX Series Switch Software Features Overview](#)
- [EX Series Virtual Chassis Software Features Overview](#)

## block (Spanning Trees)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | block;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ],<br>[edit protocols <a href="#">rstp interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ],<br>[edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ],<br>[edit protocols <a href="#">vstp vlan <i>vlan-id</i> interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ]                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure loop protection on a specific interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li> <li>• <a href="#">Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches on page 4967</a></li> <li>• <a href="#">Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches on page 4916</a></li> <li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li> </ul> |

## bpdu-block

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bpdu-block {<br/>    <b>interface</b> (all   [<i>interface-name</i>]);<br/>    <b>disable-timeout</b> <i>seconds</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit ethernet-switching-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Configure bridge protocol data unit (BPDU) protection on a specified interface or on all interfaces. If the interface receives incompatible BPDUs, it is disabled.</p> <p>If the <b>disable-timeout</b> statement is included in the BPDU configuration, the interface is automatically reenabled after the timer expires. Otherwise, you must use the operational mode command <b>clear ethernet-switching bpdu-error</b> to unblock and reenable the interface.</p>                                                                                                                                                                                    |
| <b>Default</b>                  | Not enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>drop</i></li><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">clear ethernet-switching bpdu-error on page 5031</a></li><li>• <a href="#">Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4961</a></li><li>• <a href="#">Unblocking an Interface That Receives BPDUs in Error (CLI Procedure) on page 4984</a></li><li>• <a href="#">Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches on page 4914</a></li></ul> |

## bpdu-block-on-edge

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bpdu-block-on-edge;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp</a> ],<br>[edit protocols <a href="#">rstp</a> ],<br>[edit protocols <a href="#">vstp</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.<br>Statement updated in Junos OS Release 11.1 for EX Series switches to change blocking behavior to port shutdown.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch. When the <b>bpdu-block-on-edge</b> statement is configured and the interface encounters an incompatible BPDU, the interface shuts down.<br><br>If the <a href="#">disable-timeout</a> statement is included in the BPDU configuration, the interface is automatically reenabled after the timer expires. Otherwise, you must use the operational mode command <a href="#">clear ethernet-switching bpdu-error</a> to unblock and reenable the interface.                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | Not enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">clear ethernet-switching bpdu-error on page 5031</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li> <li>• <a href="#">Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4956</a></li> <li>• <a href="#">Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches on page 4914</a></li> </ul> |

## bpdu-timeout-action (Spanning Trees)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bpdu-timeout-action {<br/>    block;<br/>    log;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <pre>[edit protocols mstp interface (all   interface-name)],<br/>[edit protocols mstp interface (all   interface-name) arp-on-stp],<br/>[edit protocols rstp interface (all   interface-name)],<br/>[edit protocols rstp interface (all   interface-name) arp-on-stp],<br/>[edit protocols stp interface (all   interface-name)],<br/>[edit protocols stp interface (all   interface-name) arp-on-stp],<br/>[edit protocols vstp vlan vlan-id interface (all   interface-name)]<br/>[edit protocols vstp vlan vlan-id interface (all   interface-name) arp-on-stp]</pre>                                                                                                                                                                |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Configure the BPDU timeout action on a specific interface. You must configure at least one action (<b>log</b>, <b>block</b>, or both).</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li><li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li><li>• <a href="#">Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches on page 4967</a></li><li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li></ul> |



## bridge-priority (Spanning Trees)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bridge-priority <i>priority</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp</a> ],<br>[edit protocols <a href="#">mstp</a> <a href="#">msti</a> <i>msti-id</i> ],<br>[edit protocols <a href="#">rstp</a> ],<br>[edit protocols <a href="#">stp</a> ],<br>[edit protocols <a href="#">vstp</a> <a href="#">vlan</a> <i>vlan-id</i> ]                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.                                                                                                                                                                                                                                  |
| <b>Default</b>                  | 32,768                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><i>priority</i>—Bridge priority. It can be set only in increments of 4096.</p> <p><b>Range:</b> 0 through 61,440</p> <p><b>Default:</b> 32,768</p>                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li> <li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li> </ul> |

## configuration-name (Spanning Trees)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | configuration-name <i>configuration-name</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit protocols mstp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Specify the configuration name. The configuration name is the MSTP region name carried in the MSTP BPDUs.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li><li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li><li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li></ul> |

## cost (Spanning Trees)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>cost cost;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">mstp interface</a> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <a href="#">mstp msti msti-id interface interface-name</a> arp-on-stp],<br>[edit protocols <a href="#">mstp msti msti-id interface interface-name</a> ],<br>[edit protocols <a href="#">rstp interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">rstp interface</a> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <a href="#">vstp vlan vlan-id interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">vstp vlan vlan-id interface</a> (all   <i>interface-name</i> ) arp-on-stp] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure the link cost to control which bridge is the designated bridge and which interface is the designated interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | The link cost is determined by the link speed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>cost</b> —Link cost associated with the port.<br><b>Range:</b> 1 through 200,000,000                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Understanding RSTP for EX Series Switches on page 4906</a></li> <li>• <a href="#">Understanding STP for EX Series Switches on page 4910</a></li> <li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li> <li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                      |

## disable (Spanning Trees)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <pre>[edit protocols mpls], [edit protocols mpls interface (all   interface-name)], [edit protocols mstp], [edit protocols mstp interface (all   interface-name)], [edit protocols mstp interface (all   interface-name) arp-on-stp], [edit protocols mstp msti msti-id vlan (vlan-id   vlan-name) interface (all   interface-name)], [edit protocols mstp msti msti-id vlan (vlan-id   vlan-name) interface interface-name   arp-on-stp], [edit protocols rstp], [edit protocols rstp interface (all   interface-name)], [edit protocols rstp interface (all   interface-name) arp-on-stp], [edit protocols stp], [edit protocols stp interface (all   interface-name)], [edit protocols stp interface (all   interface-name) arp-on-stp], [edit protocols vstp], [edit protocols vstp vlan vlan-id interface (all   interface-name)] [edit protocols vstp vlan vlan-id interface (all   interface-name) arp-on-stp]</pre> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Disable STP, MPLS, MSTP, RSTP, or VSTP on the switch or on a specific interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li><li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li><li>• <a href="#">Understanding STP for EX Series Switches on page 4910</a></li><li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li></ul>                                                                                                                                                                                                                                                                                               |

## disable-timeout (Spanning Trees)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>disable-timeout <i>timeout</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit ethernet-switching-options <a href="#">bpdu-block</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | For interfaces configured for BPDU protection, specify the amount of time an interface receiving BPDUs is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | The disable timeout is not enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b><i>timeout</i></b> —Amount of time, in seconds, the interface receiving BPDUs is disabled. Once the timeout expires, the interface is brought back into service.</p> <p><b>Range:</b> 10 through 3600 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li> <li>• <a href="#">Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4961</a></li> <li>• <a href="#">Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches on page 4914</a></li> </ul> |

## edge (Spanning Trees)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | edge;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | [edit protocols <b>mstp</b> <b>interface</b> (all   <i>interface-name</i> )],<br>[edit protocols <b>mstp</b> <b>interface</b> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <b>mstp</b> <b>msti</b> <i>msti-id</i> <b>interface</b> <i>interface-name</i> ],<br>[edit protocols <b>mstp</b> <b>msti</b> <i>msti-id</i> <b>interface</b> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <b>rstp</b> <b>interface</b> (all   <i>interface-name</i> )],<br>[edit protocols <b>rstp</b> <b>interface</b> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <b>vstp</b> <b>vlan</b> <i>vlan-id</i> <b>interface</b> (all   <i>interface-name</i> )],<br>[edit protocols <b>vstp</b> <b>vlan</b> <i>vlan-id</i> <b>interface</b> (all   <i>interface-name</i> ) arp-on-stp] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>         | For Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure interfaces as edge interfaces. Edge interfaces immediately transition to a forwarding state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



**NOTE:** Although the edge configuration statement appears in the [edit protocols stp interface (all | *interface-name*)] or [edit protocols rstp force-version stp interface (all | *interface-name*)] hierarchy on the switch, this statement has no effect on the switch operation if you configure it.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | Edge interfaces are not enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li> <li>• <a href="#">Understanding STP for EX Series Switches on page 4910</a></li> <li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li> </ul> |

## force-version (IEEE 802.1D STP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | force-version stp;                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols (rstp   vstp)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (rstp   vstp)],<br>[edit protocols (rstp   vstp)],<br>[edit routing-instances <i>routing-instance-name</i> protocols (rstp   vstp)]                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Support for logical systems added in Junos OS Release 9.6.                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Force the spanning-tree protocol version to be the original IEEE 802.1D STP.                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Spanning-Tree Protocols Supported</i></li> <li>• <a href="#">RSTP or VSTP Forced to Run as IEEE 802.1D STP on page 4912</a></li> <li>• <a href="#">Forcing RSTP or VSTP to Run as IEEE 802.1D STP (CLI Procedure) on page 4976</a></li> <li>• <a href="#">Reverting to RSTP or VSTP from Forced IEEE 802.1D STP on page 4985</a></li> </ul> |

## forward-delay (Spanning Trees)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>forward-delay seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp</a> ],<br>[edit protocols <a href="#">rstp</a> ],<br>[edit protocols <a href="#">stp</a> ],<br>[edit protocols <a href="#">vstp</a> <a href="#">vlan</a> <i>vlan-id</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify how long a bridge interface remains in the listening and learning states before transitioning to the forwarding state.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | 15 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>seconds</b> —Number of seconds the bridge interface remains in the listening and learning states.<br><b>Range:</b> 4 through 30 seconds<br><b>Default:</b> 15 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li><li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li><li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li><li>• <a href="#">Understanding STP for EX Series Switches on page 4910</a></li><li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li></ul> |



## hello-time (Spanning Trees)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | hello-time <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols <i>mstp</i> ],<br>[edit protocols <i>rstp</i> ],<br>[edit protocols <i>stp</i> ],<br>[edit protocols <i>vstp</i> <i>vlan</i> <i>vlan-id</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the time interval at which the root bridge transmits configuration BPDUs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>                  | 2 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>seconds</b> —Number of seconds between transmissions of configuration BPDUs.<br><b>Range:</b> 1 through 10 seconds<br><b>Default:</b> 2 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li> <li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li> <li>• <a href="#">Understanding STP for EX Series Switches on page 4910</a></li> <li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li> </ul> |

## interface (BPDU Block)

---

**Syntax**    interface (all | [*interface-name*]);

**Hierarchy Level**    [edit ethernet-switching-options **bpdu-block**]

**Release Information**    Statement introduced in Junos OS Release 9.1 for EX Series switches.

**Description**    Apply Bridge Protocol Data Unit (BPDU) protection on all interfaces or on one or more specified interfaces.

Spanning Tree Protocol (STP), Rapid Spanning Tree protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP) provide Layer 2 loop prevention for EX Series switches. The spanning-tree protocols use BPDU frames to communicate. Through their exchange, spanning-tree topologies determine which interfaces block traffic and which interfaces become root ports and forward traffic. User bridge applications running on a PC can also generate BPDUs. If these BPDUs are picked up by STP applications running on the switch, they can trigger STP miscalculations that can lead to network outages.

To block outside BPDUs from reaching a switch interface connected to devices that are not part of a spanning-tree topology, configure BPDU protection on interfaces in the topology.



**CAUTION:** When configuring BPDU protection on an interface without spanning trees connected to a switch with spanning trees, be careful that you do not configure BPDU protection on all interfaces. Doing so could prevent BPDUs being received on switch interfaces (such as a trunk interface) that you intended to have receive BPDUs from a switch with spanning trees.

---



**NOTE:** Interfaces that are configured as edge interfaces can transition to a forwarding state immediately because they cannot create network loops. As edge ports are connected to end devices, it is imperative that you configure BPDU protection on edge ports to protect the switch from outside BPDUs. When BPDU protection is enabled on an edge interface, the interface shuts down on encountering an outside BPDU thereby preventing any traffic from passing through the interface. To enable BPDU protection on an edge interface, see *Configuring BPDU Protection on an Interface (CLI Procedure)*

---

**Options**    all—All interfaces.

[*interface-name*]  
[*interface-name*]  
—One or more Ethernet interface names.

The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li><li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li><li>• <a href="#">Example: Configuring BPDU Protection on Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4961</a></li></ul> |

## interface (Spanning Trees)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>interface <i>interface-name</i> {<br/>    arp-on-stp;<br/>    bpdu-timeout-action<br/>        block;<br/>        log;<br/>    cost <i>cost</i>;<br/>    disable;<br/>    edge;<br/>    mode <i>mode</i>;<br/>    no-root-port;<br/>    priority <i>priority</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Hierarchy Level          | [edit protocols <a href="#">mstp</a> ],<br>[edit protocols <a href="#">mstp msti msti-id</a> ],<br>[edit protocols <a href="#">rstp</a> ],<br>[edit protocols <a href="#">stp</a> ],<br>[edit protocols <a href="#">vstp vlan</a> (all   <i>vlan-id</i>   <i>vlan-name</i> )]                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Release Information      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Description              | <p>For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure an interface.</p> <p>The <b>edge</b>, <b>mode</b>, and <b>no-root-port</b> options are not available at the [edit protocols <a href="#">mstp msti msti-id</a>] hierarchy level.</p>                                                                                                                                                                                                                                                                                                                                                      |
| Options                  | <p><b><i>interface-name</i></b>—Name of an interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li><li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP</a></li><li>• <a href="#">Configuring VSTP (CLI Procedure) on page 4982</a></li></ul> |

- [show spanning-tree bridge on page 5040](#)

## log (Spanning Trees)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | log;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ],<br>[edit protocols <a href="#">rstp interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ],<br>[edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ],<br>[edit protocols <a href="#">vstp vlan <i>vlan-id</i> interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ]                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | For interfaces configured for loop protection, configure the software to generate a message to be sent to the system log file <code>/var/log/messages</code> to record the loop-protection event.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li> <li>• <a href="#">Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on EX Series Switches on page 4967</a></li> <li>• <a href="#">Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on EX Series Switches on page 4916</a></li> <li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li> </ul> |

## max-age (Spanning Trees)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>max-age <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp</a> ],<br>[edit protocols <a href="#">rstp</a> ],<br>[edit protocols <a href="#">stp</a> ],<br>[edit protocols <a href="#">vstp</a> <a href="#">vlan</a> <i>vlan-id</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the maximum age of received protocol BPDUs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>                  | 20 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b><i>seconds</i></b> —The maximum age of received protocol BPDUs.<br><b>Range:</b> 6 through 40 seconds<br><b>Default:</b> 20 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li><li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li><li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li><li>• <a href="#">Understanding STP for EX Series Switches on page 4910</a></li><li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li></ul> |

---

## max-hops (Spanning Trees)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>max-hops hops;</code>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp</a> ]                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | For Multiple Spanning Tree Protocol (MSTP), configure the maximum number of hops a BPDU can be forwarded in the MSTP region.                                                                                                                                                                                                                                                        |
| <b>Default</b>                  | 20 hops                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>hops</b> — Number of hops the BPDU can be forwarded.<br><b>Range:</b> 1 through 255 hops<br><b>Default:</b> 20 hops                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li><li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li></ul> |

## mode (Spanning Trees)

|                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                       | <code>mode mode;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>                                                                                                                                                              | <p>[edit protocols <b>mstp</b> <b>interface</b> (all   <i>interface-name</i>)],</p> <p>[edit protocols <b>mstp</b> <b>interface</b> (all   <i>interface-name</i>) arp-on-stp],</p> <p>[edit protocols <b>mstp</b> <b>msti</b> <i>msti-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols <b>mstp</b> <b>msti</b> <i>msti-id</i> <b>interface</b> <i>interface-name</i>) arp-on-stp],</p> <p>[edit protocols <b>rstp</b> <b>interface</b> (all   <i>interface-name</i>)],</p> <p>[edit protocols <b>rstp</b> <b>interface</b> (all   <i>interface-name</i>) arp-on-stp],</p> <p>[edit protocols <b>stp</b> <b>interface</b> (all   (all   <i>interface-name</i>))],</p> <p>[edit protocols <b>stp</b> <b>interface</b> (all   <i>interface-name</i>) arp-on-stp],</p> <p>[edit protocols <b>vstp</b> <b>vlan</b> <i>vlan-id</i> <b>interface</b> (all   <i>interface-name</i>)],</p> <p>[edit protocols <b>vstp</b> <b>vlan</b> <i>vlan-id</i> <b>interface</b> (all   <i>interface-name</i>) arp-on-stp]</p> |
| <b>Release Information</b>                                                                                                                                                          | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                                                                                                                                                                  | For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure the link mode to identify point-to-point links.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                                                                                                                                                                      | For a full-duplex link, the default link mode is <b>point-to-point</b> . For a half-duplex link, the default link mode is <b>shared</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <div>  <b>NOTE:</b> For EX4300 switches, the interfaces operate in full-duplex mode only. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                                                                                                                                                      | <p><b>mode</b>—Link mode:</p> <ul style="list-style-type: none"> <li><b>point-to-point</b>—Link is point to point.</li> <li><b>shared</b>—Link is shared media.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                                                                                                                                                     | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>                                                                                                                                                        | <ul style="list-style-type: none"> <li><a href="#">show spanning-tree bridge on page 5035</a></li> <li><a href="#">show spanning-tree interface on page 5045</a></li> <li><a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li><a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li> <li><a href="#">Understanding STP for EX Series Switches on page 4910</a></li> <li><a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                         |



## msti (Spanning Trees)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> msti <i>msti-id</i> {   vlan (<i>vlan-id</i>   <i>vlan-name</i>);   interface <i>interface-name</i> {     disable;     cost <i>cost</i>;     priority <i>priority</i>;   } } </pre>                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp</a> ]                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure the Multiple Spanning Tree Instance (MSTI) identifier for Multiple Spanning Tree Protocol (MSTP). MSTI IDs are local to each region, so you can reuse the same MSTI ID in different regions.                                                                                                                                                                                   |
| <b>Default</b>                  | MSTI is disabled.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><i>msti-id</i> —MSTI identifier.</p> <p><b>Range:</b> 1 through 4094. The Common Instance Spanning Tree (CIST) is always MSTI 0.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li> </ul> |

## mstp (Spanning Trees)

---

```
Syntax  mstp {  
        bpd-block-on-edge;  
        bridge-priority priority;  
        configuration-name name;  
        disable;  
        forward-delay seconds;  
        hello-time seconds;  
        interface ( all | interface-name ){  
            arp-on-stp;  
            bpd-timeout-action {  
                block;  
                log;  
            }  
            cost cost;  
            disable;  
            edge;  
            mode mode;  
            no-root-port;  
            priority priority;  
        }  
        max-age seconds;  
        max-hops hops;  
        msti msti-id {  
            vlan (vlan-id | vlan-name);  
            interface interface-name {  
                disable;  
                cost cost;  
                priority priority;  
            }  
        }  
        revision-level revision-level;  
        traceoptions {  
            file filename <files number > <size size > <no-stamp | world-readable |  
                no-world-readable>;  
            flag flag;  
        }  
    }
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure Multiple Spanning Tree Protocol (MSTP). MSTP is defined in the IEEE 802.1Q-2003 specification and is used to create a loop-free topology in networks with multiple spanning tree regions.

The statements are explained separately.

**Default** MSTP is disabled.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li> </ul> |

## no-root-port (Spanning Trees)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-root-port;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">mstp interface</a> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <a href="#">rstp interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">rstp interface</a> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <a href="#">vstp vlan vlan-id interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">vstp vlan vlan-id interface</a> (all   <i>interface-name</i> ) arp-on-stp] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure an interface to be a spanning-tree designated port. If the bridge receives superior STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on EX Series Switches on page 4971</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li> <li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li> </ul>                                               |

## priority (Spanning Trees)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>priority</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp</a> <a href="#">interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">mstp</a> ; <a href="#">interface</a> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <a href="#">mstp</a> <a href="#">msti</a> <i>msti-id</i> <a href="#">interface</a> <i>interface-name</i> ],<br>[edit protocols <a href="#">mstp</a> <a href="#">msti</a> <i>msti-id</i> <a href="#">interface</a> <i>interface-name</i> arp-on-stp],<br>[edit protocols <a href="#">rstp</a> <a href="#">interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">rstp</a> <a href="#">interface</a> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <a href="#">stp</a> <a href="#">interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">stp</a> <a href="#">interface</a> (all   <i>interface-name</i> ) arp-on-stp],<br>[edit protocols <a href="#">vstp</a> <a href="#">vlan</a> <i>vlan-id</i> <a href="#">interface</a> (all   <i>interface-name</i> )],<br>[edit protocols <a href="#">vstp</a> <a href="#">vlan</a> <i>vlan-id</i> <a href="#">interface</a> (all   <i>interface-name</i> ) arp-on-stp] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the interface priority to control which interface is elected as the root port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | The default value is 128.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>priority</b> —Interface priority. The interface priority must be set in increments of 16.<br><b>Range:</b> 0 through 240                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li><li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li><li>• <a href="#">Understanding STP for EX Series Switches on page 4910</a></li><li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## revision-level (Spanning Trees)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>revision-level <i>revision-level</i>;</code>                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">mstp</a> ]                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | For Multiple Spanning Tree Protocol (MSTP), set the revision number of the MSTP configuration.                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | The revision level is disabled.                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>revision-level</i> —Revision number of the MSTP region configuration.<br><b>Range:</b> 0 through 65535                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li> </ul> |

## rstp (Spanning Trees)

```
Syntax  rstp {
        bpd-block-on-edge;
        bridge-priority priority;
        disable;
        forward-delay seconds;
        hello-time seconds;
        interface (all | interface-name) {
            arp-on-stp;
            bpd-timeout-action {
                block;
                log;
            }
            cost cost;
            disable;
            edge;
            mode mode;
            no-root-port;
            priority priority;
        }
        max-age seconds;
        traceoptions {
            file filename <files number > <size size > <no-stamp | no-world-readable |
            world-readable>;
            flag flag;
        }
    }
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure Rapid Spanning Tree Protocol (RSTP). RSTP is defined in the IEEE 802.1D-2004 specification and is used to prevent loops in Layer 2 networks, which results in shorter convergence times than those provided by basic Spanning Tree Protocol (STP).

VSTP and RSTP can be configured concurrently. You can selectively configure up to 253 VLANs using VSTP; the remaining VLANs will be configured using RSTP. VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on the switch. See [“Configuring VSTP \(CLI Procedure\)” on page 4982](#) for more information on configuring VSTP and RSTP concurrently.



**BEST PRACTICE:** Configure RSTP when you configure VSTP. RSTP overhead is minimal and this configuration ensures that a spanning-tree protocol is running on all VLANs on your switch, even when your switch is supporting more than 253 VLANs.

The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | RSTP is enabled on all Ethernet switching interfaces.                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919</a></li><li>• <a href="#">Understanding RSTP for EX Series Switches on page 4906</a></li></ul> |

## stp (Spanning Trees)

---

**Syntax**    `stp {  
    bpdu-block-on-edge ;  
    bridge-priority priority;  
    disable;  
    forward-delay seconds;  
    hello-time seconds;  
    interface (all | interface-name) {  
        arp-on-stp;  
        bpdu-timeout-action {  
            block;  
            log;  
        }  
        cost cost;  
        disable;  
        mode mode;  
        no-root-port;  
        priority priority;  
    }  
    max-age seconds;  
    traceoptions {  
        file filename <files number > <size size > <no-stamp | world-readable |  
        no-world-readable>;  
        flag flag;  
    }  
}`

**Hierarchy Level**    [edit protocols],  
                          [edit protocols rstp force-version]

**Release Information**    Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description**    Configure Spanning Tree Protocol (STP). When you explicitly configure STP, the switches use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP (defined in the IEEE 802.1D 1998 specification).



**NOTE:** Although the `edge` configuration statement appears in the CLI on the switch, this statement has no effect on the switch operation if you configure it.

---

The remaining statements are explained separately.

**Default**    STP is disabled. By default, RSTP is enabled on all Ethernet switching ports.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.



- Related Documentation**
- [show spanning-tree bridge on page 5035](#)
  - [show spanning-tree interface on page 5045](#)
  - [Example: Configuring BPDU Protection on Edge Interfaces to Prevent STP Miscalculations on EX Series Switches on page 4956](#)
  - [Configuring STP \(CLI Procedure\)](#)
  - [Understanding STP for EX Series Switches on page 4910](#)

## traceoptions (Spanning Trees)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>name</i> &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt; &lt;no-stamp&gt;<br/>    &lt;(world-readable   no-world-readable)&gt;;<br/>    flag <i>flag</i> &lt;<i>flag-modifier</i>&gt; &lt;disable&gt;;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | <pre>[edit protocols mpls],<br/>[edit protocols <b>mstp</b>],<br/>[edit protocols <b>mvrp</b>],<br/>[edit protocols <b>rstp</b>],<br/>[edit protocols <b>stp</b>],<br/>[edit protocols <b>vstp</b> vlan <i>vlan-id</i>]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | Set protocol-level tracing options for MPLS, MVRP, STP, RSTP, MSTP, and VSTP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default</b>             | Traceoptions is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>name</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place STP tracing output in the file <code>/var/log/stp-log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>trace-file</i></b> reaches its maximum size, it is renamed <b><i>trace-file .0</i></b>, then <b><i>trace-file .1</i></b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the size option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 1 trace file</p> <p><b><i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b><i>flag</i></b> statements:</p> <ul style="list-style-type: none"><li>• <b>all</b>—Trace all operations.</li><li>• <b>all-failures</b>—Trace all failure conditions.</li><li>• <b>bpd</b><b>u</b>—Trace BPDU reception and transmission. Note that you must also use <b>port-transmit-state-machine</b> in order to log transmit operations.</li><li>• <b>bridge-detection-state-machine</b>—Trace the bridge detection state machine.</li></ul> |

- **error**—Trace all failure conditions.
- **events**—Trace events of the protocol state machine.
- **pdu**—Trace PDUs that were received and sent.
- **port-information-state-machine**—Trace the port information state machine.
- **port-migration-state-machine**—Trace the port migration state machine.
- **port-receive-state-machine**—Trace the port receive state machine.
- **port-role-select-state-machine**—Trace the port role selection state machine.
- **port-role-transit-state-machine**—Trace the port role transit state machine.
- **port-state-transit-state-machine**—Trace the port state transit state machine.
- **port-transmit-state-machine**—Trace the port transmit state machine.
- **ppmd**—Trace the state and events for the ppm process.
- **socket**—Trace socket activity.
- **state-machine**—Trace state machine information.
- **state-machine-variables**—Trace when the state machine variables change.
- **timers**—Trace protocol timers.
- **topology-change-state-machine**—Trace the topology change state machine.

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size** *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file .0*. When the *trace-file* again reaches its maximum size, *trace-file .0* is renamed *trace-file .1* and *trace-file* is renamed *trace-file .0*. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|

**Related Documentation**

- [show spanning-tree bridge on page 5035](#)
- [show spanning-tree interface on page 5045](#)
- [Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936](#)
- [Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919](#)
- [Understanding MSTP for EX Series Switches on page 4904](#)
- [Understanding RSTP for EX Series Switches on page 4906](#)
- [Understanding STP for EX Series Switches on page 4910](#)
- [Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912](#)
- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on EX Series Switches on page 2264](#)

## vlan (Spanning Trees)

**Syntax** `vlan (all | vlan-id | vlan-name) {`  
     `bridge-priority priority;`  
     `forward-delay seconds;`  
     `hello-time seconds;`  
     `interface interface-name {`  
         `bpdu-timeout-action {`  
             `block;`  
             `log;`  
         `}`  
     `cost cost;`  
     `disable;`  
     `edge;`  
     `mode mode;`  
     `no-root-port;`  
     `priority priority;`  
     `}`  
     `max-age seconds;`  
     `traceoptions {`  
         `file filename <files number > <size size > <no-stamp | world-readable |`  
         `no-world-readable>;`  
         `flag flag;`  
     `}`  
`}`

**Hierarchy Level** [edit protocols `mstp msti msti-id`]  
                     [edit protocols `vstp`]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                           Statement updated with enhanced ? (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.

**Description** Configure the VLANs for a Multiple Spanning Tree Instance (MSTI) or VSTP instance.



**NOTE:** When you configure VSTP with the `set protocol vstp vlan all` command, `vlan-id 1` is excluded to be compatible with Cisco PVST+. If you want `vlan-id 1` to be included in VSTP, you must set it separately with the `set protocol vstp vlan 1` command.



**TIP:** To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type `?` after `vlan` or `vlangs` in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

**Default** Not enabled.

**Options**    *vlan-id*—Numeric VLAN identifier.

*vlan-name*—Name of the VLAN.

              The remaining statements are explained separately.

**Required Privilege**    routing—To view this statement in the configuration.

**Level**                routing-control—To add this statement to the configuration.

**Related**            • [Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches](#)  
**Documentation**    [on page 4936](#)

                      • [Understanding MSTP for EX Series Switches on page 4904](#)

## vlan (VSTP)

**Syntax** `vlan` (all | *vlan-id* | *vlan-name*) {  
     `bridge-priority` *priority*;  
     `forward-delay` *seconds*;  
     `hello-time` *seconds*;  
     `interface` (all | *interface-name*) {  
         `bpdu-timeout-action` {  
             `block`;  
             `log`;  
         }  
         `cost` *cost*;  
         `disable`;  
         `edge`;  
         `mode` *mode*;  
         `no-root-port`;  
         `priority` *priority*;  
     }  
     `max-age` *seconds*;  
     `traceoptions` {  
         `file` *filename* <files *number* > <size *size* > <no-stamp | world-readable |  
         no-world-readable>;  
         `flag` *flag*;  
     }  
 }

**Hierarchy Level** [edit protocols **vstp**]

**Release Information** Statement introduced in Junos OS Release 9.4 for EX Series switches.  
 Statement updated with enhanced ? (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.  
 Option **all** introduced in Junos OS Release 10.0 for EX Series switches.

**Description** Configure VSTP VLAN parameters.



**NOTE:** Option **all** is not supported in Junos OS Release 13.2X50.



**TIP:** To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after `vlan` or `vlands` in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

**Options** **all**—(Not supported on EX4300) All VLANs.

***vlan-id***—Numeric VLAN identifier.

***vlan-name***—Name of the VLAN.

The remaining statements are explained separately.

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |                                                                                                                                                  |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li></ul> |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|



## vstp

```
Syntax  vstp {
        bpd-block-on-edge;
        disable;
        force-version stp;
        vlan (all | vlan-id | vlan-name) {
            bridge-priority priority;
            forward-delay seconds;
            hello-time seconds;
            interface (all | interface-name) {
                arp-on-stp;
                bpd-timeout-action {
                    block;
                    log;
                }
                cost cost;
                disable;
                edge;
                mode mode;
                no-root-port;
                priority priority;
            }
            max-age seconds;
            traceoptions {
                file filename <files number > <size size> <no-stamp | no-world-readable |
                world-readable>;
                flag flag;
            }
        }
    }
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 9.4 for EX Series switches.

**Description** Configure VLAN Spanning Tree Protocol (VSTP). VSTP is used to prevent loops in Layer 2 networks on a per-VLAN basis.

You can have a maximum of 253 VSTP VLANs per switch.

If the number of VLANs on your switch exceeds the VSTP VLAN limit, you must use the **vlan (Spanning Trees)** statement to specify which VLANs or VLAN groups use VSTP. You also cannot use the **vlan all** option to configure VSTP when your switch has more than the maximum allowed VSTP VLANs. To ensure all VLANs are running a spanning-tree protocol, run RSTP for networks with large numbers of VLANs.



**NOTE:** When you configure VSTP with the `set protocol vstp vlan all` command, VLAN ID 1 is not set; it is excluded so that the configuration is compatible with Cisco PVST+. If you want VLAN ID 1 to be included in the VSTP configuration

on your switch, you must set it separately with the `set protocol vstp vlan 1` command.



**NOTE:** Option `vlan all` is not supported in Junos OS Release 13.2X50.



**BEST PRACTICE:** Configure RSTP when you configure VSTP. RSTP overhead is minimal and this configuration ensures that some spanning tree protocol is running on all VLANs on your switch, even when your switch has more than the maximum number of allowed VSTP VLANs.

The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | VSTP is not enabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 5035</a></li><li>• <a href="#">show spanning-tree interface on page 5045</a></li><li>• <a href="#">Configuring VSTP (CLI Procedure) on page 4982</a></li><li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li><li>• <a href="#">Example: Configuring VSTP on QFX Series Switches and EX4600 Switches</a></li></ul> |

# Administration

- [Routine Monitoring on page 5027](#)
- [Operational Commands on page 5029](#)

## Routine Monitoring

- [Monitoring Spanning-Tree Protocols on page 5027](#)

### Monitoring Spanning-Tree Protocols

Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring feature to view status and information about the spanning-tree protocol parameters on your EX Series switch.

Action

To display spanning-tree protocol parameter details in the J-Web interface, select **Monitor** > **Switching** > **STP**.

To display spanning-tree protocol parameter details in the CLI, enter the following commands:

- **show spanning-tree interface**
- **show spanning-tree bridge**

Meaning

[Table 555 on page 5027](#) summarizes the spanning-tree protocol parameters.

Table 555: Summary of Spanning-Tree Protocols Output Fields

| Field                                                         | Values                                                          |
|---------------------------------------------------------------|-----------------------------------------------------------------|
| Bridge Parameters                                             |                                                                 |
| Routing instance name                                         | Displays bridge information for the specified routing instance. |
| <b>NOTE:</b> The option is supported only on EX4300 switches. |                                                                 |
| Context ID                                                    | An internally generated identifier.                             |

Table 555: Summary of Spanning-Tree Protocols Output Fields (*continued*)

| Field                                                          | Values                                                                                                                                                 |
|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enabled Protocol                                               | Spanning-tree protocol type enabled.                                                                                                                   |
| Root ID                                                        | Bridge ID of the elected spanning-tree root bridge.<br><br>The bridge ID consists of a configurable bridge priority and the MAC address of the bridge. |
| Root cost                                                      | Calculated cost to reach the root bridge from the bridge where the command is entered.                                                                 |
| Root port                                                      | Interface that is the current elected root port for this bridge.                                                                                       |
| Bridge ID                                                      | Locally configured bridge ID.                                                                                                                          |
| Hello time                                                     | The time for which the bridge interface remains in the listening or learning state.                                                                    |
| Forward delay                                                  | The time for which the bridge interface remains in the listening or learning state before transitioning to the forwarding state.                       |
| Extended System ID                                             | The system ID.                                                                                                                                         |
| Inter Instance ID                                              | An internally generated instance identifier.                                                                                                           |
| Maximum age                                                    | Maximum age of received bridge protocol data units (BPDUs).                                                                                            |
| Number of topology changes                                     | Total number of spanning-tree protocol topology changes detected since the switch last booted.                                                         |
| Time since last topology change                                | Number of seconds elapsed since the last topology change.                                                                                              |
| <b>NOTE:</b> This option is supported only on EX4300 switches. |                                                                                                                                                        |
| Spanning Tree Interface Details                                |                                                                                                                                                        |
| Interface Name                                                 | Interface configured to participate in the spanning-tree protocol instance.                                                                            |
| Port ID                                                        | Logical interface identifier configured to participate in the spanning-tree protocol instance.                                                         |
| Designated Port ID                                             | Port ID of the designated port for the LAN segment to which the interface is attached.                                                                 |
| Designated Bridge ID                                           | ID of the designated bridge to which the interface is attached.                                                                                        |
| Port Cost                                                      | Configured cost for the interface.                                                                                                                     |

Table 555: Summary of Spanning-Tree Protocols Output Fields (*continued*)

| Field                                 | Values                                                                                                                                                                                         |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port State                            | Spanning-tree protocol port state: <ul style="list-style-type: none"> <li>• Forwarding (FWD)</li> <li>• Blocking (BLK)</li> <li>• Listening</li> <li>• Learning</li> <li>• Disabled</li> </ul> |
| Role                                  | MSTP or RSTP port role, Designated (DESG), backup (BKUP), alternate (ALT), or root.                                                                                                            |
| Spanning Tree Statistics of Interface |                                                                                                                                                                                                |
| Interface                             | Interface for which statistics is being displayed.                                                                                                                                             |
| BPDUs Sent                            | Total number of BPDUs sent.                                                                                                                                                                    |
| BPDUs Received                        | Total number of BPDUs received.                                                                                                                                                                |
| Next BPDU Transmission                | Number of seconds until the next BPDU is scheduled to be sent.                                                                                                                                 |

- Related Documentation**
- [show spanning-tree interface on page 5045](#)
  - [show spanning-tree bridge on page 5035](#)
  - [Configuring Spanning-Tree Protocols \(J-Web Procedure\) on page 4977](#)
  - [Configuring STP \(CLI Procedure\)](#)
  - [Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936](#)
  - [Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches on page 4919](#)

## Operational Commands

- [clear ethernet-switching bpd-error](#)
- [clear spanning-tree protocol-migration](#)
- [clear spanning-tree statistics](#)
- [clear spanning-tree statistics](#)
- [show spanning-tree bridge](#)
- [show spanning-tree bridge](#)
- [show spanning-tree interface](#)
- [show spanning-tree interface](#)

- [show spanning-tree mstp configuration](#)
- [show spanning-tree mstp configuration](#)
- [show spanning-tree statistics](#)
- [show spanning-tree statistics](#)

## clear ethernet-switching bpdu-error

|                                 |                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear ethernet-switching bpdu-error interface <i>interface-name</i></code>                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.1 for EX Series switches. Command updated in Junos OS Release 11.1 for EX Series switches—a BPDU error shuts down the interface and this command brings the interface back up.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                         |
| <b>Description</b>              | Clear bridge protocol data unit (BPDU) errors from an interface and bring up the interface.                                                                                                                                                                                                                        |
| <b>Options</b>                  | <i>interface-name</i> —Clear BPDU errors on the specified interface.                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree interface on page 5049</a></li> <li>• <a href="#">Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches on page 4914</a></li> <li>• <a href="#">Understanding BPDU Protection for STP, RSTP, and MSTP</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">clear ethernet-switching bpdu-error interface on page 5031</a>                                                                                                                                                                                                                                         |

### Sample Output

#### clear ethernet-switching bpdu-error interface

```
user@switch> clear ethernet-switching bpdu-error interface xe-0/0/1.0
```

## clear spanning-tree protocol-migration

---

|                                 |                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear spanning-tree protocol-migration</code><br><code>&lt;interface <i>interface-name</i>&gt;</code><br><code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0.                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Revert from the original IEEE 802.1D Spanning Tree Protocol (STP) back to the Rapid Spanning Tree Protocol after the <b>force-version</b> statement has been removed from the configuration.                                                                                                                                           |
| <b>Options</b>                  | <b>none</b> —Reset the STP protocol for all interfaces and all routing instances.<br><br><b>interface <i>interface-name</i></b> —(Optional) Reset the STP protocol for the specified interface only.<br><br><b>routing-instance <i>routing-instance-name</i></b> —(Optional) Reset the STP protocol for a particular routing instance. |
| <b>Additional Information</b>   | For information about the <b>force-version</b> statement, see the <i>Junos Routing Protocols Configuration Guide</i> . and <a href="#">“Forcing RSTP or VSTP to Run as IEEE 802.1D STP (CLI Procedure)”</a> on page 4976.                                                                                                              |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### clear spanning-tree protocol-migration

```
user@host> clear spanning-tree protocol-migration
```



## clear spanning-tree statistics

---

|                                 |                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear spanning-tree statistics<br><interface <i>interface-name</i> unit <i>logical-unit-number</i> >;                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                       |
| <b>Description</b>              | Reset STP statistics for the all interfaces or a specified interface.                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>none</b>—Reset STP counters for all interfaces.</p> <p><b><i>interface-name</i></b> —(Optional) The name of the interface for which statistics should be reset.</p> <p><b><i>logical-unit-number</i></b> —(Optional) The logical unit number of the interface.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Understanding STP for EX Series Switches on page 4910</a></li> </ul>         |
| <b>List of Sample Output</b>    | <a href="#">clear spanning-tree statistics on page 5033</a>                                                                                                                                                                                                              |
| <b>Output Fields</b>            | This command produces no output.                                                                                                                                                                                                                                         |


## Sample Output

### clear spanning-tree statistics

```
user@switch> clear spanning-tree statistics
```

## clear spanning-tree statistics

---

|                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                                                                                                                                                   | <a href="#">Syntax on page 5034</a><br><a href="#">Syntax (EX Series Switches and the QFX Series) on page 5034</a>                                                                                                                                                                                                                        |
| <b>Syntax</b>                                                                                                                                                                           | <code>clear spanning-tree statistics</code><br><code>&lt;interface <i>interface-name</i>&gt;</code><br><code>&lt;logical-system <i>logical-system-name</i>&gt;</code>                                                                                                                                                                     |
| <b>Syntax (EX Series Switches and the QFX Series)</b>                                                                                                                                   | <code>clear spanning-tree statistics</code><br><code>&lt;interface <i>interface-name</i>&gt;</code>                                                                                                                                                                                                                                       |
| <b>Release Information</b>                                                                                                                                                              | Command introduced in Junos OS Release 8.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                      | Clear Spanning Tree Protocol statistics.                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                                                                                                                                                          | <code>none</code> —Reset STP counters for all interfaces for all routing instances.<br><br><code>interface <i>interface-name</i></code> —(Optional) Clear STP statistics for the specified interface only.<br><br><code>logical-system <i>logical-system-name</i></code> —(Optional) Clear STP statistics on a particular logical system. |
| <div> <b>NOTE:</b> The <code>logical-system</code> option is not available on QFabric systems.</div> |                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b>                                                                                                                                                         | clear                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>                                                                                                                                                            | <ul style="list-style-type: none"><li><a href="#">show spanning-tree statistics on page 5061</a></li></ul>                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>                                                                                                                                                            | <a href="#">clear stp statistics on page 5034</a>                                                                                                                                                                                                                                                                                         |

### Sample Output

#### clear stp statistics

```
user@host> clear stp statistics
```

## show spanning-tree bridge

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show spanning-tree bridge<br><brief   detail><br><msti <i>msti-id</i> ><br><vlan <i>vlan-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Display the configured or calculated spanning-tree protocol (can be either STP, RSTP, MSTP, or VSTP) parameters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>none</b>—(Optional) Display brief spanning-tree protocol bridge information for all Multiple Spanning Tree Instances (MSTIs).</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>msti <i>msti-id</i></b>—(Optional) Display spanning-tree protocol bridge information for the specified MSTI or Common and Internal Spanning Tree (CIST). Specify <b>0</b> for CIST. Specify a value from <b>1</b> through <b>4094</b> for an MSTI.</p> <p><b>vlan <i>vlan-id</i></b>—(Optional) Display spanning-tree protocol bridge information for the specified VLAN. Specify a VLAN tag identifier from <b>1</b> through <b>4094</b>.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree interface on page 5045</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Understanding STP for EX Series Switches on page 4910</a></li> <li>• <a href="#">Understanding RSTP for EX Series Switches on page 4906</a></li> <li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li> </ul>                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">show spanning-tree bridge on page 5038</a><br><a href="#">show spanning-tree bridge brief on page 5038</a><br><a href="#">show spanning-tree bridge detail on page 5039</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | <p><a href="#">Table 556 on page 5035</a> lists the output fields for the <b>show spanning-tree bridge</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Table 556: show spanning-tree bridge Output Fields**

| Field Name        | Field Description                   | Level of Output                       |
|-------------------|-------------------------------------|---------------------------------------|
| <b>Context ID</b> | An internally generated identifier. | VSTP all, RSTP all, MSTP all, STP all |

Table 556: show spanning-tree bridge Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                                               | Level of Output                       |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| <b>Enabled protocol</b>                | Spanning-tree protocol type enabled.                                                                                                            | VSTP all, RSTP all, MSTP all, STP all |
| <b>Root ID</b>                         | Bridge ID of the elected spanning tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge. | VSTP all, RSTP all, MSTP all, STP all |
| <b>Root cost</b>                       | Calculated cost to reach the root bridge from the bridge where the command is entered.                                                          | VSTP all                              |
| <b>Root port</b>                       | Interface that is the current elected root port for this bridge.                                                                                | VSTP all                              |
| <b>CIST regional root</b>              | Bridge ID of the elected MSTP regional root bridge.                                                                                             | MSTP all                              |
| <b>CIST internal root cost</b>         | Calculated cost to reach the regional root bridge from the bridge where the command is entered.                                                 | MSTP all                              |
| <b>Hello time</b>                      | Configured number of seconds between transmissions of configuration BPDUs.                                                                      | VSTP all, RSTP all, MSTP all, STP all |
| <b>Maximum age</b>                     | Maximum age of received protocol BPDUs.                                                                                                         | VSTP all, RSTP all, MSTP all, STP all |
| <b>Forward delay</b>                   | Configured time an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.                   | VSTP all, RSTP all, MSTP all, STP all |
| <b>Message age</b>                     | Number of seconds elapsed since the most recent BPDU was received.                                                                              | VSTP all, RSTP all, STP all           |
| <b>Number of topology changes</b>      | Total number of STP topology changes detected since the switch last booted.                                                                     | VSTP all, RSTP all, MSTP all, STP all |
| <b>Time since last topology change</b> | Number of seconds elapsed since the most recent topology change.                                                                                | RSTP detail                           |
| <b>Topology change initiator</b>       | Interface name of the interface that received the topology change request.                                                                      | RSTP detail                           |

Table 556: show spanning-tree bridge Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                     | Level of Output                                   |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| <b>Topology change last recvd. from</b> | Bridge ID of the bridge that requested the last topology change.                                                                      | RSTP detail                                       |
| <b>Bridge ID (local)</b>                | Locally configured bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.             | VSTP all, RSTP all, MSTP all, STP all             |
| <b>Extended system ID (local)</b>       | Internally generated system identifier.                                                                                               | VSTP all, RSTP all, MSTP all, STP all             |
| <b>MSTI regional root</b>               | Bridge ID of the elected MSTP regional root bridge.                                                                                   | MSTP detail                                       |
| <b>Internal instance ID (local)</b>     | An internally generated identifier.                                                                                                   | VSTP all, RSTP all, MSTP all, STP all             |
| <b>Hello time (local)</b>               | Configured number of seconds between transmissions of configuration BPDUs.                                                            | RSTP detail, MSTP detail, STP detail              |
| <b>Maximum age (local)</b>              | Maximum age of received protocol BPDUs.                                                                                               | VSTP detail, RSTP detail, MSTP detail, STP detail |
| <b>Forward delay (local)</b>            | Configured time an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.         | RSTP detail, MSTP detail, STP detail              |
| <b>Path Cost Method (local)</b>         | Bridges supporting 802.1D (legacy) implement only 16-bit values for path cost. Newer versions of this standard support 32-bit values. | VSTP detail, RSTP detail, MSTP detail, STP detail |
| <b>Maximum Hop count (local)</b>        | Configured maximum number of hops a BPDU can be forwarded in the MSTP region.                                                         | MSTP detail                                       |

## Sample Output

### show spanning-tree bridge

```
user@switch> show spanning-tree bridge
STP bridge parameters
Context ID                : 0
Enabled protocol          : MSTP

STP bridge parameters for CIST
Root ID                   : 32768.00:11:f2:56:df:40
Root cost                  : 0
Root port                 : ge-0/0/1.0
CIST regional root        : 32768.00:11:f2:56:df:40
CIST internal root cost   : 20000
Hello time                 : 2 seconds
Maximum age                : 20 seconds
Forward delay              : 15 seconds
Hop count                  : 19
Message age                : 0
Number of topology changes : 1
Time since last topology change : 108 seconds
Topology change initiator  : ge-0/0/1.0
Topology change last recvd. from : 00:11:f2:56:df:4c
Local parameters
  Bridge ID                : 32768.00:11:f2:57:1c:00
  Extended system ID       : 0
  Internal instance ID      : 0

STP bridge parameters for MSTI 10
MSTI regional root        : 32778.00:11:f2:56:df:40
Root cost                  : 20000
Root port                 : ge-0/0/1.0
Hello time                 : 2 seconds
Maximum age                : 20 seconds
Forward delay              : 15 seconds
Hop count                  : 19
Number of topology changes : 1
Time since last topology change : 108 seconds
Topology change initiator  : ge-0/0/1.0
Topology change last recvd. from : 00:11:f2:56:df:41
Local parameters
  Bridge ID                : 32778.00:11:f2:57:1c:00
  Extended system ID       : 0
  Internal instance ID      : 1
```

### show spanning-tree bridge brief

```
user@switch> show spanning-tree bridge brief
STP bridge parameters
Context ID      : 0
Enabled protocol : RSTP
Root ID        : 32768.00:19:e2:50:95:a0
Hello time     : 2 seconds
Maximum age    : 20 seconds
Forward delay  : 15 seconds
Message age    : 0
Number of topology changes : 0
Local parameters
  Bridge ID      : 32768.00:19:e2:50:95:a0
```

```
Extended system ID      : 0
Internal instance ID    : 0
```

### show spanning-tree bridge detail

```
user@switch> show spanning-tree bridge detail
```

```
STP bridge parameters
Context ID      : 0
Enabled protocol : RSTP
Root ID        : 32768.00:19:e2:50:95:a0
Hello time     : 2 seconds
Maximum age    : 20 seconds
Forward delay  : 15 seconds
Message age    : 0
Number of topology changes : 0
Local parameters
Bridge ID      : 32768.00:19:e2:50:95:a0
Extended system ID : 0
Internal instance ID : 0
Hello time     : 2 seconds
Maximum age    : 20 seconds
Forward delay  : 15 seconds
Path cost method : 32 bit
```

## show spanning-tree bridge

**List of Syntax** [Syntax on page 5040](#)  
[Syntax \(QFX Series\) on page 5040](#)

**Syntax** `show spanning-tree bridge`  
`<brief | detail>`  
`<msti msti-id>`  
`<routing-instance routing-instance-name>`  
`<vlan-id vlan-id>`

**Syntax (QFX Series)** `show spanning-tree bridge`  
`<brief | detail>`  
`<msti msti-id>`  
`<vlan-id vlan-id>`

**Release Information** Command introduced in Junos OS Release 8.4.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display the configured or calculated Spanning Tree Protocol (STP) parameters.

**Options** **none**—(Optional) Display brief STP bridge information for all multiple spanning-tree instances (MSTIs).

**brief | detail**—(Optional) Display the specified level of output.

**msti *msti-id***—(Optional) Display STP bridge information for the specified MSTI.

**routing-instance *routing-instance-name***—(Optional) Display STP bridge information for the specified routing instance.

**vlan-id *vlan-id***—(Optional) Display STP bridge information for the specified VLAN.

**Required Privilege Level** view

**List of Sample Output** [show spanning-tree bridge routing-instance on page 5041](#)  
[show spanning-tree bridge msti on page 5042](#)  
[show spanning-tree bridge vlan-id \(MSTP\) on page 5043](#)  
[show spanning-tree bridge \(RSTP\) on page 5043](#)  
[show spanning-tree bridge vlan-id \(RSTP\) on page 5044](#)

**Output Fields** [Table 557 on page 5040](#) lists the output fields for the **show spanning-tree bridge** command. Output fields are listed in the approximate order in which they appear.

**Table 557: show spanning-tree bridge Output Fields**

| Field Name            | Field Description                                                  |
|-----------------------|--------------------------------------------------------------------|
| Routing instance name | Name of the routing instance under which the bridge is configured. |
| Enabled protocol      | Spanning Tree Protocol type enabled.                               |



Table 557: show spanning-tree bridge Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Root ID                         | Bridge ID of the elected spanning-tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge. |
| Root cost                       | Calculated cost to reach the root bridge from the bridge where the command is entered.                                                          |
| Root port                       | Interface that is the current elected root port for this bridge.                                                                                |
| CIST regional root              | Bridge ID of the elected MSTP regional root bridge.                                                                                             |
| CIST internal root cost         | Calculated cost to reach the regional root bridge from the bridge where the command is entered.                                                 |
| Hello time                      | Configured number of seconds between transmissions of configuration bridge protocol data units (BPDUs).                                         |
| Maximum age                     | Configured maximum expected arrival time of hello bridge protocol data units (BPDUs).                                                           |
| Forward delay                   | How long an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.                          |
| Hop count                       | Configured maximum number of hops a BPDU can be forwarded in the MSTP region.                                                                   |
| Message age                     | Number of elapsed seconds since the most recent BPDU was received.                                                                              |
| Number of topology changes      | Total number of STP topology changes detected since the routing device last booted.                                                             |
| Time since last topology change | Number of elapsed seconds since the most recent topology change.                                                                                |
| Bridge ID (Local)               | Locally configured bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.                       |
| Extended system ID              | System identifier.                                                                                                                              |
| MSTI regional root              | Bridge ID of the elected MSTP regional root bridge.                                                                                             |

## Sample Output

### show spanning-tree bridge routing-instance

```

user@host> show spanning-tree bridge routing-instance vs1 detail
STP bridge parameters
Routing instance name       : vs1
Enabled protocol           : MSTP

```

```
STP bridge parameters for CIST
  Root ID                : 32768.00:13:c3:9e:c8:80
  Root cost               : 0
  Root port              : ge-10/2/0
  CIST regional root     : 32768.00:13:c3:9e:c8:80
  CIST internal root cost : 22000
  Hello time             : 2 seconds
  Maximum age            : 20 seconds
  Forward delay          : 15 seconds
  Hop count              : 18
  Message age            : 0
  Number of topology changes : 1
  Time since last topology change : 1191 seconds
  Local parameters
    Bridge ID            : 32768.00:90:69:0b:7f:d1
    Extended system ID   : 1

STP bridge parameters for MSTI 1
  MSTI regional root     : 32769.00:13:c3:9e:c8:80
  Root cost              : 22000
  Root port              : ge-10/2/0
  Hello time             : 2 seconds
  Maximum age            : 20 seconds
  Forward delay          : 15 seconds
  Hop count              : 18
  Number of topology changes : 1
  Time since last topology change : 1191 seconds
  Local parameters
    Bridge ID            : 32769.00:90:69:0b:7f:d1
    Extended system ID   : 1

STP bridge parameters for MSTI 2
  MSTI regional root     : 32770.00:13:c3:9e:c8:80
  Root cost              : 22000
  Root port              : ge-10/2/0
  Hello time             : 2 seconds
  Maximum age            : 20 seconds
  Forward delay          : 15 seconds
  Hop count              : 18
  Number of topology changes : 1
  Time since last topology change : 1191 seconds
  Local parameters
    Bridge ID            : 32770.00:90:69:0b:7f:d1
    Extended system ID   : 1
```

### show spanning-tree bridge msti

```
user@host> show spanning-tree bridge msti 1 routing-instance vs1 detail
STP bridge parameters
Routing instance name      : vs1
Enabled protocol          : MSTP

STP bridge parameters for MSTI 1
  MSTI regional root     : 32769.00:13:c3:9e:c8:80
  Root cost              : 22000
  Root port              : xe-10/2/0
  Hello time             : 2 seconds
  Maximum age            : 20 seconds
  Forward delay          : 15 seconds
  Hop count              : 18
```

```

Number of topology changes      : 1
Time since last topology change : 1191 seconds
Local parameters
  Bridge ID                     : 32769.00:90:69:0b:7f:d1
  Extended system ID            : 1

```

### show spanning-tree bridge vlan-id (MSTP)

```
user@host> show spanning-tree bridge vlan-id 1101 routing-instance vs1 detail
```

```

STP bridge parameters
Routing instance name          : vs1
Enabled protocol               : MSTP

STP bridge parameters for CIST
Root ID                       : 32768.00:13:c3:9e:c8:80
Root cost                     : 0
Root port                     : xe-10/2/0
CIST regional root            : 32768.00:13:c3:9e:c8:80
CIST internal root cost       : 22000
Hello time                    : 2 seconds
Maximum age                   : 20 seconds
Forward delay                  : 15 seconds
Hop count                     : 18
Message age                   : 0
Number of topology changes    : 0
Local parameters
  Bridge ID                   : 32768.00:90:69:0b:7f:d1
  Extended system ID          : 1
  Hello time                  : 2 seconds
  Maximum age                 : 20 seconds
  Forward delay                : 15 seconds
  Path cost method             : 32 bit
  Maximum hop count           : 20

```

### show spanning-tree bridge (RSTP)

```
user@host> show spanning-tree bridge
```

```

STP bridge parameters
Routing instance name          : GLOBAL
Enabled protocol               : RSTP
Root ID                       : 28672.00:90:69:0b:3f:d0
Hello time                    : 2 seconds
Maximum age                   : 20 seconds
Forward delay                  : 15 seconds
Message age                   : 0
Number of topology changes    : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                   : 28672.00:90:69:0b:3f:d0
  Extended system ID          : 0

STP bridge parameters for bridge VLAN 10
Root ID                       : 28672.00:90:69:0b:3f:d0
Hello time                    : 2 seconds
Maximum age                   : 20 seconds
Forward delay                  : 15 seconds
Message age                   : 0
Number of topology changes    : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                   : 28672.00:90:69:0b:3f:d0

```

```
Extended system ID          : 0

STP bridge parameters for bridge VLAN 20
Root ID                     : 28672.00:90:69:0b:3f:d0
Hello time                   : 2 seconds
Maximum age                  : 20 seconds
Forward delay                : 15 seconds
Message age                  : 0
Number of topology changes   : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                  : 28672.00:90:69:0b:3f:d0
  Extended system ID        : 0
```

#### show spanning-tree bridge vlan-id (RSTP)

```
user@host> show spanning-tree bridge vlan-id 10
STP bridge parameters
Routing instance name       : GLOBAL
Enabled protocol            : RSTP

STP bridge parameters for VLAN 10
Root ID                     : 28672.00:90:69:0b:3f:d0
Hello time                   : 2 seconds
Maximum age                  : 20 seconds
Forward delay                : 15 seconds
Message age                  : 0
Number of topology changes   : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                  : 28672.00:90:69:0b:3f:d0
  Extended system ID        : 0
```

## show spanning-tree interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show spanning-tree interface<br><brief   detail><br><interface-name <i>interface-name</i> ><br><msti <i>msti-id</i> ><br><vlan-id <i>vlan-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Display the configured or calculated interface-level spanning-tree protocol (can be either STP, RSTP, or MSTP) parameters. In <b>brief</b> mode, will not display interfaces that are administratively disabled or do not have a physical link.                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>none</b>—(Optional) Display brief STP interface information.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface-name <i>interface-name</i></b>—(Optional) Name of an interface.</p> <p><b>msti <i>msti-id</i></b>—(Optional) Display STP bridge information for the specified MSTP instance ID or Common and Internal Spanning Tree (CIST). Specify <b>0</b> for CIST. Specify a value from 1 through <b>4094</b> for an MSTI.</p> <p><b>vlan-id <i>vlan-id</i></b>—(Optional) For MSTP interfaces, display interface information for the specified VLAN. Specify a value from <b>0</b> through <b>4094</b>.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Understanding STP for EX Series Switches on page 4910</a></li> <li>• <a href="#">Understanding RSTP for EX Series Switches on page 4906</a></li> <li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li> <li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li> </ul>                                                                       |
| <b>List of Sample Output</b>    | <a href="#">show spanning-tree interface on page 5046</a><br><a href="#">show spanning-tree interface brief on page 5047</a><br><a href="#">show spanning-tree interface detail on page 5047</a><br><a href="#">show spanning-tree interface ge-1/0/0 on page 5048</a>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>            | <a href="#">Table 558 on page 5046</a> lists the output fields for the <b>show spanning-tree interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

Table 558: show spanning-tree interface Output Fields

| Field Name                    | Field Description                                                                                                 |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------|
| Interface name                | Interface configured to participate in the STP, RSTP, or MSTP instance.                                           |
| Port ID                       | Logical interface identifier configured to participate in the MSTP instance.                                      |
| Designated port ID            | Port ID of the designated port for the LAN segment this interface is attached to.                                 |
| Designated bridge ID          | Bridge ID of the designated bridge for the LAN segment this interface is attached to.                             |
| Port Cost                     | Configured cost for the interface.                                                                                |
| Port State                    | STP port state. Forwarding (FWD), blocking (BLK), listening, learning, or disabled.                               |
| Port Role                     | MSTP or RSTP port role. Designated (DESG), backup (BKUP), alternate (ALT), (ROOT), or Root Prevented (Root-Prev). |
| Link type                     | MSTP or RSTP link type. Shared or point-to-point (pt-pt) and edge or non edge.                                    |
| Alternate                     | Identifies the interface as an MSTP or RSTP alternate root port (yes) or non-alternate root port (no).            |
| Boundary Port                 | Identifies the interface as an MSTP regional boundary port (yes) or non-boundary port (no).                       |
| Edge delay while expiry count | Number of times the edge delay timer expired on that interface.                                                   |
| Rcvd info while expiry count  | Number of times the rcvd info timer expired on that interface.                                                    |

## Sample Output

### show spanning-tree interface

```
user@switch> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/0.0  | 128:513 | 128:513               | 8192.0019e2500340       | 1000         | FWD   | DESG |
| ge-0/0/2.0  | 128:515 | 128:515               | 8192.0019e2500340       | 1000         | BLK   | DIS  |
| ge-0/0/4.0  | 128:517 | 128:517               | 8192.0019e2500340       | 1000         | FWD   | DESG |
| ge-0/0/23.0 | 128:536 | 128:536               | 8192.0019e2500340       | 1000         | FWD   | DESG |

```
Spanning tree interface parameters for instance 1
```

| Interface  | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/0.0 | 128:513 | 128:513               | 8193.0019e2500340       | 1000         | FWD   | DESG |
| ge-0/0/2.0 | 128:515 | 128:515               | 8193.0019e2500340       | 1000         | BLK   | DIS  |
| ge-0/0/4.0 | 128:517 | 128:517               | 8193.0019e2500340       | 1000         | FWD   | DESG |

```
ge-0/0/23.0 128:536 128:536 8193.0019e2500340 1000 FWD DESG
```

Spanning tree interface parameters for instance 2

| Interface   | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-------------|---------|--------------------|----------------------|-----------|-------|------|
| ge-0/0/0.0  | 128:513 | 128:1              | 8194.001b549fd000    | 1000      | FWD   | ROOT |
| ge-0/0/2.0  | 128:515 | 128:515            | 32770.0019e2500340   | 4000      | BLK   | DIS  |
| ge-0/0/4.0  | 128:517 | 128:1              | 16386.001b54013080   | 1000      | BLK   | ALT  |
| ge-0/0/23.0 | 128:536 | 128:536            | 32770.0019e2500340   | 1000      | FWD   | DESG |

### show spanning-tree interface brief

```
user@switch> show spanning-tree interface brief
Spanning tree interface parameters for instance 0
```

| Interface   | Port ID | Designated port ID | Designated bridge ID | Port Cost | State | Role |
|-------------|---------|--------------------|----------------------|-----------|-------|------|
| ge-1/0/0.0  | 128:625 | 128:625            | 32768.0019e25095a0   | 20000     | BLK   | DIS  |
| ge-1/0/1.0  | 128:626 | 128:626            | 32768.0019e25095a0   | 20000     | BLK   | DIS  |
| ge-1/0/2.0  | 128:627 | 128:627            | 32768.0019e25095a0   | 20000     | BLK   | DIS  |
| ge-1/0/10.0 | 128:635 | 128:635            | 32768.0019e25095a0   | 20000     | BLK   | DIS  |
| ge-1/0/20.0 | 128:645 | 128:645            | 32768.0019e25095a0   | 20000     | BLK   | DIS  |
| ge-1/0/30.0 | 128:655 | 128:655            | 32768.0019e25095a0   | 20000     | BLK   | DIS  |

### show spanning-tree interface detail

```
user@switch> show spanning-tree interface detail
Spanning tree interface parameters for instance 0
```

```
Interface name      : ge-1/0/0.0
Port identifier     : 128.625
Designated port ID  : 128.625
Port cost           : 20000
Port state          : Blocking
Designated bridge ID : 32768.00:19:e2:50:95:a0
Port role           : Disabled
Link type           : Pt-Pt/EDGE
Boundary port       : NA
Edge delay while expiry count : 0
Rcvd info while expiry count  : 0
```

```
Interface name      : ge-1/0/1.0
Port identifier     : 128.626
Designated port ID  : 128.626
Port cost           : 20000
Port state          : Blocking
Designated bridge ID : 32768.00:19:e2:50:95:a0
Port role           : Disabled
Link type           : Pt-Pt/NONEDGE
Boundary port       : NA
Edge delay while expiry count : 0
Rcvd info while expiry count  : 0
```

```
Interface name      : ge-1/0/2.0
Port identifier     : 128.627
Designated port ID  : 128.627
Port cost           : 20000
Port state          : Blocking
Designated bridge ID : 32768.00:19:e2:50:95:a0
```

```
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port   : NA
Edge delay while expiry count : 0
Rvcd info while expiry count  : 0

Interface name   : ge-1/0/10.0
Port identifier   : 128.635
Designated port ID : 128.635
Port cost        : 20000
Port state       : Blocking
Designated bridge ID : 32768.00:19:e2:50:95:a0
Port role        : Disabled
Link type        : Pt-Pt/NONEDGE
Boundary port     : NA
Edge delay while expiry count : 0
Rvcd info while expiry count  : 0

Interface name   : ge-1/0/20.0
Port identifier   : 128.645
Designated port ID : 128.645
Port cost        : 20000
Port state       : Blocking
Designated bridge ID : 32768.00:19:e2:50:95:a0
Port role        : Disabled
Link type        : Pt-Pt/NONEDGE
Boundary port     : NA
Edge delay while expiry count : 0
Rvcd info while expiry count  : 0
[output truncated]
```

#### show spanning-tree interface ge-1/0/0

```
user@switch> show spanning-tree interface ge-1/0/0
```

| Interface  | Port ID | Designated | Designated         | Port  | State | Role |
|------------|---------|------------|--------------------|-------|-------|------|
|            | port ID | bridge ID  | Cost               |       |       |      |
| ge-1/0/0.0 | 128:625 | 128:625    | 32768.0019e25095a0 | 20000 | BLK   | DIS  |



## show spanning-tree interface

|                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                 | <a href="#">Syntax on page 5049</a><br><a href="#">Syntax (EX Series Switches and the QFX Series) on page 5049</a>                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                         | <pre>show spanning-tree interface &lt;brief   detail&gt; &lt;msti <i>msti-id</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt; &lt;vlan-id <i>vlan-id</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (EX Series Switches and the QFX Series)</b> | <pre>show spanning-tree interface &lt;brief   detail&gt; &lt;msti <i>msti-id</i>&gt; &lt;vlan-id <i>vlan-id</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                            | <p>Command introduced in Junos OS Release 8.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                                    | Display the configured or calculated interface-level STP parameters.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                        | <p><b>none</b>—Display brief STP interface information.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>msti <i>msti-id</i></b>—(Optional) Display STP interface information for the specified MST instance.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Display STP interface information for the specified routing instance.</p> <p><b>vlan-id <i>vlan-id</i></b>—(Optional) Display STP interface information for the specified VLAN.</p> |
| <b>Required Privilege Level</b>                       | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                          | <a href="#">show spanning-tree interface on page 5050</a><br><a href="#">show spanning-tree interface (QFX Series) on page 5051</a><br><a href="#">show spanning-tree interface detail on page 5051</a><br><a href="#">show spanning-tree interface msti on page 5053</a><br><a href="#">show spanning-tree interface vlan-id on page 5053</a><br><a href="#">show spanning-tree interface (VSTP) on page 5054</a><br><a href="#">show spanning-tree interface vlan-id (VSTP) on page 5054</a>                   |
| <b>Output Fields</b>                                  | <p><a href="#">Table 559 on page 5049</a> lists the output fields for the <b>show spanning-tree interface</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                               |

**Table 559: show spanning-tree Interface Output Fields**

| Field Name            | Field Description                                                             |
|-----------------------|-------------------------------------------------------------------------------|
| <b>Interface name</b> | Interface configured to participate in the STP, RSTP, VSTP, or MSTP instance. |

Table 559: show spanning-tree Interface Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                     |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Port ID</b>              | Logical interface identifier configured to participate in the MSTP or VSTP instance.                                                                                  |
| <b>Designated port ID</b>   | Port ID of the designated port for the LAN segment to which this interface is attached.                                                                               |
| <b>Designated bridge ID</b> | Bridge ID of the designated bridge for the LAN segment to which this interface is attached.                                                                           |
| <b>Port Cost</b>            | Configured cost for the interface.                                                                                                                                    |
| <b>Port State</b>           | STP port state: forwarding ( <b>FWD</b> ), blocking ( <b>BLK</b> ), listening, learning, or disabled.                                                                 |
| <b>Port Role</b>            | MSTP, VSTP, or RSTP port role: designated ( <b>DESG</b> ), backup ( <b>BKUP</b> ), alternate ( <b>ALT</b> ), ( <b>ROOT</b> ), or Root Prevented ( <b>Root-Prev</b> ). |
| <b>Link type</b>            | MSTP, VSTP, or RSTP link type. Shared or point-to-point (pt-pt) and edge or nonedge.                                                                                  |
| <b>Alternate</b>            | Identifies the interface as an MSTP, VSTP, or RSTP alternate root port ( <b>Yes</b> ) or nonalternate root port ( <b>No</b> ).                                        |
| <b>Boundary Port</b>        | Identifies the interface as an MSTP regional boundary port ( <b>Yes</b> ) or nonboundary port ( <b>No</b> ).                                                          |

## Sample Output

### show spanning-tree interface

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1       | 128:1   | 128:1                 | 32768.0090690b47d1      | 1000         | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2                 | 32768.0090690b47d1      | 20000        | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3                 | 32768.0090690b47d1      | 29999        | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26                | 32768.0013c39ec880      | 20000        | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5                 | 32768.0090690b47d1      | 2000         | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6                 | 32768.0090690b47d1      | 2000         | FWD   | DESG |

```
Spanning tree interface parameters for instance 1
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1       | 128:1   | 128:1                 | 32769.0090690b47d1      | 1000         | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2                 | 32769.0090690b47d1      | 20000        | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3                 | 32769.0090690b47d1      | 29999        | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26                | 32769.0013c39ec880      | 20000        | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5                 | 32769.0090690b47d1      | 2000         | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6                 | 32769.0090690b47d1      | 2000         | FWD   | DESG |

Spanning tree interface parameters for instance 2

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1       | 128:1   | 128:1                 | 32770.0090690b47d1      | 1000         | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2                 | 32770.0090690b47d1      | 20000        | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3                 | 32770.0090690b47d1      | 29999        | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26                | 32770.0013c39ec880      | 20000        | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5                 | 32770.0090690b47d1      | 2000         | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6                 | 32770.0090690b47d1      | 2000         | FWD   | DESG |

### show spanning-tree interface (QFX Series)

```
user@1f0> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1       | 128:1   | 128:1                 | 32768.0090690b47d1      | 1000         | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2                 | 32768.0090690b47d1      | 20000        | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3                 | 32768.0090690b47d1      | 29999        | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26                | 32768.0013c39ec880      | 20000        | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5                 | 32768.0090690b47d1      | 2000         | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6                 | 32768.0090690b47d1      | 2000         | FWD   | DESG |

Spanning tree interface parameters for instance 1

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1       | 128:1   | 128:1                 | 32769.0090690b47d1      | 1000         | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2                 | 32769.0090690b47d1      | 20000        | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3                 | 32769.0090690b47d1      | 29999        | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26                | 32769.0013c39ec880      | 20000        | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5                 | 32769.0090690b47d1      | 2000         | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6                 | 32769.0090690b47d1      | 2000         | FWD   | DESG |

Spanning tree interface parameters for instance 2

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ae1       | 128:1   | 128:1                 | 32770.0090690b47d1      | 1000         | FWD   | DESG |
| ge-2/1/2  | 128:2   | 128:2                 | 32770.0090690b47d1      | 20000        | FWD   | DESG |
| ge-2/1/5  | 128:3   | 128:3                 | 32770.0090690b47d1      | 29999        | FWD   | DESG |
| ge-2/2/1  | 128:4   | 128:26                | 32770.0013c39ec880      | 20000        | FWD   | ROOT |
| xe-9/2/0  | 128:5   | 128:5                 | 32770.0090690b47d1      | 2000         | FWD   | DESG |
| xe-9/3/0  | 128:6   | 128:6                 | 32770.0090690b47d1      | 2000         | FWD   | DESG |

### show spanning-tree interface detail

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

```
Interface name           : ae1
Port identifier          : 128.1
Designated port ID      : 128.1
Port cost                : 1000
Port state               : Forwarding
Designated bridge ID     : 32768.00:90:69:0b:47:d1
Port role                : Designated
Link type                : Pt-Pt/NONEDGE
```

```
Boundary port                : No

Interface name                : ge-2/1/2
Port identifier               : 128.2
Designated port ID           : 128.2
Port cost                     : 20000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                : No

Interface name                : ge-2/1/5
Port identifier               : 128.3
Designated port ID           : 128.3
Port cost                     : 29999
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                : No

Interface name                : ge-2/2/1
Port identifier               : 128.4
Designated port ID           : 128.26
Port cost                     : 20000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:13:c3:9e:c8:80
Port role                     : Root
Link type                     : Pt-Pt/NONEDGE
Boundary port                : No

Interface name                : xe-9/2/0
Port identifier               : 128.5
Designated port ID           : 128.5
Port cost                     : 2000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                : No

Interface name                : xe-9/3/0
Port identifier               : 128.6
Designated port ID           : 128.6
Port cost                     : 2000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                : No
```

#### Spanning tree interface parameters for instance 1

```
Interface name                : ae1
Port identifier               : 128.1
Designated port ID           : 128.1
Port cost                     : 1000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
```

```

Port role           : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

Interface name       : ge-2/1/2
Port identifier      : 128.2
Designated port ID   : 128.2
Port cost            : 20000
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role            : Designated
Link type            : Pt-Pt/NONEDGE
Boundary port        : No

Interface name       : ge-2/1/5
Port identifier      : 128.3
Designated port ID   : 128.3
Port cost            : 29999
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role            : Designated
Link type            : Pt-Pt/NONEDGE
Boundary port        : No

Interface name       : ge-2/2/1
Port identifier      : 128.4
Designated port ID   : 128.26
Port cost            : 20000
Port state           : Forwarding
Designated bridge ID : 32768.00:13:c3:9e:c8:80
Port role            : Root
Link type            : Pt-Pt/NONEDGE
Boundary port        : No

...

```

### show spanning-tree interface msti

```

user@host> show spanning-tree interface msti 1 routing-instance vs1 detail
Spanning tree interface parameters for instance 1

```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-7/0/0  | 128:1   | 128:1                 | 32769.0090690b4fd1      | 2000         | FWD   | DESG |
| ge-5/1/0  | 128:2   | 128:2                 | 32769.0090690b4fd1      | 20000        | FWD   | DESG |
| ge-5/1/1  | 128:3   | 128:3                 | 32769.0090690b4fd1      | 20000        | FWD   | DESG |
| ae1       | 128:4   | 128:1                 | 32769.0090690b47d1      | 10000        | BLK   | ALT  |
| ge-5/1/4  | 128:5   | 128:3                 | 32769.0090690b47d1      | 20000        | BLK   | ALT  |
| xe-7/2/0  | 128:6   | 128:6                 | 32769.0090690b47d1      | 2000         | FWD   | ROOT |

### show spanning-tree interface vlan-id

```

user@host> show spanning-tree interface vlan-id 101 routing-instance vs1 detail
Spanning tree interface parameters for instance 0

```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-----------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-11/0/5 | 128:1   | 128:1                 | 32768.0090690b7fd1      | 20000        | FWD   | DESG |
| ge-11/0/6 | 128:2   | 128:1                 | 32768.0090690b7fd1      | 20000        | BLK   | BKUP |
| ge-11/1/0 | 128:3   | 128:2                 | 32768.0090690b4fd1      | 20000        | BLK   | ALT  |
| ge-11/1/1 | 128:4   | 128:3                 | 32768.0090690b4fd1      | 20000        | BLK   | ALT  |

|           |       |       |                    |       |     |      |
|-----------|-------|-------|--------------------|-------|-----|------|
| ge-11/1/4 | 128:5 | 128:1 | 32768.0090690b47d1 | 20000 | BLK | ALT  |
| xe-10/0/0 | 128:6 | 128:5 | 32768.0090690b4fd1 | 2000  | BLK | ALT  |
| xe-10/2/0 | 128:7 | 128:4 | 32768.0090690b47d1 | 2000  | FWD | ROOT |

**show spanning-tree interface (VSTP)**

```
user@host> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Cost  | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1  | 128:1   | 128:1                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |
| ge-1/0/2  | 128:2   | 128:2                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |

```
Spanning tree interface parameters for VLAN 10
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Cost  | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1  | 128:1   | 128:1                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |
| ge-1/0/2  | 128:2   | 128:2                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |

```
Spanning tree interface parameters for VLAN 20
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Cost  | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1  | 128:1   | 128:1                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |
| ge-1/0/2  | 128:2   | 128:2                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |

**show spanning-tree interface vlan-id (VSTP)**

```
user@host> show spanning-tree interface vlan-id 10
```

```
Spanning tree interface parameters for VLAN 10
```

| Interface | Port ID | Designated<br>port ID | Designated<br>bridge ID | Cost  | State | Role |
|-----------|---------|-----------------------|-------------------------|-------|-------|------|
| ge-1/0/1  | 128:1   | 128:1                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |
| ge-1/0/2  | 128:2   | 128:2                 | 28672.0090690b3fe0      | 20000 | FWD   | DESG |

## show spanning-tree mstp configuration

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show spanning-tree mstp configuration<br><brief   detail>                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display the MSTP configuration.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b>none</b> —Display MSTP configuration information.<br><b>brief   detail</b> —(Optional) Display the specified level of output.                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">show spanning-tree statistics on page 5059</a></li> <li>• <a href="#">mstp (Spanning Trees) on page 5010</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show spanning-tree mstp configuration on page 5056</a>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | <a href="#">Table 560 on page 5055</a> lists the output fields for the <b>show spanning-tree mstp configuration</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                    |

**Table 560: show spanning-tree mstp configuration Output Fields**

| Field Name           | Field Description                                                |
|----------------------|------------------------------------------------------------------|
| Context identifier   | Internally generated identifier.                                 |
| Region name          | MSTP region name carried in the MSTP BPDUs.                      |
| Revision             | Revision number of the MSTP configuration.                       |
| Configuration digest | Numerical value derived from the VLAN-to-instance mapping table. |
| MSTI                 | MSTI instance identifier.                                        |
| Member VLANs         | Identifiers for VLANs associated with the MSTI.                  |

## Sample Output

### show spanning-tree mstp configuration

```
user@host> show spanning-tree mstp configuration
MSTP configuration information
Context identifier      : 0
Region name            : region1
Revision               : 0
Configuration digest    : 0xc92e7af9febb44d8df928b87f16b
```

```
MSTI      Member VLANs
0 0-100,105-4094
1 101-102
2 103-104
```



## show spanning-tree mstp configuration

|                                                     |                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5057</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5057</a>                                                                                                                                                                        |
| <b>Syntax</b>                                       | show spanning-tree mstp configuration<br><brief   detail><br><routing-instance <i>routing-instance-name</i> >                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show spanning-tree mstp configuration<br><brief   detail>                                                                                                                                                                                                                               |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 8.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                    |
| <b>Description</b>                                  | Display the MSTP configuration.                                                                                                                                                                                                                                                         |
| <b>Options</b>                                      | <b>none</b> —Display MSTP configuration information.<br><br><b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>routing-instance <i>routing-instance-name</i></b> —(Optional) Display MSTP configuration information for the specified routing instance. |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>                        | <a href="#">show spanning-tree mstp configuration detail on page 5058</a><br><a href="#">show spanning-tree mstp configuration detail (QFX Series) on page 5058</a>                                                                                                                     |
| <b>Output Fields</b>                                | <a href="#">Table 561 on page 5057</a> lists the output fields for the <b>show spanning-tree mstp configuration</b> command. Output fields are listed in the approximate order in which they appear.                                                                                    |

**Table 561: show spanning-tree mstp configuration Output Fields**

| Field Name                  | Field Description                                                |
|-----------------------------|------------------------------------------------------------------|
| <b>Context id</b>           | Internally generated identifier.                                 |
| <b>Region name</b>          | MSTP region name carried in the MSTP BPDUs.                      |
| <b>Revision</b>             | Revision number of the MSTP configuration.                       |
| <b>Configuration digest</b> | Numerical value derived from the VLAN-to-instance mapping table. |
| <b>MSTI</b>                 | MST instance identifier.                                         |
| <b>Member VLANs</b>         | VLAN identifiers associated with the MSTI.                       |

## Sample Output

### show spanning-tree mstp configuration detail

```
user@host> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier      : 1
Region name            : henry
Revision               : 3
Configuration digest    : 0x6da4b5c4fd587757eef35675365e1

MSTI      Member VLANs
  0 0-99,101-199,201-4094
  1 100
  2 200
```

### show spanning-tree mstp configuration detail (QFX Series)

```
user@1f0> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier      : 1
Region name            : henry
Revision               : 3
Configuration digest    : 0x6da4b5c4fd587757eef35675365e1

MSTI      Member VLANs
  0 0-99,101-199,201-4094
  1 100
  2 200
```

## show spanning-tree statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show spanning-tree statistics<br>interface <i>interface-name</i><br>vlan <i>vlan-id</i><br><brief   detail>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Option <b>vlan</b> <i>vlan-id</i> introduced in Junos OS Release 10.1 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Display STP statistics on an interface, or for a VLAN when VSTP is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>none</b>—Display brief STP statistics.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface</b> <i>interface-name</i>—(Optional) The name of the interface.</p> <p><b>vlan</b> <i>vlan-id</i>—(Optional) The name of a VLAN.</p>                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 5035</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches on page 4936</a></li> <li>• <a href="#">Understanding STP for EX Series Switches on page 4910</a></li> <li>• <a href="#">Understanding RSTP for EX Series Switches on page 4906</a></li> <li>• <a href="#">Understanding MSTP for EX Series Switches on page 4904</a></li> <li>• <a href="#">Understanding VSTP for EX Series Switches and QFX Series Switches on page 4912</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show spanning-tree statistics interface on page 5060</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>            | <a href="#">Table 562 on page 5059</a> lists the output fields for the <b>show spanning-tree statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                     |

**Table 562: show spanning-tree statistics Output Fields**

| Field Name             | Field Description                                              |
|------------------------|----------------------------------------------------------------|
| BPDUs sent             | Total number of BPDUs sent.                                    |
| BPDUs received         | Total number of BPDUs received.                                |
| Interface              | Interface for which the statistics are being displayed.        |
| Next BPDU transmission | Number of seconds until the next BPDU is scheduled to be sent. |

## Sample Output

### show spanning-tree statistics interface

```
user@switch> show spanning-tree statistics interface ge-0/0/4
Interface  BPDUs sent  BPDUs received  Next BPDU
              transmission
ge-0/0/4    7    190    0
```

## show spanning-tree statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax on page 5061</a><br><a href="#">Syntax (EX Series Switch and the QFX Series) on page 5061</a>                                                                                                                                                                                                                                                               |
| <b>Syntax</b>                                       | <pre>show spanning-tree statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>                                                                                                                                                                                                                  |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show spanning-tree statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>   vlan <i>vlan-id</i>&gt;</pre>                                                                                                                                                                                                                                                  |
| <b>Release Information</b>                          | <p>Command introduced in Junos OS Release 8.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series switches.</p>                                                                                                                                                                       |
| <b>Description</b>                                  | Display STP statistics.                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                      | <p><b>none</b>—Display brief STP statistics.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display STP statistics for the specified interface.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Display STP statistics for the specified routing instance.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                        | <a href="#">show spanning-tree statistics routing-instance on page 5062</a><br><a href="#">show spanning-tree statistics interface routing-instance detail on page 5062</a>                                                                                                                                                                                                    |
| <b>Output Fields</b>                                | <p><a href="#">Table 563 on page 5061</a> lists the output fields for the <b>show spanning-tree statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                            |

**Table 563: show spanning-tree statistics Output Fields**

| Field Name                      | Field Description                                     |
|---------------------------------|-------------------------------------------------------|
| Message type                    | Type of message being counted.                        |
| BPDUs sent                      | Total number of BPDUs sent.                           |
| BPDUs received                  | Total number of BPDUs received.                       |
| BPDUs sent in last interval     | Number of BPDUs sent within a specified interval.     |
| BPDUs received in last interval | Number of BPDUs received within a specified interval. |

Table 563: show spanning-tree statistics Output Fields (*continued*)

| Field Name                    | Field Description                                              |
|-------------------------------|----------------------------------------------------------------|
| <b>Interface</b>              | Interface for which the statistics are being displayed.        |
| <b>Next BPDU transmission</b> | Number of seconds until the next BPDU is scheduled to be sent. |

## Sample Output

### show spanning-tree statistics routing-instance

```

user@host> show spanning-tree statistics routing-instance vs1 detail
Routing instance level STP statistics
Message type           : bpdus
BPDUs sent             : 1396
BPDUs received         : 1027
BPDUs sent in last interval : 5      (duration: 4 sec)
BPDUs received in last interval: 4    (duration: 4 sec)

```

### show spanning-tree statistics interface routing-instance detail

```

user@host> show spanning-tree statistics interface ge-11/1/4 routing-instance vs1 detail
Interface  BPDUs sent  BPDUs received  Next BPDU
                                     transmission
ge-11/1/4      7           190           0

```

## PART 26

# Virtual Chassis

- [Overview on page 5065](#)
- [Configuration on page 5097](#)
- [Administration on page 5149](#)
- [Troubleshooting on page 5235](#)





## CHAPTER 84

# Overview

- [Virtual Chassis Overview on page 5065](#)

### Virtual Chassis Overview

---

- [Understanding EX4300 Virtual Chassis on page 5065](#)
- [Understanding EX Series Virtual Chassis Components on page 5067](#)
- [Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric on page 5073](#)
- [Understanding How the Master in a Virtual Chassis Is Elected on page 5078](#)
- [Understanding Software Upgrade in an EX Series Virtual Chassis on page 5079](#)
- [Understanding Global Management of a Virtual Chassis on page 5080](#)
- [Understanding Nonvolatile Storage in a Virtual Chassis on page 5083](#)
- [Understanding EX Series Virtual Chassis Port Link Aggregation on page 5083](#)
- [Understanding EX Series Virtual Chassis Configuration on page 5086](#)
- [Understanding Split and Merge in a Virtual Chassis on page 5087](#)
- [Understanding Automatic Software Update on Virtual Chassis Member Switches on page 5090](#)
- [Understanding MAC Address Assignment on a Virtual Chassis on page 5093](#)
- [Understanding High Availability on an EX Series Virtual Chassis on page 5094](#)

### Understanding EX4300 Virtual Chassis

EX4300 switches can act as member switches in a non-mixed Virtual Chassis—a Virtual Chassis composed entirely of EX4300 switches—as well as participate as member switches in a QFX Series Virtual Chassis or as leaf nodes in a Virtual Chassis Fabric (VCF). This topic provides information regarding EX4300 member switches in a non-mixed EX4300 Virtual Chassis only. For information on EX4300 member switches in a QFX Series Virtual Chassis, see *Understanding QFX Series Virtual Chassis*. For information on EX4300 switches in a VCF, see [“Virtual Chassis Fabric Overview” on page 5241](#).

EX4300 Virtual Chassis brings the Virtual Chassis flexible, scaling switch solution to the Juniper Networks EX4300 Ethernet Switch. You can connect up to ten EX4300 switches together to form one EX4300 Virtual Chassis and manage the unit as a single chassis. The advantages of connecting multiple switches into a Virtual Chassis include

better-managed bandwidth at a network layer, simplified configuration and maintenance because multiple switches can be managed as a single switch, increased fault tolerance and high availability (HA) because a Virtual Chassis can remain active and network traffic can be redirected to other member switches when a single member switch fails, and a simplified Layer 2 network topology that minimizes or eliminates the need for loop prevention protocols such as Spanning Tree Protocol (STP).

The Virtual Chassis also provides a flexible model for expanding your network. If you are using an EX4300 switch or EX4300 Virtual Chassis at the access layer, for instance, and need additional access ports to support more servers, computers, phones, or other devices, you can add an EX4300 switch as a Virtual Chassis member to increase the number of access ports on your network with minimal complications to the existing network topology and switch configuration. You can add this new switch to a Virtual Chassis if the switch is installed in the same building or at a different site because the long-distance optical ports can be used to interconnect EX4300 switches into the a Virtual Chassis.

You configure an EX4300 Virtual Chassis by configuring optical interfaces connecting EX4300 switches into Virtual Chassis ports (VCPs). VCPs connect switches together to form a Virtual Chassis, and are responsible for passing all data and control traffic between member switches in the Virtual Chassis. All 40-Gigabit QSFP+ optical ports on an EX4300 switch are configured as VCPs by default. All 10-Gigabit optical ports on an EX4300 switch can be configured into VCPs. You can increase the VCP bandwidth between any two member switches by connecting multiple VCP links between the switches. When multiple VCP links are interconnecting the same two member switches, a Link Aggregation Group (LAG) bundle is formed when the links are identical speeds. For instance, if you have four 40-Gigabit links configured as VCPs between member switches, a LAG with four member links at 160Gbps of bandwidth is formed. 10-Gigabit and 40-Gigabit links configured as VCPs cannot be members of the same LAG, however.

All models of EX4300 switches can be interconnected into the same EX4300 Virtual Chassis. EX4300 switches cannot be interconnected into a Virtual Chassis with any other Juniper Networks product, including other EX Series Ethernet Switches.

An EX4300 Virtual Chassis is configured, monitored, and maintained like other EX Series Virtual Chassis. See [“Understanding EX Series Virtual Chassis Components” on page 5067](#).

**Related  
Documentation**

- [Configuring an EX4300 Virtual Chassis \(CLI Procedure\) on page 5097](#)
- [Understanding EX Series Virtual Chassis Components on page 5067](#)
- [Understanding How the Master in a Virtual Chassis Is Elected on page 5078](#)
- [Understanding EX Series Virtual Chassis Port Link Aggregation on page 5083](#)
- [Understanding EX Series Virtual Chassis Configuration on page 5086](#)

## Understanding EX Series Virtual Chassis Components



**NOTE:** This topic applies to all EX Series Virtual Chassis except EX8200 Virtual Chassis. See *Understanding EX8200 Virtual Chassis Components* for information about EX8200 Virtual Chassis.

This topic describes the components of EX Series Virtual Chassis—including the components of any mixed Virtual Chassis that contains EX series member switches—except EX8200 Virtual Chassis.

This topic covers:

- [Maximum Number of Switches per Virtual Chassis on page 5067](#)
- [Virtual Chassis Ports \(VCPs\) on page 5068](#)
- [Master Role on page 5070](#)
- [Backup Role on page 5070](#)
- [Linecard Role on page 5071](#)
- [Member Switch and Member ID on page 5071](#)
- [Mastership Priority on page 5072](#)
- [Mixed Virtual Chassis on page 5072](#)
- [Virtual Chassis Identifier \(VCID\) on page 5073](#)

### Maximum Number of Switches per Virtual Chassis

The maximum number of switches that a Virtual Chassis supports varies by Virtual Chassis and can depend on the Junos OS release running on the Virtual Chassis. [Table 564 on page 5067](#) lists the maximum member switch support by Virtual Chassis and Junos OS release.

**Table 564: Maximum Member Switch Support for Virtual Chassis by Junos OS Release**

| Maximum Member Switch Support | Initial Junos OS Release                                                                                                        |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| EX2200 Virtual Chassis        | 12.2R1—Initial release. Support for up to four EX2200 member switches.                                                          |
| EX3300 Virtual Chassis        | 11.3R1—Initial release. Support for up to six EX3300 member switches<br><br>12.2R1—Support for up to ten EX3300 member switches |
| EX4200 Virtual Chassis        | 9.0R1—Initial release. Support for up to ten EX4200 member switches                                                             |

**Table 564: Maximum Member Switch Support for Virtual Chassis by Junos OS Release (*continued*)**

| Maximum Member Switch Support                    | Initial Junos OS Release                                                                                                                                                                                                                                                                                                               |
|--------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EX4300 Virtual Chassis                           | <p>13.2X50-D10—Initial release. Support for up to ten EX4300 member switches</p> <p>13.2X50-D20—EX4300 switch support added for QFX Series Virtual Chassis and for VCF. See <a href="#">Understanding QFX Series Virtual Chassis Components</a> or <a href="#">“Understanding Virtual Chassis Fabric Components”</a> on page 5243.</p> |
| EX4500 Virtual Chassis                           | <p>11.1R1—Initial release. Support for up to two EX4500 switches</p> <p>11.4R1—Support for up to ten EX4500 member switches</p>                                                                                                                                                                                                        |
| EX4550 Virtual Chassis                           | 12.2R1—Initial release. Support for up to ten EX4550 switches                                                                                                                                                                                                                                                                          |
| EX4600 Virtual Chassis                           | 13.2X51-D25—Initial release. Support for up to ten EX4600 switches                                                                                                                                                                                                                                                                     |
| Mixed EX4200 and EX4500 Virtual Chassis          | <p>11.1R1—Initial release. Support for up to two EX4500 switches and up to eight EX4200 switches</p> <p>11.2R1—Support for up to nine EX4200 switches</p> <p>11.4R1—Support for up to nine EX4500 switches</p>                                                                                                                         |
| Mixed EX4200 and EX4550 Virtual Chassis          | 12.2R1—Initial release. Support for up to ten total EX4200 and EX4550 switches                                                                                                                                                                                                                                                         |
| Mixed EX4200, EX4500, and EX4550 Virtual Chassis | 12.2R1—Initial release. Support for up to ten total EX4200, EX4500, and EX4550 switches                                                                                                                                                                                                                                                |
| Mixed EX4300 and EX4600 Virtual Chassis          | 13.2X51-D25—Initial release. Support for up to ten total EX4300 and EX4600 switches. EX4600 switches must assume routing engine role.                                                                                                                                                                                                  |
| Mixed EX4500 and EX4550 Virtual Chassis          | 12.2R1—Initial release. Support for up to ten total EX4500 and EX4550 switches                                                                                                                                                                                                                                                         |
| EX9200 Virtual Chassis                           | 13.2R2—Initial release. Support for up to two EX9200 switches.                                                                                                                                                                                                                                                                         |

### Virtual Chassis Ports (VCPs)

You use Virtual Chassis ports (VCPs) to interconnect the member switches in a Virtual Chassis.

Some switches have dedicated VCPs. Dedicated VCPs allow you to interconnect switches without requiring any additional interface configuration.

These switches have dedicated VCPs:

- EX4200 switches, on the rear panel
- EX4500 switches, on the Virtual Chassis module
- EX4550 switches, on the Virtual Chassis module

To interconnect switches that do not have dedicated VCPs or to interconnect switches across greater distances than allowed by a dedicated-VCP connection, you configure an optical port as a VCP. You can configure those VCPs on these switches:

- EX2200 switches, through an uplink port



**NOTE:** All RJ-45 interfaces, including built-in network ports with 10/100/1000BASE-T Gigabit Ethernet connectors and 1000BASE-T RJ-45 transceivers, on EX2200 and EX2200-C switches, can also be configured into VCPs.

- EX3300 switches, through an uplink port



**NOTE:** Uplink ports 2 and 3 on EX3300 switches are configured as VCPs by default.

- EX4200 switches, through uplink module ports (SFP, SFP+, or XFP) or through an SFP+ port on the EX4200-24F switch
- EX4300 switches, through uplink ports



**NOTE:** All QSFP+ ports on an EX4300 switch are configured as VCPs by default.

- EX4500 switches, through any SFP+ port
- EX4550 switches, through any SFP+ port
- EX4600 switches, through SFP+ and QSFP+ ports.

All supported SFP, SFP+, and XFP uplink connections between EX4200, EX4500, and EX4550 switches can be configured as VCPs.

You can increase the Virtual Chassis bandwidth between member switches by configuring multiple optical ports connecting the switches as VCPs. The optical ports configured as VCPs automatically form a Link Aggregation Group (LAG) bundle. See [“Understanding EX Series Virtual Chassis Port Link Aggregation” on page 5083](#).

## Master Role

---

The member that functions in the master role in the Virtual Chassis:

- Manages the member switches.
- Runs Junos OS for EX Series switches in a master role.
- Runs the chassis management processes and control protocols.
- Represents all the member switches interconnected within the Virtual Chassis configuration. (The hostname and other properties that you assign to this switch during setup apply to all members of the Virtual Chassis configuration.)

When an EX Series switch that supports Virtual Chassis is powered on as a standalone switch, it is considered the master member. In a Virtual Chassis, one member functions as the master and a second member functions as the backup:

- In a preprovisioned configuration, one of the two members assigned as **routing-engine** functions as the master member. The selection of which member assigned as **routing-engine** functions as master and which as backup is determined by the software based on the master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 5078](#).
- In a configuration that is not preprovisioned, the selection of the master and backup is determined by the mastership priority value and secondary factors in the master election algorithm.

In a mixed EX4300 and EX4600 Virtual Chassis, an EX4600 switch must assume the master role.

In any mixed Virtual Chassis configuration that includes EX4200 switches, EX4500 switches, or EX4550 switches, any switch can be configured in any role in any configuration.

## Backup Role

---

The member that functions in the backup role in the Virtual Chassis:

- Maintains a state of readiness to take over the master role if the master fails.
- Runs Junos OS for EX Series switches in a backup role.
- Synchronizes with the master in terms of protocol states, forwarding tables, and so forth, so that it is prepared to preserve routing information and maintain network connectivity without disruption in case the master is unavailable.

You must have at least two member switches in the Virtual Chassis configuration in order to have a backup member.

- In a preprovisioned configuration, one of the two members assigned as **routing-engine** functions in the backup role. The selection of which member assigned as **routing-engine** functions as master and which as backup is determined by the software based on the

master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 5078](#).

- In a configuration that is not preprovisioned, the selection of the master and backup is determined by the mastership priority value and secondary factors in the master election algorithm.

In a mixed EX4300 and EX4600 Virtual Chassis, we strongly recommend configuring an EX4600 switch into the backup role. A mixed EX4300 and EX4600 Virtual Chassis must use an EX4600 member switch in the master role, and configuring an EX4600 switch into the backup role ensures that the Virtual Chassis remains up after a switchover event.

In any mixed Virtual Chassis configuration that includes EX4200 switches, EX4500 switches, or EX4550 switches, any switch can be configured in any role in any configuration.

### Linecard Role

A member that functions in the linecard role in the Virtual Chassis:

- Runs only a subset of Junos OS for EX Series switches.
- Does not run the chassis control protocols.
- Can detect certain error conditions (such as an unplugged cable) on any interfaces that have been configured on it through the master.

The Virtual Chassis configuration must have at least three members in order to include a linecard member.

- In a preprovisioned configuration, you can explicitly configure a member with the linecard role, which makes it ineligible for functioning as a master or backup.
- In a configuration that is not preprovisioned, the members that are not selected as master or backup function as linecard members of the Virtual Chassis configuration. The selection of the master and backup is determined by the mastership priority value and secondary factors in the master election algorithm. A switch with a mastership priority of 0 is always in the linecard role.

### Member Switch and Member ID

Each standalone EX Series switch that supports Virtual Chassis is a potential member of a Virtual Chassis configuration. When one of those switches is powered on, it receives a member ID that can be seen by viewing the front-panel LCD or by entering the **show virtual-chassis** command. If the switch is powered on as a standalone switch, that member's member ID is always 0. When the switch is interconnected with other switches in a Virtual Chassis configuration, its member ID is assigned by the master based on various factors, such as the order in which the switch was added to the Virtual Chassis configuration or the member ID assigned by a preprovisioned configuration. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 5078](#).

If the Virtual Chassis configuration previously included a member switch and that member was physically disconnected or removed from the Virtual Chassis configuration, its

member ID is not available for assignment as part of the standard sequential assignment by the master. For example, you might have a Virtual Chassis configuration composed of member 0, member 2, and member 3, because member 1 was removed. When you add another member switch and power it on, the master assigns it as member 4.

The member ID distinguishes the member switches from one another. You use the member ID:

- To assign a mastership priority value to a member switch
- To configure interfaces for a member switch (The function is similar to that of a slot number on Juniper Networks routers.)
- To apply some operational commands to a member switch
- To display status or characteristics of a member switch

### Mastership Priority

---

In a configuration that is not preprovisioned, you can designate the role (master, backup, or linecard) that a member switch assumes by configuring its mastership priority (from 0 through 255). The mastership priority value is the factor in the master election algorithm with the highest precedence for selecting the master of the Virtual Chassis configuration. A switch with a mastership priority of 0 never assumes the backup or master role.

The default value for mastership priority is 128 for EX2200, EX3300, EX4200, EX4500, EX4550, and EX4600 switches. When a standalone switch is powered on, it receives the default mastership priority value. Because it is the only member of the Virtual Chassis configuration, it is also the master. When you interconnect a standalone switch to an existing Virtual Chassis configuration (which implicitly includes its own master), we recommend that you explicitly configure the mastership priority of the members that you want to function as the master and backup.



**NOTE:** Configuring the same mastership priority value for both the master and backup helps to ensure a smooth transition from master to backup when the master becomes unavailable. It prevents the original master from preempting control from the backup when the backup has taken control of the Virtual Chassis configuration because the original master became unavailable.

In a preprovisioned configuration, you assign the role of each member switch. An EX9200 switch does not have a mastership priority and an EX9200 Virtual Chassis, therefore, must be preprovisioned.

### Mixed Virtual Chassis

---

EX4200 switches, EX4500 switches, and EX4550 switches can be interconnected into the same Virtual Chassis to form a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, mixed EX4500 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, and EX4550 Virtual Chassis. The mixed Virtual Chassis supports up to 10 member switches regardless of whether the switches are EX4200 switches,



EX4500 switches, or EX4550 switches. Any model of EX4200, EX4500, or EX4550 switch can be interconnected into the same mixed Virtual Chassis. The master election process that decides member switch roles in a mixed Virtual Chassis is identical to the master election process in a non-mixed Virtual Chassis, so any member switch in a mixed Virtual Chassis can assume the master, backup, or linecard role.

EX4300 and EX4600 switches can act as member switches in a mixed EX4300 and EX4600 Virtual Chassis. An EX4600 switch must assume the master role in a mixed EX4300 and EX4600 Virtual Chassis. We strongly recommend configuring EX4600 switches only into the Routing Engine role when you are configuring a mixed EX4300 and EX4600 Virtual Chassis to ensure only EX4600 switches assume the master role.

EX2200 and EX3300 switches cannot be a part of any mixed Virtual Chassis.

### Virtual Chassis Identifier (VCID)

All members of a Virtual Chassis configuration share one Virtual Chassis identifier (VCID). This identifier is derived from internal parameters. When you are monitoring a Virtual Chassis configuration, the VCID is displayed in certain interface views and is also part of the **show virtual-chassis** output.

#### Related Documentation

- [Understanding EX8200 Virtual Chassis Components](#)
- [EX Series Virtual Chassis Overview](#)
- [Understanding EX4300 Virtual Chassis on page 5065](#)
- [Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet](#)
- [Example: Configuring an EX4500 Virtual Chassis with a Master and Backup in a Single Wiring Closet](#)
- [Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File](#)
- [Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port \(CLI Procedure\) on page 5109](#)
- [Setting an Uplink Port as a Virtual Chassis Port on an EX4500 or EX4550 Switch \(CLI Procedure\)](#)

## Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric

This topic describes the requirements for a mixed Virtual Chassis or a mixed Virtual Chassis Fabric (VCF).

A mixed Virtual Chassis includes two or more types of EX Series switches, two or more types of QFX Series switches, or a mix of EX and QFX Series switches.

A mixed VCF is any VCF that includes two or more types of member switches. Because a VCF must use a QFX5100 switch as a spine device, a mixed VCF is any VCF that includes EX4300, QFX3500, or QFX3600 member switches in addition to the required QFX5100 switches.



**NOTE:** The optimal VCF topology is to use QFX5100 devices only. A VCF composed entirely of QFX5100 devices supports the largest breadth of features at the highest scalability while also supporting the highest number of high-speed interfaces.

This topic covers:

- [Virtual Chassis Fabric Summary on page 5074](#)
- [Understanding Mixed Virtual Chassis Fabric on page 5075](#)
- [Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches on page 5075](#)
- [Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches on page 5076](#)
- [Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis on page 5077](#)
- [Understanding Mixed EX4300 and EX4600 Virtual Chassis on page 5077](#)
- [Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis on page 5077](#)

### Virtual Chassis Fabric Summary

[Table 565 on page 5074](#) provides a high-level overview of the permitted hardware allowed in the routing engine and line card roles of a mixed and a non-mixed VCF. The table also includes license requirements and supported configuration methods.

**Table 565: Virtual Chassis Fabric Summary**

| Category  | Allowed Routing Engines | Allowed Line Cards | License Requirement                                                               | Configuration Methods                                                    |
|-----------|-------------------------|--------------------|-----------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| Non-mixed | QFX5100                 | QFX5100            | Yes (on two QFX5100 switches operating in master and backup Routing Engine roles) | Autoprovisioning<br>Preprovisioning<br>Nonprovisioning (not recommended) |

Table 565: Virtual Chassis Fabric Summary (*continued*)

| Category | Allowed Routing Engines | Allowed Line Cards                      | License Requirement                                                               | Configuration Methods                                                    |
|----------|-------------------------|-----------------------------------------|-----------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| Mixed    | QFX5100                 | QFX5100<br>QFX3600<br>QFX3500<br>EX4300 | Yes (on two QFX5100 switches operating in master and backup Routing Engine roles) | Autoprovisioning<br>Preprovisioning<br>Nonprovisioning (not recommended) |

### Understanding Mixed Virtual Chassis Fabric

A VCF must use a QFX5100 switch in the spine role. A mixed VCF is, therefore, any VCF that includes EX4300, QFX3500, or QFX3600 member switches in addition to the required QFX5100 switch.

The optimal method of configuring a VCF is to use QFX5100 devices only. A non-mixed VCF composed entirely of QFX5100 devices supports the largest breadth of features at the highest scalability while also supporting the highest number of high-speed interfaces. You can, however, also configure a mixed VCF.

If you use QFX3600, QFX3500, or EX4300 devices as leaf devices in your VCF, you must configure all devices in your VCF into mixed mode. If you are turning a non-mixed VCF into a mixed VCF, you have to reboot the VCF to change the mixed mode setting.

### Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches

[Table 566 on page 5076](#) provides a high-level overview of the permitted hardware allowed in the routing engine and line card roles of a mixed and a non-mixed Virtual Chassis for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 switches. The table also includes license requirements and supported configuration methods.

Table 566: Virtual Chassis Summary

| Category  | Allowed Routing Engines | Allowed Line Cards                      | License Requirement | Configuration Methods              |
|-----------|-------------------------|-----------------------------------------|---------------------|------------------------------------|
| Non-mixed | QFX5100                 | QFX5100                                 | No                  | Nonprovisioning<br>Preprovisioning |
|           | QFX3600<br>QFX3500      | QFX3600<br>QFX3500                      | No                  | Nonprovisioning<br>Preprovisioning |
|           | EX4600                  | EX4600                                  | No                  | Nonprovisioning<br>Preprovisioning |
|           | EX4300                  | EX4300                                  | No                  | Nonprovisioning<br>Preprovisioning |
| Mixed     | QFX5100                 | QFX5100<br>QFX3600<br>QFX3500<br>EX4300 | No                  | Nonprovisioning<br>Preprovisioning |
|           | QFX3600<br>QFX3500      | QFX3600<br>QFX3500<br>EX4300            | No                  | Nonprovisioning<br>Preprovisioning |
|           | EX4600                  | EX4600<br>EX4300                        | No                  | Nonprovisioning<br>Preprovisioning |

#### Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches

In a mixed Virtual Chassis, the switch in the master Routing Engine role determines which switches are supported in the line card role of the mixed Virtual Chassis.

When a mixed Virtual Chassis is using a QFX5100 switch in the master Routing Engine role, you can use QFX5100, QFX3600, QFX3500, or EX4300 switches in the line card role.

When a mixed Virtual Chassis is using a QFX3600 or QFX3500 switch in the master Routing Engine role, you can use QFX3600, QFX3500, or EX4300 switches in the line card role.

In a mixed EX4300 and EX4600 Virtual Chassis, an EX4600 switch automatically assumes the Routing Engine role.

EX4600 switches can only be in a mixed Virtual Chassis with EX4300 switches. EX4600 switches cannot be in a mixed Virtual Chassis with QFX5100, QFX3600, or QFX3500 switches.

We recommend always configuring the same type of switch into the master and backup Routing Engine role, to ensure that the switch operating in the master role remains the same type of switch in the event of a switchover.

In most mixed Virtual Chassis, you must configure your Virtual Chassis to ensure a switch that supports the master Routing Engine assumes the master Routing Engine role. Without user configuration, any switch—with the exception of the EX4300 switch, which can never assume the master or backup Routing Engine role in a mixed Virtual Chassis or VCF—can assume the master or backup Routing Engine role.

### **Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis**

Up to ten EX4300 switches, QFX3500 switches, QFX3600 switches, and QFX5100 switches can be interconnected using Virtual Chassis ports (VCPs) to form a mixed or non-mixed Virtual Chassis. The mixed Virtual Chassis supports up to ten member switches regardless of the switches that compose the mixed Virtual Chassis.

EX4300 switches can also be interconnected into a mixed Virtual Chassis with EX4600 switches. See the following section for information on mixed EX4300 and EX4600 Virtual Chassis.

### **Understanding Mixed EX4300 and EX4600 Virtual Chassis**

EX4300 switches and EX4600 switches can be interconnected into the same Virtual Chassis. An EX4600 switch automatically assumes the master Routing Engine role in a mixed EX4300 and EX4600 Virtual Chassis, since EX4300 switches cannot assume the Routing Engine role in a mixed Virtual Chassis. EX4600 switches cannot be in a mixed Virtual Chassis with any other type of switch.

The mixed Virtual Chassis supports up to ten member switches.

### **Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis**

EX4200 switches, EX4500 switches, and EX4550 switches can be interconnected into the same Virtual Chassis to form a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, mixed EX4500 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, and EX4550 Virtual Chassis. The mixed Virtual Chassis supports up to 10 member switches regardless of whether the switches are EX4200 switches, EX4500 switches, or EX4550 switches. Any model of EX4200, EX4500, or EX4550 switch can be interconnected into the same mixed Virtual Chassis. The master election process that decides member switch roles in a mixed Virtual Chassis is identical to the master election process in a non-mixed Virtual Chassis, so any member switch in a mixed Virtual Chassis can assume the master, backup, or linecard role.

EX4200 switches, EX4500 switches, and EX4550 switches cannot be interconnected into a Virtual Chassis with any other switches.

#### **Related Documentation**

- [Virtual Chassis Fabric Overview on page 5241](#)
- [Understanding QFX Series Virtual Chassis](#)
- [EX Series Virtual Chassis Overview](#)
- [Understanding Virtual Chassis Fabric Components on page 5243](#)
- [Understanding QFX Series Virtual Chassis Components](#)
- [Understanding EX Series Virtual Chassis Components on page 5067](#)

## Understanding How the Master in a Virtual Chassis Is Elected



**NOTE:** This topic does not apply to EX8200 Virtual Chassis. See *EX8200 Virtual Chassis Overview*.

All switches that are interconnected in a Virtual Chassis configuration are member switches of that Virtual Chassis. Each Virtual Chassis configuration has one member that functions as the *master* and controls the Virtual Chassis configuration.

When a Virtual Chassis configuration boots, the Juniper Networks Junos operating system (Junos OS) on the switches automatically runs a master election algorithm to determine which member switch assumes the role of master.

The algorithm proceeds from the top condition downward until the stated condition is satisfied:

1. Choose the member with the highest user-configured mastership priority (255 is the highest possible value). A switch with a mastership priority of 0 will always stay in the linecard role.
2. Choose the member that was master the last time the Virtual Chassis configuration booted.
3. Choose the member that has been included in the Virtual Chassis configuration for the longest period of time. (For this to be a deciding factor, there has to be a minimum time lapse of 1 minute between the power-ons of the individual interconnected member switches.)
4. Choose the member with the lowest MAC address.

The variations among switches and switch models do not impact the master election algorithm.

To ensure that a specific member is elected as the master:

1. Power on only the switch that you want to configure as master of the Virtual Chassis configuration.
2. Configure the mastership priority of that member to have the highest possible value (255).
3. Continue to configure other members through the master member.
4. Power on the other members.

You can also specify the switch roles by preprovisioning your Virtual Chassis. Preprovisioning a Virtual Chassis allows you to manually assign the member ID and role for each switch in the Virtual Chassis. See *Configuring an EX3300 Virtual Chassis (CLI Procedure)*, *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*, or *Configuring a QFX Series Virtual Chassis (CLI Procedure)*.

#### Related Documentation

- [EX8200 Virtual Chassis Overview](#)
- [EX Series Virtual Chassis Overview](#)
- [Understanding QFX Series Virtual Chassis](#)
- [Understanding EX Series Virtual Chassis Components on page 5067](#)
- [Understanding QFX Series Virtual Chassis Components](#)
- [Understanding EX Series Virtual Chassis Configuration on page 5086](#)

## Understanding Software Upgrade in an EX Series Virtual Chassis



**NOTE:** This topic applies to all EX Series Virtual Chassis except EX8200 Virtual Chassis. See *Understanding Software Upgrades in an EX8200 Virtual Chassis* for information about EX8200 Virtual Chassis.

In a Virtual Chassis, each member switch must be running the same version of Juniper Networks Junos operating system (Junos OS).

You can install a new Junos OS release on the entire Virtual Chassis or on a particular member in the Virtual Chassis by using the same CLI command that you use to install Junos OS on standalone switches—the **request system software add** command.

You can use the automatic software update feature to automatically update the Junos OS version on member switches as you add them to a Virtual Chassis. See “[Understanding Automatic Software Update on Virtual Chassis Member Switches](#)” on page 5090. If you are not configuring the automatic software update feature, we recommend that you update the new member switch to the version of Junos OS running on the Virtual Chassis before adding the member switch to the Virtual Chassis.

In any mixed Virtual Chassis, the member switches must be running the same version of Junos OS. You can upgrade all member switches simultaneously by specifying a path to multiple Junos OS images in the same **request system software add** command. Multiple Junos OS images are needed to upgrade a mixed Virtual Chassis because an EX4200 switch runs a different version of Junos OS than an EX4500 or EX4550 switch. See *Installing Software on a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Switches (CLI Procedure)*.



**NOTE:** EX4500 and EX4550 switches run the same Junos OS image. You do not need to use multiple Junos OS images when updating a mixed EX4500 and EX4550 Virtual Chassis.

You can also use nonstop software upgrade (NSSU) to upgrade Junos OS on all members. NSSU provides an orderly upgrade of each member of the Virtual Chassis and takes advantage of graceful Routing Engine switchover, nonstop active routing, and link aggregation to minimize traffic disruption during the upgrade. For more information about NSSU, see “[Understanding Nonstop Software Upgrade on EX Series Switches](#)” on page 2488.

**Related  
Documentation**

- [Understanding Software Upgrades in an EX8200 Virtual Chassis](#)
- [Understanding EX Series Virtual Chassis Components on page 5067](#)
- [Understanding Automatic Software Update on Virtual Chassis Member Switches on page 5090](#)
- [Installing Software on an EX Series Switch with a Single Routing Engine \(CLI Procedure\)](#)
- [Upgrading Software on an EX3300, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and Mixed Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 2508](#)

## Understanding Global Management of a Virtual Chassis

---



**NOTE:** This topic does not apply to EX8200 Virtual Chassis. See [Understanding Global Management of an EX8200 Virtual Chassis](#).


---

A Virtual Chassis is composed of multiple switches, and it, therefore, has multiple console ports and multiple out-of-band management Ethernet ports located on the switches.

You can connect a PC or laptop directly to a console port of any member switch to set up and configure the Virtual Chassis. When you connect to the console port of any member switch, the console session is redirected to the master switch, as shown in [Figure 79 on page 5081](#).



Figure 79: Console Session Redirection (EX4200 Virtual Chassis Pictured)




If the master becomes unavailable, the console session is disconnected from the old master and a new session is established with the newly elected master.

An out-of-band management Ethernet port is often referred to simply as a management Ethernet port. It uses a dedicated management channel for device maintenance and allows a system administrator to monitor and manage the switch by remote control.

The Virtual Chassis configuration can be managed remotely through SSH or Telnet using a global management interface called the virtual management Ethernet (VME) interface. The VME interface is a logical interface representing all of the out-of-band management ports on the member switches. When you connect to the Virtual Chassis configuration using the VME interface's IP address, the connection is redirected to the master member as shown in [Figure 80 on page 5082](#).

Figure 80: Management Ethernet Port Redirection to the VME Interface



If the master management Ethernet link is unavailable, the session is redirected through the backup management Ethernet link. If there is no active management Ethernet link on the backup, the VME interface chooses a management Ethernet link on one of the linecard members, selecting the linecard member with the lowest member ID as its first choice.

You can configure an IP address for the VME global management interface at any time.

You can perform remote configuration and administration of all members of the Virtual Chassis configuration through the VME interface.

#### Related Documentation

- *Understanding Global Management of an EX8200 Virtual Chassis*
- [Understanding EX Series Virtual Chassis Components on page 5067](#)
- *Understanding QFX Series Virtual Chassis Components*

- *Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet*
- *Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)*

## Understanding Nonvolatile Storage in a Virtual Chassis



**NOTE:** This topic applies to all EX Series Virtual Chassis except EX8200 Virtual Chassis. See *Understanding File Storage in an EX8200 Virtual Chassis* for information about EX8200 Virtual Chassis.

The EX Series or QFX Series switches store the Juniper Networks Junos operating system (Junos OS) system files in internal flash memory. In the Virtual Chassis configurations, both the master and the backup switch store the configuration information for all the member switches.

- [Nonvolatile Memory Features on page 5083](#)

### Nonvolatile Memory Features

Junos OS optimizes the way the Virtual Chassis stores its configuration if a member switch or the Virtual Chassis configuration is shut down improperly:

- If the master is not available, the backup switch takes on the role of the master and its internal flash memory takes over as the alternate location for maintaining nonvolatile configuration memory.
- If a member switch is taken offline for repair, the master stores the configuration of the member switch.

#### Related Documentation

- [Understanding File Storage in an EX8200 Virtual Chassis](#)
- [Understanding QFX Series Virtual Chassis](#)
- [Command Forwarding Usage with an EX Series Virtual Chassis on page 5149](#)

## Understanding EX Series Virtual Chassis Port Link Aggregation



**NOTE:** This topic applies to all EX Series Virtual Chassis except EX8200 Virtual Chassis. See *Understanding Virtual Chassis Port Link Aggregation in an EX8200 Virtual Chassis* for information about EX8200 Virtual Chassis.

You can combine physical Ethernet ports belonging to different member switches of a Virtual Chassis configuration to form a logical point-to-point link, known as a *link aggregation group (LAG)* or *bundle*. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation provides network redundancy by

load-balancing traffic across all available links. If one of the links fails, the system automatically load-balances traffic across all remaining links.

Similarly, if a Virtual Chassis member switch that has LAG member interfaces on multiple member switches fails for any reason, the traffic traversing the LAG can be redirected through the active member switch. This setup has benefits for failover purposes and can be especially beneficial in cases when a member switch needs to be inactive for some time, such as during a software upgrade using NSSU.

You can configure any optical uplink port that can be used to connect EX2200, EX3300, EX4200, EX4300, EX4500, or EX4550 switches together into a Virtual Chassis port (VCP). You can configure multiple optical uplink interfaces between two member switches in the same Virtual Chassis as VCPs. If you have configured two or more optical ports as VCPs connecting the same member switches, the optical uplink ports configured as VCPs automatically form a LAG provided the optical uplink ports are configured to operate at the same link speeds. Each LAG is assigned a positive-integer identifier called a *trunk ID*.

On EX2200 and EX2200-C switches only, you can also configure the RJ-45 interfaces, including built-in network ports with 10/100/1000BASE-T Gigabit Ethernet connectors and 1000BASE-T RJ-45 transceivers, into VCPs. On EX2200 and EX2200-C switches, a LAG that includes up to 8 interfaces configured as VCPs automatically forms. The LAG bundles includes all interfaces configured as VCPs, regardless of whether the interfaces are optical transceiver interfaces, RJ-45 transceiver interfaces, or built-in network ports with 10/100/1000BASE-T Gigabit Ethernet connectors.

You can create an optical VCP LAG connecting any two member switches in any Virtual Chassis, including VCP LAG connections interconnecting different switch models in a mixed Virtual Chassis.

On an EX2200 switch only, you can also configure all RJ-45 interfaces, including built-in network ports with 10/100/1000BASE-T Gigabit Ethernet connectors and 1000BASE-T RJ-45 transceivers, on EX2200 and EX2200-C switches, into VCPs. The RJ-45 interfaces also automatically form a LAG when configured into VCPs.

[Table 567 on page 5084](#) provides the maximum member link limit for each optical VCP LAG.

**Table 567: Maximum Member Links in LAGs Over Optical Interface VCPs**

| Member Switch 1 | Member Switch 2 | Maximum Member Links in VCP LAG |
|-----------------|-----------------|---------------------------------|
| EX2200 Switch   | EX2200 Switch   | 8                               |
| EX3300 Switch   | EX3300 Switch   | 8                               |
| EX4200 Switch   | EX4200 Switch   | 8                               |
| EX4200 Switch   | EX4500 Switch   | 8                               |
| EX4200 Switch   | EX4550 Switch   | 8                               |
| EX4500 Switch   | EX4500 Switch   | 8                               |

**Table 567: Maximum Member Links in LAGs Over Optical Interface VCPs (continued)**

| Member Switch 1 | Member Switch 2 | Maximum Member Links in VCP LAG |
|-----------------|-----------------|---------------------------------|
| EX4500 Switch   | EX4550 Switch   | 8                               |
| EX4550 Switch   | EX4550 Switch   | 8                               |

A LAG over uplink VCPs provides higher overall bandwidth for forwarding traffic between the member switches connected by the optical VCPs, faster management communications, and greater redundancy of operations among the members than would be available without the LAG. A LAG over uplink VCPs provides an additional Virtual Chassis link throughput for the switches.

See [“Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port \(CLI Procedure\)” on page 5109](#) for information about configuring uplink ports as VCPs.



**NOTE:** The interfaces that are included within a bundle or LAG are sometimes referred to as *member interfaces*. Do not confuse this term with *member switches*, which refers to switches that are interconnected as a Virtual Chassis. It is possible to create a LAG that is composed of member interfaces that are located in different member switches of a Virtual Chassis.

#### Related Documentation

- [Understanding Virtual Chassis Port Link Aggregation in an EX8200 Virtual Chassis](#)
- [EX Series Virtual Chassis Overview](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2582](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch](#)
- [Example: Configuring an EX4200 Virtual Chassis Interconnected Across Multiple Wiring Closets](#)
- [Example: Connecting EX4500 Member Switches in a Virtual Chassis Across Wiring Closets](#)
- [Example: Configuring Link Aggregation Groups Using EX4200 Uplink Virtual Chassis Ports](#)

## Understanding EX Series Virtual Chassis Configuration



**NOTE:** This topic applies to all EX Series Virtual Chassis except EX8200 Virtual Chassis. See *Understanding EX8200 Virtual Chassis Components* for information about EX8200 Virtual Chassis.

You configure and manage almost all aspects of an EX Series Virtual Chassis configuration through the master switch of the Virtual Chassis. However, you can also configure Virtual Chassis parameters when a switch is a standalone switch not interconnected with other members.

EX Series switches that support Virtual Chassis have some innate characteristics of a Virtual Chassis by default. A standalone switch is assigned member ID 0 and is the master of itself. Therefore, you can edit its Virtual Chassis configuration. When the standalone switch is interconnected with an existing Virtual Chassis configuration, the Virtual Chassis configuration statements and any uplink Virtual Chassis port (VCP) settings that you previously specified on the standalone switch remain part of its configuration.

A switch is not recognized by the Virtual Chassis as a member switch until it is interconnected with the master or interconnected with an existing member of the Virtual Chassis.

EX2200, EX3300, and EX4300 switches do not have dedicated VCPs. These switches form a Virtual Chassis by configuring optical ports connecting the member switches into VCPs. You can expand the uplink connections on these switches over long distances to connect switches at different locations into the same Virtual Chassis.



**NOTE:** Uplink ports 2 and 3 on any EX3300 switch are configured, by default, as VCPs. You can change this default configuration or configure another uplink port as a VCP on an EX3300 switch.

QSFP+ ports on an EX4300 switch are also configured, by default, as VCPs.

On all other EX Series switches, the optical ports are not configured as VCPs by default.

When an EX4200, EX4500, or EX4550 switch is located too far away to be interconnected through dedicated VCPs, you can configure an optical port as a VCP.

A link aggregation group (LAG) will be formed automatically when the new switch is added to the configuration if more than one such link with the same speed is detected between optical VCPs on the new member and an existing member. See [“Understanding EX Series Virtual Chassis Port Link Aggregation” on page 5083](#).

When an optical port is set as a VCP, it cannot be used for any other purpose. If you want to use the optical port for another purpose, you must delete the VCP setting. You can

execute this command directly on the member whose uplink VCP setting you want to delete or through the master of the Virtual Chassis configuration.



**CAUTION:** Deleting a VCP in a Virtual Chassis configuration can cause the Virtual Chassis configuration to split. For more information, see [“Understanding Split and Merge in a Virtual Chassis” on page 5087](#).

You can create a preprovisioned configuration. This type of configuration allows you to deterministically control the member ID and role assigned to a member switch by associating the switch with its serial number. For an example of a preprovisioned configuration, see *Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File*.



**NOTE:** If a switch is interconnected with other switches in a Virtual Chassis configuration, each individual switch that is included as a member of the configuration is identified with a member ID. The member ID functions as an FPC slot number. When you are configuring interfaces for a Virtual Chassis configuration, you specify the appropriate member ID as the *slot* element of the interface name.

The default factory settings for a Virtual Chassis configuration include FPC 0 as a member of the default VLAN because FPC 0 is configured as part of the `ethernet-switching` family. To include the FPC in the default VLAN, add the `ethernet-switching` family to the configurations for those interfaces.

#### Related Documentation

- [Understanding EX8200 Virtual Chassis Components](#)
- [Understanding EX Series Virtual Chassis Components on page 5067](#)
- [Understanding How the Master in a Virtual Chassis Is Elected on page 5078](#)
- [Configuring an EX3300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)

## Understanding Split and Merge in a Virtual Chassis

In a Virtual Chassis, two or more switches are connected together to form a unit that is managed as a single chassis. If there is a disruption to the Virtual Chassis configuration due to member switches failing or being removed from the configuration, the Virtual Chassis configuration splits into two separate Virtual Chassis. This situation could cause disruptions in the network if the two separate configurations share common resources, such as global IP addresses. The split and merge feature provides a method to prevent the separate Virtual Chassis configurations from adversely affecting the network and also allows the two parts to merge back into a single Virtual Chassis configuration.



**NOTE:** If a Virtual Chassis configuration splits into separate parts, we recommend that you resolve the problem that caused the Virtual Chassis configuration to split as soon as possible.

You can also use this feature to merge two active but separate Virtual Chassis that have not previously been part of the same configuration into one Virtual Chassis configuration.



**NOTE:** The split and merge feature is enabled by default on EX Series and QFX Series Virtual Chassis. You can disable the split and merge feature by using the `set virtual-chassis no-split-detection` command.

This topic describes:

- [What Happens When a Virtual Chassis Configuration Splits on page 5088](#)
- [Merging Virtual Chassis Configurations on page 5089](#)

### [What Happens When a Virtual Chassis Configuration Splits](#)

---

When a Virtual Chassis configuration splits into two separate Virtual Chassis configurations, the individual member switches detect this topology change and run the master election algorithm to select a new master for each of the two Virtual Chassis configurations. The new masters then determine whether their Virtual Chassis configuration remains active. One of the configurations remains active based on the following:

- It contains both the stable master and the stable backup (that is, the master and backup from the original Virtual Chassis configuration before the split).
- It contains the stable master and the configuration is greater than half the Virtual Chassis size.
- It contains the stable backup and is at least half the Virtual Chassis size.

In accordance with the rules given in the second and third list items, if the Virtual Chassis configuration splits into two equal parts and the stable master and stable backup are in different parts, then the part that contains the stable backup becomes active.



**NOTE:** The number of members in the Virtual Chassis configuration includes all member switches connected to date minus the number whose Virtual Chassis member IDs have been recycled (that is, made available for reassignment). Therefore, the size of the Virtual Chassis configuration increases when a new member switch is detected and decreases when a member switch's ID is recycled.

These rules ensure that only one of the two separate Virtual Chassis configurations created by the split remains active. The member switches in the inactive Virtual Chassis



configuration remain in a linecard role. For the inactive members to become active again, one of the following things must happen:

- The problem that caused the original Virtual Chassis configuration to split is resolved, allowing the two Virtual Chassis configurations to merge.
- You load the factory default configuration on the inactive members, which causes the inactive members to function as standalone switches or become part of a different Virtual Chassis configuration.



**NOTE:** When you remove a member switch from a Virtual Chassis configuration, we recommend that you recycle the member ID using the `request virtual-chassis recycle` command.

### Merging Virtual Chassis Configurations

There are two scenarios in which separate Virtual Chassis merge:

- A Virtual Chassis configuration that had split into two is now merging back into a single configuration because the problem that had caused it to split has been resolved.
- You want to merge two Virtual Chassis that had not previously been configured together.

Every Virtual Chassis configuration has a unique ID (VCID) that is automatically assigned when the Virtual Chassis configuration is formed. You can also explicitly assign a VCID using the `set virtual-chassis id` command. A VCID that you assign takes precedence over automatically assigned VCIDs.

When you reconnect the separate Virtual Chassis configurations or connect them for the first time, the members determine whether or not the separate Virtual Chassis configurations can merge. The members use the following rules to determine whether a merge is possible:

- If the Virtual Chassis configurations have the same VCID, then the configurations can merge. If the two Virtual Chassis were formed as the result of a split, they have the same VCID.
- If the VCIDs are different, then the two configurations can merge only if both are active (inactive configurations cannot merge, ensuring that members removed from one Virtual Chassis configuration do not become members of another Virtual Chassis configuration). If the configurations to merge are both active and one of them has a user-configured VCID, this ID becomes the ID of the merged Virtual Chassis. If neither Virtual Chassis has a user-configured VCID, then the VCID of the configuration with the highest mastership priority becomes the ID of the merged Virtual Chassis. The resulting merged Virtual Chassis configuration is active.

When you connect two Virtual Chassis configurations, the following events occur:

1. Connecting the two split Virtual Chassis configurations triggers the shortest-path-first (SPF) algorithm. The SPF algorithm computes the network topology and then triggers the master election algorithm. The master election algorithm waits for the members to synchronize the topology information before running.
2. The master election algorithm merges the VCIDs of all the members.
3. Each member runs the master election algorithm to select a master and a backup from among all members with the same VCIDs. For more information, see [“Understanding How the Master in a Virtual Chassis Is Elected”](#) on page 5078.
4. The master determines whether the Virtual Chassis configuration is active or inactive. (See [“What Happens When a Virtual Chassis Configuration Splits”](#) on page 5088.)
5. If the Virtual Chassis configuration is active, the master assigns roles to all members. If the Virtual Chassis configuration is inactive, the master assigns all members the role of linecard.
6. When the other members receive their role from the master, they change their role to backup or linecard. They also use the active or inactive state information sent by the master to set their own state to active or inactive and to construct the Virtual Chassis member list from the information sent by the master.
7. If the Virtual Chassis state is active, the master waits for messages from the members indicating that they have changed their roles to the assigned roles, and then the master changes its own role to master.



**NOTE:** When you merge two Virtual Chassis that had not previously been part of the same Virtual Chassis configuration, any configuration settings (such as the settings for Telnet and FTP services, graceful Routing Engine switchover (GRES), fast failover, VLANs, and so on) that exist on the new master become the configuration settings for all members of the new Virtual Chassis, overwriting any other configuration settings.

---

**Related  
Documentation**

- [Disabling Split and Merge in a Virtual Chassis \(CLI Procedure\)](#) on page 5114
- [Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge \(CLI Procedure\)](#) on page 5116
- [Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge](#)
- [Understanding EX Series Virtual Chassis Configuration](#) on page 5086
- [Understanding QFX Series Virtual Chassis](#)

## Understanding Automatic Software Update on Virtual Chassis Member Switches

You can use the automatic software update feature to automatically update the Juniper Networks Junos operating system (Junos OS) version on prospective member switches as you add them to an EX Series or QFX Series Virtual Chassis.

This topic includes:

- [Automatic Software Update Basics on page 5091](#)
- [Automatic Software Update Restrictions on page 5091](#)

### Automatic Software Update Basics

---

When you have configured automatic software update on a Virtual Chassis, the Junos OS version is updated on the new member switch when you add it to the Virtual Chassis. The new member switch immediately joins the Virtual Chassis configuration and is put in the active state.

For a standalone switch to join an existing Virtual Chassis, it must be running the same version of Junos OS that is running on the Virtual Chassis master. When the master in a Virtual Chassis detects that a new switch has been added to the configuration, it checks the software version on the new switch. If the software version on the new switch is not the same as the version running on the master, the master keeps the new switch in the inactive state. If you have not enabled the automatic software update feature, you have to manually install the correct software version on each prospective member switch as it is added to the Virtual Chassis.

### Automatic Software Update Restrictions

---

You cannot use automatic software update in certain scenarios, and you must ensure that the software release version on the Virtual Chassis is supported by the release on the prospective member switch.

You cannot use the automatic software update feature to update software for a prospective member switch in the following scenarios:

- The Virtual Chassis was preprovisioned and is running Junos OS Release 10.4R2 or earlier.
- You configured the **mastership-priority** command to manually configure the mastership priority of at least one Virtual Chassis member switch and the Virtual Chassis was running Junos OS Release 10.4R2 or earlier when you committed this configuration.
- The Junos OS versions on the Virtual Chassis and the prospective member switch are different versions of the same major Junos OS release. For instance, if a Virtual Chassis is running Junos OS Release 10.4R1, the prospective member switch cannot be updated using automatic software update if it is running Junos OS Release 10.4R2, 10.4R3, or any other Junos OS Release 10.4 release version.

The automatic software update feature also has a Junos OS release dependency between the release that is already running on the Virtual Chassis and the release that is running on the prospective member switch.

[Table 568 on page 5092](#) summarizes automatic software update support for each Junos OS release combination.

Table 568: Automatic Software Update Support

| Virtual Chassis Junos OS Release                | Supported Junos OS Releases for Prospective Member Switches                                                                                                                                                        |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| All versions of Junos OS 9.0 through 9.6        | All versions of Junos OS 9.0 through 9.6<br>Junos OS Releases 10.0R1 through 10.0R4<br>All versions of Junos OS Release 10.1<br>Junos OS Releases 10.2R1 through 10.2R3<br>Junos OS Releases 10.3R1 through 10.3R3 |
| Junos OS Releases 10.0R1 through 10.0R4         | All versions of Junos OS 9.0 through 9.6<br>All versions of Junos OS Release 10.1<br>Junos OS Releases 10.2R1 through 10.2R3<br>Junos OS Releases 10.3R1 through 10.3R3                                            |
| Junos OS Release 10.0R5 and later 10.0 releases | Junos OS Release 10.2R4 and later 10.2 releases<br>Junos OS Release 10.3R4 and later 10.3 releases<br>All versions of Junos OS Release 10.4<br>All versions of Junos OS Release 11.1                               |
| All versions of Junos OS Release 10.1           | All versions of Junos OS 9.0 through 9.6<br>Junos OS Releases 10.0R1 through 10.0R4<br>Junos OS Releases 10.2R1 through 10.2R3<br>Junos OS Releases 10.3R1 through 10.3R3                                          |
| Junos OS Releases 10.2R1 through 10.2R3         | All versions of Junos OS 9.0 through 9.6<br>Junos OS Releases 10.0R1 through 10.0R4<br>All versions of Junos OS Release 10.1<br>Junos OS Releases 10.3R1 through 10.3R3                                            |
| Junos OS Release 10.2R4 and later 10.2 releases | Junos OS Release 10.0R5<br>Junos OS Release 10.3R4 and later 10.3 releases<br>All versions of Junos OS Release 10.4<br>All versions of Junos OS Release 11.1                                                       |
| Junos OS Releases 10.3R1 through 10.3R3         | All versions of Junos OS 9.0 through 9.6<br>Junos OS Releases 10.0R1 through 10.0R4<br>All versions of Junos OS Release 10.1<br>Junos OS Releases 10.2R1 through 10.2R3                                            |
| Junos OS Release 10.3R4 and later 10.3 releases | Junos OS Release 10.0R5<br>All versions of Junos OS Release 10.4<br>All versions of Junos OS Release 11.1                                                                                                          |
| Junos OS Releases 10.4R1 through 10.4R3         | All versions of Junos OS 9.0 through 9.6<br>Junos OS Releases 10.0R1 through 10.0R4<br>All versions of Junos OS Release 10.1<br>Junos OS Releases 10.2R1 through 10.2R3<br>Junos OS Releases 10.3R1 through 10.3R3 |
| Junos OS Release 10.4R4 and later 10.4 releases | Junos OS Release 10.0R5<br>Junos OS Release 10.2R4 and later 10.2 releases<br>Junos OS Release 10.3R4 and later 10.3 releases<br>All versions of Junos OS Release 11.1                                             |

**Table 568: Automatic Software Update Support (*continued*)**

| Virtual Chassis Junos OS Release                    | Supported Junos OS Releases for Prospective Member Switches                                                                                                                        |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Junos OS Release 11.1R1                             | All versions of Junos OS Release 10.4<br>Junos OS Release 11.2 and later Junos OS releases                                                                                         |
| Junos OS Release 11.1R2 and later Junos OS releases | Junos OS Release 10.0R5<br>Junos OS Release 10.2R4 and later 10.2 releases<br>Junos OS Release 10.3R4 and later 10.3 releases<br>Junos OS Release 11.2 and later Junos OS releases |

**Related Documentation**

- [Understanding Software Upgrade in an EX Series Virtual Chassis on page 5079](#)
- [Understanding Software Upgrades in a QFX Series Virtual Chassis](#)
- [Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches](#)
- [Configuring Automatic Software Update on Virtual Chassis Member Switches \(CLI Procedure\) on page 5114](#)

**Understanding MAC Address Assignment on a Virtual Chassis**

In a Virtual Chassis, multiple switches—each with its own set of interfaces with unique MAC addresses—are connected together to form one chassis that can be managed as a single switch. The MAC address assigned to each network-facing interface on the switch changes when the switch joins a Virtual Chassis. Because all Layer 2 traffic decisions are based on an interface's MAC address, understanding MAC address assignment is important to understanding how network traffic is forwarded and received by the Virtual Chassis. For additional information about how a network uses MAC addresses to forward and receive traffic, see “[Understanding Bridging and VLANs on EX Series Switches](#)” on [page 2245](#).

When a Virtual Chassis is formed, the MAC address of the switch in the master role becomes the system MAC base address. The Virtual Chassis assigns the system MAC base address as the MAC address for all Layer 3 interfaces within the Virtual Chassis. The Virtual Chassis also assigns the system MAC base address to the virtual management Ethernet (VME) interface and to all of the virtual LANs (VLANs) in the Virtual Chassis.

The system MAC base address does not change in the event of a switchover if the switch that was originally configured in the master role remains a member of the Virtual Chassis. If the switch that was originally configured in the master role is removed from the Virtual Chassis, the MAC address of the current member switch in the master role is assigned as the system MAC base address after the MAC persistence timer interval has expired. You can configure the MAC persistence timer interval.

For Layer 2 and aggregated Ethernet interfaces, the Virtual Chassis assigns a unique MAC address that is derived from the member switch MAC address to each interface. The assignment of a unique MAC address to each network interface helps ensure that functions that require MAC address differentiation—such as redundant trunk groups

(RTGs), Link Aggregation Control Protocol (LACP), and general monitoring functions—can function properly.



**NOTE:** Unique MAC address assignment for Layer 2 and aggregated Ethernet interfaces in a Virtual Chassis was introduced in Junos OS Release 11.3. The same MAC address could be assigned to interfaces on different member switches in the same Virtual Chassis prior to this release.

If you reconfigure a Layer 2 interface into a Layer 3 interface, or the reverse, within a Virtual Chassis, the MAC address of that interface changes accordingly.

MAC addresses are assigned to interfaces in a Virtual Chassis automatically—no user configuration is possible or required. You can view the MAC addresses that are assigned to the interfaces by using the **show interfaces** command.

**Related  
Documentation**

- *Understanding MAC Address Assignment in an EX Series Switch*
- [Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis \(CLI Procedure\) on page 5113](#)
- *EX Series Virtual Chassis Overview*
- *EX8200 Virtual Chassis Overview*
- *Understanding QFX Series Virtual Chassis*

## Understanding High Availability on an EX Series Virtual Chassis

You increase your network's high availability (HA) when you interconnect a Juniper Networks EX Series Ethernet switch into a Virtual Chassis. A Virtual Chassis is more fault tolerant than a standalone EX series switch because it remains up when a single member switch fails.

You can further improve HA by configuring the HA features available for your EX Series Virtual Chassis. You can, for instance, configure Link Aggregation Groups (LAG) bundles to include member links on multiple member switches in the same Virtual Chassis. This configuration increases fault tolerance because traffic traversing the LAG can be redirected to an active member switch when a single member switch fails.

A Virtual Chassis has dual Routing Engines—the switch in the master role and the switch in the backup role—and therefore supports many HA features not supported on standalone EX Series switches. For a complete list of High Availability features available for your EX Series Virtual Chassis, see *EX Series Virtual Chassis Software Features Overview*.

Many HA features for the EX Series Virtual Chassis are designed to improve network resiliency after a Routing Engine switchover. [Table 569 on page 5095](#) describes the effects of a Routing Engine switchover when no high availability features are enabled and when some High Availability features are enabled.

Table 569: Effects of a Routing Engine Switchover

| High Availability Feature                                             | Effect of Routing Engine Switchover                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No HA features enabled                                                | Kernel and forwarding state information is not preserved to the backup Routing Engine. A convergence process that requires all interfaces on the Virtual Chassis to be taken offline has to be performed before the Virtual Chassis returns online. The switchover can take several minutes and the Virtual Chassis does not send or receive traffic until the switchover is complete.                                                                                                      |
| Graceful Routing Engine switchover (GRES) enabled                     | Kernel and forwarding state information is preserved on both Routing Engines, so the convergence process does not occur and the switchover happens quickly with minimal traffic loss.                                                                                                                                                                                                                                                                                                       |
| Nonstop active routing (NSR), Nonstop bridging (NSB), or both enabled | <p>Layer 2 protocols that are supported by NSB are not disrupted by a Routing Engine switchover when NSB is enabled. Layer 2 protocol information for all active Layer 2 protocols is stored on both Routing Engines when NSB is enabled.</p> <p>Layer 3 protocols that are supported by NSR are not disrupted by a Routing Engine switchover when NSR is enabled. Layer 3 protocol information for all active Layer 3 protocols is stored on both Routing Engines when NSR is enabled.</p> |
| Graceful Protocol Restart enabled                                     | Traffic is not interrupted during the switchover. Interface and kernel information is preserved. Graceful restart protocol extensions quickly collect and restore routing information for supported protocols from the neighboring devices.                                                                                                                                                                                                                                                 |

**Related  
Documentation**

- *EX Series Virtual Chassis Overview*
- *High Availability Features for EX Series Switches Overview*





# Configuration

- [Configuration Tasks on page 5097](#)
- [Configuration Statements on page 5117](#)

## Configuration Tasks

---

- [Configuring an EX4300 Virtual Chassis \(CLI Procedure\) on page 5097](#)
- [Adding a New Switch to an Existing EX4300 Virtual Chassis \(CLI Procedure\) on page 5103](#)
- [Replacing a Member Switch of a Virtual Chassis Configuration \(CLI Procedure\) on page 5104](#)
- [Configuring Mastership of a Virtual Chassis \(CLI Procedure\) on page 5107](#)
- [Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port \(CLI Procedure\) on page 5109](#)
- [Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis \(CLI Procedure\) on page 5113](#)
- [Disabling Split and Merge in a Virtual Chassis \(CLI Procedure\) on page 5114](#)
- [Configuring Automatic Software Update on Virtual Chassis Member Switches \(CLI Procedure\) on page 5114](#)
- [Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge \(CLI Procedure\) on page 5116](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 5116](#)

## Configuring an EX4300 Virtual Chassis (CLI Procedure)

This topic provides information about configuring a non-mixed EX4300 Virtual Chassis only. For information on configuring an QFX Virtual Chassis that includes EX4300 member switches, see *Configuring a QFX Series Virtual Chassis (CLI Procedure)*. For information on configuring EX4300 switches as leaf nodes in a VCF, see [“Preprovisioning a Virtual Chassis Fabric” on page 5264](#) or [“Autoprovisioning a Virtual Chassis Fabric” on page 5261](#).

In an EX4300 Virtual Chassis, you interconnect the EX4300 switches by using uplink ports configured as Virtual Chassis ports (VCPs). The QSFP+ uplink ports, which support 40Gbps speeds and can connect switches that are up to 492 feet (150 m) apart, are configured as VCPs by default. The SFP+ uplink ports, which support 10Gbps speeds

and can connect switches that are up to 6.2 miles (10 km) apart, can be configured into VCPs.

The simplest way to interconnect EX4300 switches into a Virtual Chassis is to install the switches within 492 feet of one another and interconnect them into a Virtual Chassis by using the QSFP+ ports.

If the member switches must be installed in locations that are more than 492 feet apart—such as at a different site or at a distant location within the same site—or if you are using the QSFP+ ports for another purpose, you must configure the SFP+ uplink module ports into VCPs. If you need additional VCP bandwidth between two member switches, you can configure additional ports as VCPs between the member switches. The ports that have identical speeds become links in a link aggregation group (LAG) to provide resiliency to the Virtual Chassis; for instance, if you had two QSFP+ ports and two SFP+ ports configured as VCPs connecting to the same switch, you would have two LAGs—one LAG with two 40Gbps QSFP+ port member links and another with two 10Gbps SFP+ port member links—between the member switches.



**NOTE:** A Virtual Chassis configuration has two Routing Engines—the master switch and the backup switch. Therefore, we recommend that you always use `commit synchronize` rather than simply `commit` to save configuration changes made for a Virtual Chassis. This ensures that the configuration changes are saved on both Routing Engines.

An EX4300 Virtual Chassis can be configured with either:

- A nonprovisioned configuration—The master sequentially assigns a member ID to other member switches. The role is determined by the mastership priority value and other factors in the master election algorithm.
- A preprovisioned configuration—You can deterministically control the member ID and role assigned to a member switch by tying the member switch to its serial number.



**NOTE:** You must configure a VLAN on all interfaces in an EX4300 Virtual Chassis, with the exception of the interfaces on member switch 0, before the interfaces can send or receive traffic. The interfaces on member switch 0 are initially placed into the default VLAN; the interfaces on all other member switches are not placed into any VLAN. See [“Configuring VLANs for EX Series Switches \(CLI Procedure\)” on page 2337](#).



**NOTE:** On an EX4300 Virtual Chassis, STP is disabled on all interfaces except the interfaces on member switch 0 until a type of STP is enabled. See *Configuring STP (CLI Procedure)* to enable STP on the interfaces in your EX4300 Virtual Chassis.

This topic includes:

- [Configuring an EX4300 Virtual Chassis with a Nonprovisioned Configuration File on page 5099](#)
- [Configuring an EX4300 Virtual Chassis with a Preprovisioned Configuration File on page 5101](#)

### [Configuring an EX4300 Virtual Chassis with a Nonprovisioned Configuration File](#)

You can use a nonprovisioned configuration to configure an EX4300 Virtual Chassis.

To configure the Virtual Chassis using a nonprovisioned configuration:



**NOTE:** We recommend that you physically cable the optical ports as the final step of this procedure.

You can, however, configure an EX4300 Virtual Chassis while the cables are physically connected.

1. Power on only the switch that you will use as the master switch.
2. Run the EZSetup program on the master switch, specifying the identification parameters. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* for details.



**NOTE:** The properties that you specify for the master switch apply to the entire Virtual Chassis configuration.

3. (Optional) Configure the master switch with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis:

[edit]

```
user@switch# set interfaces vme unit 0 family inet address /ip-address/mask/
```

4. (Optional) Configure mastership priority for the other member switches.

[edit virtual-chassis]

```
user@switch# set member 0 mastership-priority 255
```

```
user@switch# set member 1 mastership-priority 255
```

```
user@switch# set member 2 mastership-priority 10
```

```
user@switch# set member 3 mastership-priority 9
```

```
user@switch# set member 4 mastership-priority 8
```

```
user@switch# set member 5 mastership-priority 7
```

```
user@switch# set member 6 mastership-priority 6
```

```
user@switch# set member 7 mastership-priority 5
```

```
user@switch# set member 8 mastership-priority 4
```

```
user@switch# set member 9 mastership-priority 3
```

The mastership priority value determines the roles in a non-provisioned Virtual Chassis configuration. The switches with the highest mastership priority values assume the master and backup roles. All other switches assume the linecard role.

If you do not configure the mastership priority for any switch in your Virtual Chassis, including when you do not configure the Virtual Chassis, all switches assume the default mastership priority of 128. The master election algorithm selects the roles for the member switches. In most cases, the switches that have been powered on the longest assume the master and backup roles when all Virtual Chassis member switches are configured with the same mastership priority. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 5078](#) for additional information on the master election algorithm.

A switch with a mastership priority of 0 never assumes the master or backup role.



**NOTE:** We recommend that you specify the same mastership priority value for the intended master and backup members.

5. (Optional. Recommended for a two-member Virtual Chassis) On the master switch, disable the split and merge feature:

```
[edit virtual-chassis]
user@switch# set no-split-detection
```

6. Power on the other member switches.
7. On each individual member switch, configure the SFP+ optical ports that will be used to interconnect the EX4300 member switches into VCPs.



**NOTE:** Because QSFP+ ports are configured into VCPs by default, you do not usually have to perform this step when you are using a QSFP+ port as a VCP.

You only need to configure a QSFP+ port as a VCP if you previously configured the QSFP+ port into a network port. If you previously configured the QSFP+ port into a network port, perform this step to configure the QSFP+ port into a VCP.

```
user@switch-0> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-0> request virtual-chassis vc-port set pic-slot 1 port 1
user@switch-1> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-1> request virtual-chassis vc-port set pic-slot 1 port 1
user@switch-2> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-2> request virtual-chassis vc-port set pic-slot 1 port 1
user@switch-3> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-3> request virtual-chassis vc-port set pic-slot 1 port 1
```



**NOTE:** If you want to change the member ID that the master has assigned to a member switch, use the [request virtual-chassis renumber](#) command.

### Configuring an EX4300 Virtual Chassis with a Preprovisioned Configuration File

Preprovisioning a Virtual Chassis configuration allows you to assign the member ID and role for each switch in the Virtual Chassis.

To configure a Virtual Chassis using a preprovisioned configuration:



**NOTE:** We recommend that you physically cable the optical ports as the final step of this procedure.

You can, however, configure an EX4300 Virtual Chassis while the cables are physically connected.

1. Make a list of the serial numbers of all the switches to be connected in a Virtual Chassis configuration.
2. Note the intended role (**routing-engine** or **line-card**) of each switch. If you configure the member with a **routing-engine** role, it is eligible to function in the master or backup role. If you configure the member with a **line-card** role, it is not eligible to function in the master or backup role.
3. Power on only the switch that you plan to use as the master switch.
4. Run the EZSetup program on the master switch, specifying the identification parameters. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* for details.



**NOTE:** The properties that you specify for the master switch apply to the entire Virtual Chassis configuration.

5. (Optional) Configure the master switch with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis:

```
[edit]
```

```
user@switch# set interfaces vme unit 0 family inet address /ip-address/mask/
```

6. Specify the preprovisioned configuration mode:

```
[edit virtual-chassis]
```

```
user@switch# set preprovisioned
```

7. Specify all the members that you want included in the Virtual Chassis, listing each switch's serial number with the desired member ID and role.



**NOTE:** You can retrieve the switch's serial number using the **show chassis hardware** command output or by viewing the serial number ID label on the switch. See *Locating the Serial Number on an EX4300 Switch or Component*.

```
[edit virtual-chassis]
```

```
user@switch# set member 0 serial-number abc123 role routing-engine
```

```
user@switch# set member 1 serial-number def456 role routing-engine
```

```

user@switch# set member 2 serial-number ghi789 role line-card
user@switch# set member 3 serial-number jkl012 role line-card
user@switch# set member 4 serial-number mno013 role line-card
user@switch# set member 5 serial-number pqr014 role line-card
user@switch# set member 6 serial-number stu015 role line-card
user@switch# set member 7 serial-number vwx016 role line-card
user@switch# set member 8 serial-number yzz017 role line-card
user@switch# set member 9 serial-number aaa018 role line-card

```

8. (Optional. Recommended for a two-member Virtual Chassis) Disable the split and merge feature:

```

[edit virtual-chassis]
user@switch# set no-split-detection

```

9. Power on the other member switches. The member IDs and roles have been determined by the configuration, so you can power on the member switches in any order.
10. On each individual member switch, configure the SFP+ optical ports that will be used to interconnect the EX4300 member switches into VCPs.



**NOTE:** You can also use the [request virtual-chassis vc-port](#) command to configure a QSFP+ port into a VCP.

Because QSFP+ ports are configured into VCPs by default, you do not usually have to perform this step when you are using a QSFP+ port as a VCP.

You only need to configure a QSFP+ port as a VCP if you previously configured the QSFP+ port into a network port. If you previously configured the QSFP+ port into a network port, perform this step to configure the QSFP+ port into a VCP.

```

user@switch-0> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-0> request virtual-chassis vc-port set pic-slot 1 port 1
user@switch-1> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-1> request virtual-chassis vc-port set pic-slot 1 port 1
user@switch-2> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-2> request virtual-chassis vc-port set pic-slot 1 port 1
user@switch-3> request virtual-chassis vc-port set pic-slot 1 port 0
user@switch-3> request virtual-chassis vc-port set pic-slot 1 port 1

```



**NOTE:** You cannot modify the mastership priority when you are using a preprovisioned configuration. The mastership priority values are generated automatically and controlled by the role that is assigned to the member switch in the configuration file. The two Routing Engines are assigned the same mastership priority value. However, the member that was powered on first has higher prioritization according to the master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected”](#) on page 5078.

#### Related Documentation

- [Configuring Mastership of a Virtual Chassis \(CLI Procedure\) on page 5107](#)
- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157](#)

## Adding a New Switch to an Existing EX4300 Virtual Chassis (CLI Procedure)

You can use this procedure to add an EX4300 switch to an EX4300 Virtual Chassis. For information on adding an EX4300 switch to a QFX Series Virtual Chassis or to a Virtual Chassis Fabric (VCF), see *Adding a New Switch to an Existing QFX Series Virtual Chassis (CLI Procedure)* or *“Adding a Device to a Virtual Chassis Fabric” on page 5270*.

Before you begin, be sure you have:

- Mounted the new switch in a rack.
- Confirmed that the new switch is powered off.
- If you are expanding a preprovisioned configuration, made a note of the serial number (the number is on the back of the switch). You will need to edit the Virtual Chassis configuration to include the serial number of the new member switch.
- If you are expanding a preprovisioned configuration, edited the existing Virtual Chassis configuration to include the serial number of the new member switch. The parameters specified in the master Virtual Chassis configuration file are applied to the new switch after it has been interconnected to an existing member switch.



**NOTE:** If you are expanding a preprovisioned Virtual Chassis configuration, you can use the autoprovisioning feature to add member switches to that configuration.

- (Optional) Configured Ethernet interfaces on different member switches into the same LAG. See *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*.

An active member switch might temporarily go down before coming back up as part of this procedure. Having traffic load-balanced across member switches using a LAG helps alleviate traffic loss during this procedure.

To add a new member switch to an existing Virtual Chassis configuration:

1. If the new member switch has been previously configured, revert that switch's configuration to the factory defaults before interconnecting it into the Virtual Chassis. See *“Reverting to the Default Factory Configuration for the EX Series Switch” on page 620*.

2. Interconnect the unpowered new switch to one member of the existing Virtual Chassis configuration using a QSFP+ or SFP+ port.

Connect only one VCP on the unpowered new switch to a VCP on a member switch in the existing Virtual Chassis at this point of the procedure.

3. Power on the new switch.
4. (SFP+ interface only) Set the SFP+ interface as a Virtual Chassis Port (VCP):

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port port-number
```



**NOTE:** QSFP+ ports are configured as VCPs, by default. You do not, therefore, typically need to perform this step on QSFP+ ports.

You can use the [request virtual-chassis vc-port](#) to set a QSFP+ port as a VCP if the QSFP+ port had previously been configured as a network port.

5. Confirm that the new member switch is now included within the Virtual Chassis configuration by entering the **show virtual-chassis** command. The new member switch should be listed in the output and the **Status** is **Prsnt**.
6. Cable the next port into the Virtual Chassis, using Steps 2 through 5.



**CAUTION:** If you immediately cable both VCPs on the new switch into the existing Virtual Chassis at the same time, a member switch that was already part of the Virtual Chassis might become nonoperational for several seconds. Network traffic to this switch is dropped during the downtime.

The member switch will return to the normal operational state with no user intervention, and normal operation of the Virtual Chassis will resume after this downtime.

**Related Documentation**

- [Configuring an EX4300 Virtual Chassis \(CLI Procedure\)](#) on page 5097

## Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure)



**NOTE:** This topic does not apply to Virtual Chassis Fabric (VCF) or EX8200 Virtual Chassis. See [“Removing a Device From a Virtual Chassis Fabric” on page 5277](#) or [Adding or Replacing a Member Switch or an External Routing Engine in an EX8200 Virtual Chassis \(CLI Procedure\)](#).

You can replace a member switch in a Virtual Chassis without disrupting network service on the other members. You can retain the existing configuration of the member switch and apply it to a new member switch, or you can free up the member ID and make it available for assignment to a new member switch.

If you want to replace a member switch of a *mixed* Virtual Chassis that contains EX4200, EX4500, or EX4550 switches, see [Removing an EX4200, EX4500, or EX4550 Switch From a Mixed Virtual Chassis \(CLI Procedure\)](#).



To replace a member switch, use the procedure that matches what you need to accomplish:

- [Remove, Repair, and Reinstall the Same Switch on page 5105](#)
- [Remove a Member Switch, Replace It with a Different Switch, and Reapply the Old Configuration on page 5105](#)
- [Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch on page 5106](#)

### **Remove, Repair, and Reinstall the Same Switch**

If you need to repair a member switch, you can remove it from the Virtual Chassis configuration without disrupting network service for the other members. The master stores the configuration for the member ID so that it can be reapplied when the member switch (with the same base MAC address) is reconnected.

To remove, repair, and reinstall the member switch:

1. Power off and disconnect the member switch to be repaired.
2. Repair, as necessary.
3. Reconnect the switch and power it on.

### **Remove a Member Switch, Replace It with a Different Switch, and Reapply the Old Configuration**

If you are unable to repair a member switch, you can replace it with a different member switch while retaining the previous configuration. The master stores the configuration of the member that was removed. When you connect a different member switch, the master assigns a new member ID. But the old configuration is still stored under the previous member ID of the previous member switch.



**NOTE:** If you have used a preprovisioned configuration, you can use the `replace` command to change the serial number in the Virtual Chassis configuration file. Substitute the serial number of the replacement member switch (on the back of the switch) for the serial number of the member switch that was removed.

To remove and replace a switch and reapply the old configuration:

1. Power off and disconnect the member switch to be replaced.
2. If the replacement member switch has been previously configured, revert that switch's configuration to the factory defaults. See [“Reverting to the Default Factory Configuration for the EX Series Switch” on page 620](#) for information about reverting to the factory default configuration on an EX Series switch or *Reverting to the Default Factory Configuration* for information about reverting to the factory default configuration on a QFX Series switch.

3. If you are interconnecting a switch using a dedicated VCP, connect one VCP on the replacement member switch to a VCP of another Virtual Chassis member switch.

If you are interconnecting a switch using an optical port configured as a VCP, cable the optical ports together then configure the port on the Virtual Chassis as a VCP:

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port port-number
```

4. Power on the new member switch.
5. Confirm that the new member switch is now included in the Virtual Chassis configuration by checking the front-panel LCD or the for the member ID. It should display a member ID in the range from 0 through 9.

If you are using a switch that does not have an LCD interface, confirm the switch is part of the Virtual Chassis configuration by entering the **show virtual-chassis** and reviewing the output.

6. Cable the other VCP on the new member switch into the Virtual Chassis. Use the instruction in Step 3 to complete this step.



**CAUTION:** If you immediately cable both VCPs on the new switch into the existing Virtual Chassis at the same time, a member switch that was already part of the Virtual Chassis might become nonoperational for several seconds. Network traffic to this switch is dropped during the downtime.

The member switch will return to the normal operational state with no user intervention, and normal operation of the Virtual Chassis will resume after this downtime.

7. On the master switch, Issue the **request virtual-chassis renumber** command from the Virtual Chassis master to change the member switch's current member ID to the member ID of the member switch that was removed from the Virtual Chassis configuration.

---

### Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch

---

When you remove a member switch from the Virtual Chassis configuration, the master keeps that member switch's member ID in reserve. To make that member switch's member ID available for reassignment, issue the **request virtual-chassis recycle** command from the Virtual Chassis master.



**NOTE:** When you add or delete members in a Virtual Chassis configuration, internal routing changes might cause temporary traffic loss for a few seconds.

#### Related Documentation

- *Adding or Replacing a Member Switch or an External Routing Engine in an EX8200 Virtual Chassis (CLI Procedure)*
- *Adding a New Switch to an Existing QFX Series Virtual Chassis (CLI Procedure)*

- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157](#)
- [Adding a New EX4200 Switch to an Existing EX4200 Virtual Chassis \(CLI Procedure\)](#)
- [Adding an EX4200 Switch to a Preprovisioned EX4500 Virtual Chassis or a Preprovisioned Mixed EX4200 and EX4500 Virtual Chassis \(CLI Procedure\)](#)
- [Adding an EX4500 Switch to a Preprovisioned EX4200 Virtual Chassis \(CLI Procedure\)](#)
- [Adding an EX4500 Switch to a Nonprovisioned EX4200 Virtual Chassis \(CLI Procedure\)](#)

## Configuring Mastership of a Virtual Chassis (CLI Procedure)



**NOTE:** This topic applies to all EX Series Virtual Chassis except EX8200 Virtual Chassis. See [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#) for information about EX8200 Virtual Chassis.

You can designate the role (master, backup, or linecard) that a member switch performs within any Virtual Chassis, whether or not you are using a preprovisioned configuration.



**NOTE:** A Virtual Chassis configuration has two Routing Engines—one is the switch in the master role and the other is the switch in the backup role. Therefore, we recommend that you always use `commit synchronize` rather than `commit` to save configuration changes made for a Virtual Chassis. This ensures that the configuration changes are saved in both Routing Engines.

This topic describes:

- [Configuring Mastership Using a Preprovisioned Configuration File on page 5107](#)
- [Configuring Mastership Using a Configuration File That Is Not Preprovisioned on page 5108](#)

### Configuring Mastership Using a Preprovisioned Configuration File

To configure mastership using a preprovisioned configuration:

1. Note the serial numbers of the switches that you want to function in the master role and backup role.
2. Power on only the switch that you want to function in the master role.
3. Edit the configuration to specify the preprovisioned configuration mode:
 

```
[edit virtual-chassis]
user@switch# set preprovisioned
```
4. Specify the serial numbers of the member switches that you want to function as master and backup, specifying their role as **routing-engine**:
 

```
[edit]
user@switch# set virtual-chassis member 0 serial-number abc123 role routing-engine
user@switch# set virtual-chassis member 1 serial-number def456 role routing-engine
```



**NOTE:** You cannot directly modify the mastership priority value when you are using a preprovisioned configuration. The mastership priority values are generated automatically and controlled by the role that is assigned to the member switch in the configuration file. The two members assigned the **routing-engine** role are assigned the same mastership priority value (128). However, the member that was powered on first has higher priority for the master role election according to the master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 5078](#). Only two members can be configured with the **routing-engine** role.

5. Specify the serial numbers of any other member switches that you are including in the Virtual Chassis configuration. You can also explicitly configure their role as **line-card**.

### Configuring Mastership Using a Configuration File That Is Not Preprovisioned

To configure mastership of the Virtual Chassis through a configuration that is not preprovisioned:

1. Power on only the switch that you want to function in the master role.
2. Configure the highest possible mastership priority value (**255**) for the member that you want to function in the master role:
3. Configure the same mastership priority value (continue to edit the Virtual Chassis configuration on the master) for the member that you want to be in the backup role:

```
[edit virtual-chassis]
user@switch# set member 0 mastership-priority 255
```

```
[edit virtual-chassis]
user@switch# set member 1 mastership-priority 255
```



**NOTE:** We recommend that the master and backup have the same mastership priority value to prevent the master and backup status from switching back and forth between master and backup members in failover conditions.

4. Use the default mastership priority value (**128**) for the remaining member switches or configure the mastership priority to a value that is lower than the value specified for members functioning in the master and backup roles.

#### **Related Documentation**

- [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\)](#)
- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157](#)
- [Adding a New EX4200 Switch to an Existing EX4200 Virtual Chassis \(CLI Procedure\)](#)

- *Adding an EX4200 Switch to a Preprovisioned EX4500 Virtual Chassis or a Preprovisioned Mixed EX4200 and EX4500 Virtual Chassis (CLI Procedure)*
- *Adding an EX4500 Switch to a Preprovisioned EX4200 Virtual Chassis (CLI Procedure)*
- *Adding an EX4500 Switch to a Nonprovisioned EX4200 Virtual Chassis (CLI Procedure)*

## Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)

The procedure described in this topic can be used to connect two EX series switches together within the same Virtual Chassis.

You typically configure an uplink port as a Virtual Chassis Port (VCP) for one of the following reasons:

- You want to interconnect two EX series switches into a Virtual Chassis that are located in different wiring closets or sites, and the switches are farther apart than the maximum length of the dedicated VCP cable.
- You are configuring an EX Series Virtual Chassis composed of switches that support Virtual Chassis but do not have dedicated VCPs. EX2200, EX3300, and EX4300 switches support Virtual Chassis but do not have dedicated VCPs.

You must manually configure VCPs to connect EX2200 switches together to form an EX2200 Virtual Chassis. See *Setting a Port on an EX2200 Switch as a Virtual Chassis Port (CLI Procedure)*.

This procedure is usually not needed to configure an EX3300 Virtual Chassis. Uplink ports 2 and 3 on an EX3300 switch are configured as VCPs by default and, therefore, do not require user configuration to be set as VCPs. We recommend that you use this procedure to configure an uplink port on an EX3300 switch as a VCP only if you configured ports 2 and 3 as network uplink ports and the ports need to be reconfigured as VCPs, or when ports 2 and 3 cannot be used as VCPs for an unexpected reason. You can use this procedure to configure any uplink port on an EX3300 switch as a VCP.

QSFP+ ports on EX4300 switches are configured as VCPs, by default. You must use this procedure if you want to configure an SFP+ port on an EX4300 switch as a VCP, or if you want to configure a QSFP+ port that had been configured into a network port back into a VCP.

You can interconnect EX4200, EX4500, and EX4550 switches that are beyond the reach of the dedicated Virtual Chassis cables as members of a Virtual Chassis by using the uplink ports—including the ports on the SFP uplink module, SFP+ uplink module, or XFP uplink module—and connecting the uplink ports. To use the uplink ports or SFP network ports for interconnecting member switches, you must explicitly set the uplink ports as VCPs.



**NOTE:** You cannot set a 1000BASE-T copper SFP transceiver (EX-SFP-1GE-T) connection as a VCP on EX4200, EX4500, and EX4550 switches.



**NOTE:** When an uplink port is set as a VCP, it cannot be used for any other purpose. You can set one port as a VCP and configure the other port in trunk mode as an uplink to another switch.

Before you set an uplink port as a VCP:

1. Install the uplink module in the member switches that you want to interconnect, if you are configuring an uplink module port as a VCP.
2. Power on and connect to the switch that you plan to designate as the master of the Virtual Chassis.



**NOTE:** Do not power on the other switches at this point.

3. Run EZSetup on the switch that you are configuring to be the master. Follow the prompts to specify the hostname and other identification, time zone, and network properties. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* for details. The properties that you specify for the master apply to the entire Virtual Chassis, including all the member switches that you later interconnect with the master.
4. If you want to configure and manage the Virtual Chassis remotely, specify the VME global management interface. You can configure the VME global management interface when you are setting up the master or you can do it after completing the other configuration steps for the Virtual Chassis. See *Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)*.
5. Configure mastership of the Virtual Chassis by using either the nonprovisioned or preprovisioned configuration. See [“Configuring Mastership of a Virtual Chassis \(CLI Procedure\)” on page 5107](#) for details.



**NOTE:** A Virtual Chassis has two Routing Engines, one in the master and the other in the backup. Therefore, we recommend that you always use **commit synchronize** rather than simply **commit** to save configuration changes made for a Virtual Chassis. This ensures that the configuration changes are saved in both Routing Engines.

Before you begin to interconnect a Virtual Chassis across long distances, such as between wiring closets:

- Prepare the existing Virtual Chassis for interconnecting with a potential member switch that is beyond the reach of a dedicated Virtual Chassis cable by setting at least one uplink VCP on an existing member of the Virtual Chassis.
- Prepare the potential member switch for interconnecting with the existing Virtual Chassis by setting at least one uplink VCP on the standalone switch.



**NOTE:** We recommend that you set two uplink VCPs within each wiring closet for redundancy.

This topic describes:

1. [Setting an Uplink VCP Between the Member Switches on page 5111](#)
2. [Setting an Uplink VCP on a Standalone Switch on page 5111](#)

### Setting an Uplink VCP Between the Member Switches

You can set an uplink port as a VCP.



**NOTE:** If you use the SFP+ uplink module, you must configure all member switches to support either 1-gigabit SFP transceivers or 10-gigabit SFP+ transceivers on EX4200 switches. See *Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module (CLI Procedure)*.

To set the uplink ports for the local member switch (for example, member 0) and for a different member switch (for example, member 1) to function as VCPs:

1. Set one uplink port of member 0 as a VCP. You do not need to specify the **member member-id** option, because the command applies by default on the member where it is executed.

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port 0
```

2. Set one uplink port of member 1 as a VCP.

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port 0 member 1
```

This step includes the member *member-id* option, because it is executed on a different member switch than the local member switch.

### Setting an Uplink VCP on a Standalone Switch

You can set an uplink VCP on a standalone switch. You must set an uplink port on the standalone switch as a VCP prior to physically interconnecting the switch with the existing Virtual Chassis. Otherwise, the master cannot detect that the switch is a member of the Virtual Chassis.

To set one uplink VCP on the potential member, which is currently operating as a standalone switch:

1. Power on the standalone switch.
2. Set one uplink port as a VCP. You do not need to specify the **member member-id** option, because the command applies by default on the member where it is executed.

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port 0
```



**NOTE:** If you do specify the member *member-id* option, use member ID 0. Because the switch is not yet interconnected with the other members of the Virtual Chassis, its current member ID is 0. Its member ID will change when it is interconnected with the Virtual Chassis. It does not impact the functioning of the uplink VCP that its VCP is set with 0 as the member ID. The VCP has significance only on the local switch.

3. After you have set the uplink VCP on the standalone switch, physically interconnect its uplink port with the VCP uplink ports of the members in the existing Virtual Chassis.

The new member switch reboots and joins the now expanded Virtual Chassis with a different member ID.



**NOTE:** The setting for the new member switch's uplink VCP remains intact and is not affected by the change of member ID.

4. If you have additional members in the second wiring closet, set a redundant VCP uplink on another member switch by issuing the **request virtual-chassis vc-port** command.

#### Related Documentation

- [Configuring an EX3300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a Virtual Chassis on an EX Series Switch \(J-Web Procedure\)](#)
- [Example: Configuring an EX4200 Virtual Chassis Interconnected Across Multiple Wiring Closets](#)
- [Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File](#)
- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157](#)



## Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis (CLI Procedure)

When a backup member takes control of a Virtual Chassis because of a reset or other temporary failure, the backup member uses the MAC address of the old master switch as the system MAC base address. This process helps ensure a smooth transition of mastership with no disruption to network connectivity.

The MAC persistence timer is used in situations in which the master switch is no longer a member of the Virtual Chassis because it has been physically disconnected or removed. If the old master switch does not rejoin the Virtual Chassis before the timer elapses, the new master switch starts using its own MAC address as the system's MAC base address. For information regarding how the system MAC base address is used to assign MAC addresses to ports in a Virtual Chassis, see [“Understanding MAC Address Assignment on a Virtual Chassis” on page 5093](#).

The default timer value is 10 minutes. The maximum timer value is 60 minutes.

You can disable the MAC persistence timer starting in Junos OS Release 12.2. When the MAC persistence timer is disabled, the MAC address of the old master switch is used as the system MAC base address; no MAC address changes occur within the Virtual Chassis even when the old master switch is no longer a member of the Virtual Chassis because it has been physically disconnected or removed.

To configure or modify the MAC persistence timer:

```
[edit virtual-chassis]
user@switch# set mac-persistence-timer minutes
```

To disable the MAC persistence timer:

```
[edit virtual-chassis]
user@switch# set mac-persistence-timer disable
```

### Related Documentation

- [Configuring an EX3300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\)](#)
- [Understanding EX Series Virtual Chassis Components on page 5067](#)

## Disabling Split and Merge in a Virtual Chassis (CLI Procedure)

The split and merge feature is enabled by default on all EX Series switches and QFX Series devices in a Virtual Chassis. You can disable the split and merge feature. If you disable the split and merge feature and the Virtual Chassis splits, both parts of the split Virtual Chassis configuration remain active.

In a preprovisioned Virtual Chassis, if both of the Routing Engines end up in the same Virtual Chassis configuration after a split, the other part of the split Virtual Chassis configuration remains inactive. If the Routing Engines end up in different parts of the split Virtual Chassis configuration and the rest of the member switches are configured as having linecard roles, then a backup Routing Engine might not be selected for either part.

We recommend disabling split and merge on a Virtual Chassis with two member switches. A two-member switch Virtual Chassis that has disabled split and merge can reform more quickly and with less complications as a result of the feature being disabled.

To disable the split and merge feature in a Virtual Chassis:

```
[edit]
user@switch# set virtual-chassis no-split-detection
```

### Related Documentation

- [Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge](#)
- [Understanding Split and Merge in a Virtual Chassis on page 5087](#)

## Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure)

The automatic software update feature allows you to automatically update the software version on prospective member switches as they are added so that they can join the Virtual Chassis.



**NOTE:** The version of Junos OS running on the Virtual Chassis must be compatible with the software running on the prospective member switch for an automatic software update to occur. For information on Junos OS compatibility and other automatic software update restrictions, see [“Understanding Automatic Software Update on Virtual Chassis Member Switches” on page 5090](#).

Before you begin, ensure that you know the name or the URL of the software package to be used by the automatic software update feature.

To configure the automatic software update feature for an EX Series or QFX Series Virtual Chassis with the exception of a mixed Virtual Chassis containing at least one EX4200 switch and at least one EX4500 or EX4550 switch:

[edit]

```
user@switch# set virtual-chassis auto-sw-update package-name package-name
```

To configure the automatic software update feature on a mixed Virtual Chassis containing at least one EX4200 switch and at least one EX4500 or EX4550 switch:

[edit]

```
user@switch# set virtual-chassis auto-sw-update ex-4200 package-name package-name
```

```
user@switch# set virtual-chassis auto-sw-update ex-4500 package-name package-name
```

If the software package is located on a local directory on the switch, use the following format for **package-name**:

***/pathname/package-name***

If the software package is to be downloaded and installed from a remote location, use one of the following formats:

***ftp://hostname/pathname/package-name***

***ftp://username:prompt@ftp.hostname.net/package-name***

***http://hostname/pathname/package-name***

If you are configuring a mixed Virtual Chassis containing at least one EX4200 switch and at least one EX4500 or EX4550 switch, use the **ex-4200** keyword when you are specifying a path to a package for the EX4200 switches and the **ex-4500** when you are specifying a path to a package for the EX4500 or EX4550 switches. You do not need to specify the **ex4500** keyword when configuring automatic software update for a mixed EX4500 and EX4550 Virtual Chassis, however, because the Junos OS package for an EX4500 switch updates the software for both EX4500 and EX4550 switches.

#### Related Documentation

- *Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches*
- [Understanding Automatic Software Update on Virtual Chassis Member Switches on page 5090](#)

## Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure)

Every Virtual Chassis has a unique ID that is automatically assigned when the Virtual Chassis configuration is formed. You can also explicitly assign a Virtual Chassis ID using the **set virtual-chassis id** command. When two Virtual Chassis configurations attempt to merge, the Virtual Chassis ID that you assigned takes precedence over the automatically assigned Virtual Chassis IDs and becomes the ID for the newly merged Virtual Chassis configuration.

To configure the Virtual Chassis ID:

```
[edit]
user@switch# set virtual-chassis id id
```

### Related Documentation

- *Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge*
- [Understanding Split and Merge in a Virtual Chassis on page 5087](#)

## Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)

In a Virtual Chassis, one member switch is assigned the master role and has the master Routing Engine. Another member switch is assigned the backup role and has the backup Routing Engine. Graceful Routing Engine switchover (GRES) enables the master and backup Routing Engines in a Virtual Chassis configuration to switch from the master to backup without interruption to packet forwarding. When you configure graceful Routing Engine switchover, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and the forwarding state.

To set up the Virtual Chassis configuration to use graceful Routing Engine switchover (GRES):

1. Set up a minimum of two switches in a Virtual Chassis configuration with mastership priority of 255:

```
[edit]
user@switch# set virtual-chassis member 0 mastership-priority 255
[edit]
user@switch# set virtual-chassis member 1 mastership-priority 255
```

2. Set up graceful Routing Engine switchover:

```
[edit]
user@switch# set chassis redundancy graceful-switchover
```

Commit the configuration.



**NOTE:** We recommend that you use the **commit synchronize** command to save any configuration changes that you make to a multimember Virtual Chassis.

---

- Related Documentation**
- *Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet*
  - *High Availability Features for EX Series Switches Overview*
  - [Understanding EX Series Virtual Chassis Configuration on page 5086](#)
  - *Understanding QFX Series Virtual Chassis*

## Configuration Statements

- [\[edit virtual-chassis\] Configuration Statement Hierarchy on page 5117](#)
- [aliases \(Virtual Chassis\) on page 5120](#)
- [alias-name \(Virtual Chassis aliases\) on page 5121](#)
- [auto-sw-update on page 5122](#)
- [graceful-restart \(Enabling Globally\) on page 5124](#)
- [graceful-switchover on page 5125](#)
- [id on page 5126](#)
- [location \(Virtual Chassis\) on page 5127](#)
- [mac-persistence-timer on page 5128](#)
- [mastership-priority on page 5129](#)
- [member on page 5131](#)
- [no-management-vlan on page 5132](#)
- [no-split-detection on page 5133](#)
- [package-name on page 5134](#)
- [preprovisioned on page 5135](#)
- [redundancy \(Graceful Switchover\) on page 5136](#)
- [role on page 5137](#)
- [serial-number on page 5140](#)
- [serial-number \(Virtual Chassis aliases\) on page 5141](#)
- [traceoptions \(Virtual Chassis\) on page 5142](#)
- [vcp-no-hold-time on page 5145](#)
- [virtual-chassis on page 5147](#)

### [edit virtual-chassis] Configuration Statement Hierarchy

This topic lists supported and unsupported configuration statements in the **[edit virtual-chassis]** hierarchy level on EX Series and QFX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.

- Not all features are supported on all switch platforms.

For detailed information about feature support on specific EX Series or QFX Series switch platforms, see [Feature Explorer](#).

.

This topic lists:

- [Supported Statements in the \[edit virtual-chassis\] Hierarchy Level on page 5118](#)
- [Unsupported Statements in the \[edit virtual-chassis\] Hierarchy Level on page 5119](#)

### [Supported Statements in the \[edit virtual-chassis\] Hierarchy Level](#)

---

The following hierarchy shows the **[edit virtual-chassis]** configuration statements supported on EX Series or QFX Series switches:

```
virtual-chassis {
  aliases {
    serial-number serial-number {
      alias-name alias-name;
    }
  }
  auto-provisioned;
  auto-sw-update {
    (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
    package-name package-name;
  }
  fast-failover (ge | vcp disable | xe);
  graceful-restart {
    disable;
  }
  id id;
  mac-persistence-timer [minutes | disable];;
  member member-id {
    location location;
    mastership-priority number;
    no-management-vlan;
    role (line-card | routing-engine);
    serial-number;
  }
  no-split-detection;
  preprovisioned;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
      regex>;
    flag flag ;
  }
  vc-port {
    lag-hash (packet-based | source-port-based);
  }
  vcp-no-hold-time;
}
```

---

### Unsupported Statements in the `[edit virtual-chassis]` Hierarchy Level

---

All statements in the `[edit virtual-chassis]` hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4300 Virtual Chassis \(CLI Procedure\) on page 5097](#)
- [Configuring an EX2200 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX3300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches \(CLI Procedure\)](#)
- [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#)

## aliases (Virtual Chassis)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>aliases {<br/>    serial-number serial-number {<br/>        alias-name alias-name;<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). An alias allows you to more clearly identify the member switches in your Virtual Chassis or VCF by assigning a text label to a member switch's serial number.</p> <p>An alias is not specified for a device until the alias name is specified using the <b>alias-name</b> keyword.</p> <p>The alias appears in the <b>Alias-Name</b> field in the <b>show virtual-chassis</b> command.</p> <p>Alias usage is optional and aliases are used for administrative purposes only. Setting an alias has no effect on the operation of the member switch.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Understanding Virtual Chassis Fabric Components on page 5243</a></li><li>• <a href="#">Understanding QFX Series Virtual Chassis Components</a></li></ul>                                                                                                                                                                                                                                              |



## alias-name (Virtual Chassis aliases)

**Syntax** `alias-name alias-name;`

**Hierarchy Level** `[edit virtual-chassis aliases serial-number serial-number]`

**Release Information** Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series switches.

**Description** Create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). An alias allows you to more clearly identify the member switches in your Virtual Chassis or VCF by assigning a text label to a member switch's serial number.

The alias appears in the **Alias-Name** field in the **show virtual-chassis** command.

Alias usage is optional and aliases are used for administrative purposes only. Setting an alias has no effect on the operation of the member switch.

In the following example, the **dc-floor-1** alias name is assigned to the member switch with the serial number AB0123456789.

### set serial-number

```
[edit virtual-chassis aliases]
user@switch# set serial-number AB0123456789 alias-name dc-floor-1
```

### show virtual-chassis

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 9d5d.5556.919a
Fabric Mode: Enabled

Member ID  Status  Serial No  Alias-Name  Model  Mstr  prio  Role
0 (FPC 0)  Prsnt    AB0123456789  dc-floor-1  qfx5100-48s-6q  129  Master
<additional output removed for brevity>
```

**Options** *alias-name*—The text label, or alias, assigned to the member switch by the user.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\)](#)
- [Understanding Virtual Chassis Fabric Components on page 5243](#)
- [Understanding QFX Series Virtual Chassis Components](#)

## auto-sw-update

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>auto-sw-update {<br/>    (ex-4200   ex-4300   ex-4500   ex-4600   qfx-3   qfx-5)<br/>    package-name package-name;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>     | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>The <b>ex-4200</b> and <b>ex-4500</b> options introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>The <b>ex-4300</b>, <b>qfx-3</b>, and <b>qfx-5</b> options introduced in Junos OS Release 13.2X51-D20.</p> <p>The <b>ex-4600</b> option introduced in Junos OS Release 13.2X51-D25.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>         | <p>Enable the automatic software update feature for Virtual Chassis or Virtual Chassis Fabric (VCF) configurations.</p> <p>You should only use the keywords that specify a device—<b>ex-4300</b>, <b>ex-4600</b>, <b>qfx-3</b>, and <b>qfx-5</b>—when configuring automatic software update on a mixed Virtual Chassis or Virtual Chassis Fabric (VCF). You can simply specify the <i>package-name</i> without specifying the device keywords in non-mixed Virtual Chassis or VCF topologies.</p> <p>You must enter the <b>auto-sw-update</b> statement multiple times—once for each device family in your mixed Virtual Chassis or VCF—in most scenarios when enabling the automatic software update for a mixed Virtual Chassis or VCF.</p> <p>The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. You do not, therefore, need to specify the <b>ex-4500</b> keyword when configuring automatic software update for a mixed Virtual Chassis that include EX4500 and EX4550 switches only. You also only have to enter the <b>ex-4500</b> keyword once to configure automatic software update for all EX4500 and EX4550 member switches in the same mixed Virtual Chassis.</p> <p>The Junos OS package for a QFX3500 device updates the software for QFX3500 and QFX3600 devices. You do not, therefore, need to specify the <b>qfx-3</b> keyword when configuring automatic software update for a Virtual Chassis composed entirely of QFX3500 and QFX3600 devices. You also have to enter the <b>qfx-3</b> keyword only once to configure automatic software update for all QFX3500 and QFX3600 member devices in the same mixed Virtual Chassis.</p> <p>The remaining statement is explained separately.</p> |
| <b>Default</b>             | The automatic software update feature is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b>package-name package-name</b>—Specify a path to a Junos OS software image.</p> <p><b>ex-4200</b>—Specify a path to a Junos OS image for an EX4200 switch when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

**ex-4300**—Specify a path to a Junos OS image for an EX4300 switch when enabling automatic software update for a mixed Virtual Chassis or VCF.

**ex-4500**—Specify a path to a Junos OS image for an EX4500 switch, an EX4550 switch, or both types of switches when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. Therefore, you only enter this command once to upgrade the EX4500 and EX4550 member switches in the same mixed Virtual Chassis.

The **ex-4500** keyword also does not need to be specified when configuring automatic software update for a mixed EX4500 and EX4550 Virtual Chassis.

**ex-4600**—Specify a path to a Junos OS image for an EX4600 switch when enabling automatic software update for a mixed Virtual Chassis.

**qfx-3**—Specify a path to a Junos OS image for a QFX3500, QFX3600, or both types of devices when enabling automatic software update for a mixed VCF.

**qfx-5**—Specify a path to a Junos OS image for a QFX5100 device when enabling automatic software update for a mixed VCF.

|                           |                                                            |
|---------------------------|------------------------------------------------------------|
| <b>Required Privilege</b> | system—To view this statement in the configuration.        |
| <b>Level</b>              | system-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches</i></li> <li>• <a href="#">Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure) on page 5114</a></li> <li>• <a href="#">Understanding Software Upgrades in a Virtual Chassis Fabric on page 5258</a></li> </ul> |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## graceful-restart (Enabling Globally)

---

|                            |                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>graceful-restart {<br/>  disable;<br/>  helper-disable;<br/>  maximum-helper-recovery-time <i>seconds</i>;<br/>  maximum-helper-restart-time <i>seconds</i>;<br/>  notify-duration <i>seconds</i>;<br/>  recovery-time <i>seconds</i>;<br/>  restart-duration <i>seconds</i>;<br/>  stale-routes-time <i>seconds</i>;<br/>}</pre> |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options]                                            |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                         |
| <b>Description</b>         | Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.                                                                               |



### NOTE:

- For VPNs, the **graceful-restart** statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.
  - For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.
  - LDP sessions flap when **graceful-restart** configurations change.
- 

|                                 |                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | Graceful restart is disabled by default.                                                                                                           |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling Graceful Restart</i></li><li>• <i>Configuring Routing Protocols Graceful Restart</i></li></ul> |

- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Graceful Restart Configuration Statements*
- *Configuring Graceful Restart for QFabric Systems*

## graceful-switchover

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | graceful-switchover;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis redundancy</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | For switches with more than one Routing Engine, including those in a Virtual Chassis or a Virtual Chassis Fabric, configure the master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | Graceful Routing Engine switchover (GRES) is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Nonstop Active Routing on Switches on page 2514</a></li> <li>• <i>Configuring Graceful Routing Engine Switchover</i></li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 5116</a></li> <li>• <a href="#">Configuring Nonstop Active Routing on Switches on page 2505</a></li> <li>• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i></li> </ul> |

## id

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>id id;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure the alphanumeric string that identifies a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <i>id</i> —Virtual Chassis ID (VCID), which uses the ISO family address format—for example, <b>9622.6ac8.5345</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge</i></li><li>• <a href="#">Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure) on page 5116</a></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li><li>• <a href="#">Configuring an EX8200 Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Understanding Virtual Chassis Member ID Numbering in an EX8200 Virtual Chassis</a></li></ul> |

## location (Virtual Chassis)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>location location;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis member member-id</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Set a description of the location of the Virtual Chassis or VCF member switch or external Routing Engine.</p> <p>The <b>Location</b> field is visible to users who enter the <b>show virtual-chassis status detail</b> command.</p> <p>Setting this description has no effect on the operation of the member device.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>location</b> —Location of the current member switch or external Routing Engine. The <b>location</b> can be any single word.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</a></li> <li>• <a href="#">Example: Configuring a Preprovisioned Mixed EX4200 and EX4500 Virtual Chassis</a></li> <li>• <a href="#">Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</a></li> <li>• <a href="#">Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX8200 Virtual Chassis (CLI Procedure)</a></li> </ul> |

## mac-persistence-timer

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mac-persistence-timer [<i>minutes</i>   <b>disable</b>];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>disable</b> introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>The maximum timer limit changed from no maximum timer limit to 60 minutes in Junos OS Release 12.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                           |
| <b>Description</b>              | <p>Specify how long the Virtual Chassis or VCF continues to use the MAC address of the switch that was originally configured in the master role as the system MAC base address after the original master switch is removed from the Virtual Chassis or VCF. The system MAC base address does not change in the event of a switchover provided the switch originally configured in the master role remains a member of the Virtual Chassis or VCF.</p> <p>The maximum timer limit is 60 minutes starting in Junos OS Release 12.2. There are no minimum or maximum timer limits in prior Junos OS releases.</p> |
| <b>Default</b>                  | The MAC persistence timer is set to 10 minutes by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>minutes</b>—Time in minutes that the member switch in the backup role continues to use the system MAC base address of the old master before using its own system MAC base address after the switch in the master role is physically disconnected or removed from the Virtual Chassis or VCF.</p> <p><b>disable</b>—Disable the MAC persistence timer. The system MAC base address never changes when the MAC persistence timer is disabled, even when the switch in the master role is physically disconnected or removed from the Virtual Chassis or VCF.</p>                                           |
| <b>Required Privilege Level</b> | <p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis (CLI Procedure) on page 5113</a></li><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li></ul>                                                                                                                                                                                                                                 |



## mastership-priority

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mastership-priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis member</a> <i>member-id</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Mastership priority option <b>0</b> introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>The mastership priority value is the most important factor in determining the role of the member switch within a nonprovisioned Virtual Chassis or VCF configuration. Other factors (see <a href="#">“Understanding How the Master in a Virtual Chassis Is Elected” on page 5078</a>) also affect the election of the master.</p> <p>The mastership priority value takes the highest precedence in the master election algorithm. The member switch with highest mastership priority assumes the master Routing Engine role of the Virtual Chassis or VCF. Toggling back and forth between master and backup status in failover conditions is undesirable, so we recommend that you assign the same mastership priority value to both the master and the backup. Secondary factors in the master election algorithm determine which of these two members (that is, the two members that are assigned the highest mastership priority value) functions as the master of the Virtual Chassis or VCF.</p> <p>This statement is not used for the EX8200 Virtual Chassis, which determines mastership by external Routing Engine uptime. See <i>Understanding Virtual Chassis Roles in an EX8200 Virtual Chassis</i>.</p> <p>A switch with a mastership priority of <b>0</b> never takes the master or backup role.</p> |
| <b>Default</b>                  | 128                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><i>number</i>—Mastership priority value.</p> <p><b>Range:</b> 0 through 255</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX4300 Virtual Chassis (CLI Procedure) on page 5097</a></li> <li>• <a href="#">Example: Configuring an EX3300 Virtual Chassis with a Master and Backup</a></li> <li>• <a href="#">Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

- *Example: Configuring an EX4200 Virtual Chassis Interconnected Across Multiple Wiring Closets*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*

## member

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>member <i>member-id</i> {   location <i>location</i>;   mastership-priority <i>number</i>;   no-management-vlan;   serial-number <i>serial-number</i>;   role <i>role</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure a switch or an XRE200 External Routing Engine as a member of a Virtual Chassis or a Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | <p>When an EX Series switch or a QFX Series devices configured in standalone mode is powered on but not interconnected through its Virtual Chassis ports (VCPs) with other member switches, its default member ID is 0.</p> <p>There is no default member ID in an EX8200 or EX9200 Virtual Chassis. An EX8200 or EX9200 Virtual Chassis must be preprovisioned, and that process configures the member IDs.</p>                                                                                                                                                             |
| <b>Options</b>                  | <p><b><i>member-id</i></b>—Identifies a specific member switch of a Virtual Chassis or VCF configuration.</p> <p>The exact range for a specific Virtual Chassis or VCF depends on the number of switches allowed in the Virtual Chassis or VCF.</p> <p>In an EX8200 Virtual Chassis, member IDs 0 through 7 are reserved for EX8200 member switches and member IDs 8 and 9 are reserved for the master and backup external Routing Engines.</p> <p>The remaining statements are explained separately.</p>                                                                    |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</a></li> <li>• <a href="#">Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</a></li> </ul> |

- *Configuring an EX3300 Virtual Chassis (CLI Procedure)*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- *Configuring an EX8200 Virtual Chassis (CLI Procedure)*
- *Configuring an EX9200 Virtual Chassis*
- *Configuring a QFX Series Virtual Chassis (CLI Procedure)*

---

## no-management-vlan

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-management-vlan;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <b>virtual-chassis member</b> <i>member-id</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Remove the specified member's out-of-band management port from the virtual management Ethernet (VME) global management VLAN of the Virtual Chassis or VCF configuration.</p> <p>For a member that is functioning in a linecard role, you can use this configuration to reserve the member's management Ethernet port for local troubleshooting:</p> <pre>virtual-chassis {<br/>  member 2 {<br/>    no-management-vlan;<br/>  }<br/>}</pre> <p>You cannot configure the IP address for a local management Ethernet port using the CLI or the J-Web interface. To do this, you need to use the shell <b>ifconfig</b> command.</p> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Setting Up a Multimember EX4200 Virtual Chassis Access Switch with a Default Configuration</i></li><li>• <i>Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)</i></li><li>• <a href="#">Understanding Global Management of a Virtual Chassis on page 5080</a></li><li>• <a href="#">Understanding Virtual Chassis Fabric Configuration on page 5251</a></li></ul>                                                                                                                                     |

## no-split-detection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-split-detection;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Disable the split and merge feature in a Virtual Chassis or VCF configuration.</p> <p>We recommend using this statement to disable the split and merge feature when configuring a two-member Virtual Chassis. Enabling this statement on a two-member Virtual Chassis ensures that both switches remain in the correct Virtual Chassis roles in the event of a Virtual Chassis split.</p>                                                                                                                          |
| <b>Default</b>                  | The split and merge feature is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge</i></li> <li>• <a href="#">Disabling Split and Merge in a Virtual Chassis (CLI Procedure) on page 5114</a></li> <li>• <a href="#">Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure) on page 5116</a></li> <li>• <a href="#">Understanding Split and Merge in a Virtual Chassis on page 5087</a></li> </ul> |

## package-name

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>package-name <i>package-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit virtual-chassis <a href="#">auto-sw-update</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify the software package name or location of the software package to be used by the automatic software update feature for Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>                  | No package name is specified.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b><i>package-name</i></b>—Name of the software package or the URL to the software package to be used.</p> <ul style="list-style-type: none"><li>• If the software package is located on a local directory on the switch, use the following format for <b><i>package-name</i></b>:<br/><br/><b><i>/pathname/package-name</i></b></li><li>• If the software package is to be downloaded and installed from a remote location, use one of the following formats:<br/><br/><b><i>ftp://hostname/pathname/package-name</i></b><br/><b><i>ftp://username:prompt@ftp.hostname.net/package-name</i></b><br/><b><i>http://hostname/pathname/package-name</i></b></li></ul> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches</a></li><li>• <a href="#">Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure) on page 5114</a></li><li>• <a href="#">Understanding Software Upgrades in a Virtual Chassis Fabric on page 5258</a></li></ul>                                                                                                                                                                                                                                                                         |

## preprovisioned

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | preprovisioned;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Enable the preprovisioned configuration mode for a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.</p> <p>When the preprovisioned configuration mode is enabled, you cannot use the CLI or the J-Web interface to change the mastership priority or member ID of member switches.</p> <p>You must use this statement to configure an EX8200 Virtual Chassis. Nonprovisioned configuration of an EX8200 Virtual Chassis is not supported.</p>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <i>Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</i></li> <li>• <i>Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</i></li> <li>• <i>Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</i></li> <li>• <i>Configuring an EX8200 Virtual Chassis (CLI Procedure)</i></li> <li>• <i>Configuring an EX9200 Virtual Chassis</i></li> <li>• <i>Configuring a QFX Series Virtual Chassis (CLI Procedure)</i></li> <li>• <a href="#">Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 5104</a></li> </ul> |

## redundancy (Graceful Switchover)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>redundancy {<br/>    failover {<br/>        on-disk-failure;<br/>        on-loss-of-keepalives;<br/>    }<br/>    graceful-switchover;<br/>}</pre>                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Enable redundant Routing Engines on a Virtual Chassis with two or more member switches or on a Virtual Chassis Fabric, on a standalone EX6200 or EX8200 switch with more than one Routing Engine.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                   |
| <b>Default</b>                  | Redundancy is enabled for the Routing Engines.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">graceful-switchover on page 2527</a></li><li>• <a href="#">Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 5116</a></li><li>• <i>Configuring Graceful Routing Engine Switchover</i></li><li>• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i></li><li>• <i>High Availability Features for EX Series Switches Overview</i></li></ul> |



## role

|                            |                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>role (line-card   routing-engine);</code>                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | [edit <b>virtual-chassis</b> <b>preprovisioned member</b> <i>member-id</i> ]                                                                                                                                                               |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF). |
| <b>Description</b>         | Specify the roles of the members of the Virtual Chassis or a Virtual Chassis Fabric (VCF) in a preprovisioned Virtual Chassis.                                                                                                             |

### Virtual Chassis Fabric

Specify the role to be performed by each switch. In a VCF, the spine devices are configured into the Routing Engine role and the leaf devices are configured into the line card role. You can configure several devices into the Routine Engine role, but only two will operate in the Routing Engine role at a time. The role must be associated with the member's serial number.

### EX Series (except EX8200 Virtual Chassis) and QFX Series Virtual Chassis

Specify the role to be performed by each member switch. Associate the role with the member's serial number.

When you use a preprovisioned configuration, you cannot modify the mastership priority or member ID of member switches through the user interfaces. The mastership priority value is generated by the software, based on the assigned role:

- A member configured as **routing-engine** is assigned the mastership priority **129**.
- A member configured as **line-card** is assigned the mastership priority **0**.
- A member listed in the preprovisioned configuration without an explicitly specified role is assigned the mastership priority **128**.

The configured role specifications are permanent. If both **routing-engine** members fail, a **line-card** member cannot take over as master of the Virtual Chassis configuration. You must delete the preprovisioned configuration to change the specified roles in a Virtual Chassis.

Explicitly configure two members as **routing-engine** and configure additional switches as members of the preprovisioned Virtual Chassis by specifying only their serial numbers. If you do not explicitly configure the role of the additional members, they function in a linecard role by default. In that case, a member that is functioning in a linecard role can take over mastership if the members functioning as master and backup (**routing-engine** role) both fail.

### EX8200 Virtual Chassis

Specify the role to be performed by each XRE200 External Routing Engine and each EX8200 member switch. Associate the role with the member's serial number. An EX8200 Virtual Chassis cannot function when both external Routing Engines, which must be configured in the **routing-engine** role, have failed.

- Options**
- **line-card**—Enables the member to be eligible to function only in the linecard role. Any member of the Virtual Chassis or VCF configuration other than the master or backup functions in the linecard role and runs only a subset of Junos OS for EX Series switches. A member functioning in the linecard role does not run the control protocols or the chassis management processes.

A Virtual Chassis must have at least three members for one member to function in the linecard role.

In an EX8200 Virtual Chassis configuration, all member switches must be in the linecard role.

- **routing-engine**—Enables the member to function as a master or backup of the Virtual Chassis or VCF configuration. The master manages all members and runs the chassis management processes and control protocols. The backup synchronizes with the master in terms of protocol states, forwarding tables, and so forth, so that it is prepared to preserve routing information and maintain network connectivity without disruption in case the master is unavailable.

(All Virtual Chassis composed of EX Series switches, except EX8200 switches, or QFX Series devices) Specify two and only two members as **routing-engine**. The software determines which of the two members assigned the **routing-engine** role functions as master, based on the master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 5078](#). In these Virtual Chassis, the **routing-engine** role is associated with a switch.

(EX8200 Virtual Chassis) All XRE200 External Routing Engines must be in the **routing-engine** role.

|                                 |                                                            |
|---------------------------------|------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.        |
|                                 | system-control—To add this statement to the configuration. |

**Related  
Documentation**

- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- *Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File*
- *Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines*
- *Configuring an EX3300 Virtual Chassis (CLI Procedure)*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- *Configuring an EX8200 Virtual Chassis (CLI Procedure)*
- *Configuring an EX9200 Virtual Chassis*
- *Configuring a QFX Series Virtual Chassis (CLI Procedure)*
- *Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)*
- *Adding a New EX4200 Switch to an Existing EX4200 Virtual Chassis (CLI Procedure)*
- [Replacing a Member Switch of a Virtual Chassis Configuration \(CLI Procedure\) on page 5104](#)

## serial-number


---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>serial-number serial-number;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis preprovisioned member member-id</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>In a preprovisioned Virtual Chassis or Virtual Chassis Fabric (VCF), specify the serial number of each member switch to be included in the configuration. If you do not include the serial number within the configuration, the switch cannot be recognized as a member of a preprovisioned configuration.</p> <p>In an EX8200 Virtual Chassis configuration, specify the serial number of each XRE200 External Routing Engine and each EX8200 member switch to be included in the Virtual Chassis configuration. If you do not include the serial number within the Virtual Chassis configuration, the external Routing Engine or switch cannot be recognized as a member of the configuration.</p>                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>serial-number</b> —Permanent serial number for the external Routing Engine or for the member switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li><li>• <a href="#">Configuring an EX2200 Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Configuring an EX3300 Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Configuring an EX4300 Virtual Chassis (CLI Procedure) on page 5097</a></li><li>• <a href="#">Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Configuring an EX8200 Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Configuring an EX9200 Virtual Chassis</a></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)</a></li></ul> |

## serial-number (Virtual Chassis aliases)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>serial-number <i>serial-number</i> {<br/>    <i>alias-name alias-name</i>;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis aliases</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series Virtual Chassis and Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Specify the serial number that will be labeled with an alias in a Virtual Chassis or Virtual Chassis Fabric (VCF).<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b><i>serial-number</i></b> —Permanent serial number for the member switch in the Virtual Chassis or VCF.<br><br>You can retrieve the serial number for any device in your Virtual Chassis or VCF by entering the <b>show virtual-chassis</b> command and reviewing the output in the <b>Serial No</b> field.                                                                                                                                                          |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Understanding Virtual Chassis Fabric Components on page 5243</a></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis Components</a></li> </ul> |

## traceoptions (Virtual Chassis)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i> &lt;detail&gt; &lt;disable&gt; &lt;receive&gt; &lt;send&gt;; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>     | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>detail</b> added in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>         | Define tracing operations for the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>             | Tracing operations are disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>             | <p><b>detail</b>—(Optional) Generate detailed trace information for a flag.</p> <p><b>disable</b>—(Optional) Disable a flag.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> <li><b>all</b>—All tracing operations.</li> </ul> |
|                            | <div>  <p><b>TIP:</b> The <b>all</b> flag displays a subset of logs that are useful in debugging most issues. For more detailed information, use <b>all detail</b>.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|                            | <ul style="list-style-type: none"> <li><b>auto-configuration</b>—Trace Virtual Chassis ports (VCPs) that have been automatically configured.</li> <li><b>csn</b>—Trace Virtual Chassis complete sequence number (CSN) packets.</li> <li><b>error</b>—Trace Virtual Chassis errored packets.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- **hello**—Trace Virtual Chassis hello packets.
- **krt**—Trace Virtual Chassis KRT events.
- **lsp**—Trace Virtual Chassis link-state packets.
- **lsp-generation**—Trace Virtual Chassis link-state packet generation.
- **me**—Trace Virtual Chassis ME events.
- **normal**—Trace normal events.
- **packets**—Trace Virtual Chassis packets.
- **parse**—Trace reading of the configuration.
- **psn**—Trace partial sequence number (PSN) packets.
- **route**—Trace Virtual Chassis routing information.
- **spf**—Trace Virtual Chassis SPF events.
- **state**—Trace Virtual Chassis state transitions.
- **task**—Trace Virtual Chassis task operations.

**no-stamp**—(Optional) Do not place a timestamp on any trace file.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**receive**—(Optional) Trace received packets.

**replace**—(Optional) Replace a trace file rather than appending information to it.

**send**—(Optional) Trace transmitted packets.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.


|                           |                                                            |
|---------------------------|------------------------------------------------------------|
| <b>Required Privilege</b> | system—To view this statement in the configuration.        |
| <b>Level</b>              | system-control—To add this statement to the configuration. |

**Related  
Documentation**

- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157](#)
- [Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 5154](#)
- [Verifying That Virtual Chassis Ports Are Operational on page 5155](#)
- [Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis](#)
- [Troubleshooting an EX Series Virtual Chassis on page 5235](#)



## vcp-no-hold-time

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | vcp-no-hold-time;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Disable the Virtual Chassis port (VCP) holddown timer for all VCPs in the Virtual Chassis or Virtual Chassis Fabric (VCF).</p> <p>The VCP holddown timer is an internal mechanism that delays a Virtual Chassis reconvergence for several seconds when a VCP becomes inactive. The purpose of this delay is to provide the VCP time to return online without having to reconverge the Virtual Chassis to adjust to the inactive VCP. All traffic to the VCP is dropped while the VCP is inactive. If the VCP remains down for a time that exceeds the VCP holddown timer, a Virtual Chassis reconvergence occurs.</p> <p>When this statement is enabled, the VCP holddown timer is disabled and the Virtual Chassis reconvergence occurs when a VCP becomes inactive. The period of time where traffic is dropped waiting for the VCP to return online is avoided.</p> <p>We recommend enabling this statement after a Virtual Chassis is operational. We recommend disabling this statement when you are adding or removing member switches from your Virtual Chassis.</p> <p>The VCP holddown timer cannot be viewed and is not user-configurable. You can only control whether the VCP holddown timer is enabled or disabled by configuring this statement.</p> |
|                                 | <p> <b>NOTE:</b> For the EX4300 Virtual Chassis, you should enable the <code>vcp-no-hold-time</code> statement before performing a software upgrade using NSSU. If you do not enable the <code>vcp-no-hold-time</code> statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see <a href="#">“Understanding Split and Merge in a Virtual Chassis”</a> on page 5087</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>                  | The VCP holddown timer is enabled by default on all devices that support this statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding EX4300 Virtual Chassis on page 5065</a></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

- [Understanding EX Series Virtual Chassis Components on page 5067](#)
- *Understanding QFX Series Virtual Chassis Components*

## virtual-chassis

```
Syntax  virtual-chassis {
        aliases {
            serial-number serial-number {
                alias-name alias-name;
            }
        }
        auto-provisioned
        auto-sw-update {
            (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
            package-name package-name;
        }
        fast-failover (ge | vcp disable | xe);
        graceful-restart {
            disable;
        }
        id id;
        mac-persistence-timer [minutes | disable];;
        member member-id {
            location location;
            mastership-priority number;
            no-management-vlan;
            serial-number;
            role;
        }
        no-split-detection;
        preprovisioned;
        traceoptions (Virtual Chassis) {
            file filename <files number> <size size> <world-readable | no-world-readable> <match
                regex>;
            flag flag ;
        }
        vc-port {
            lag-hash (packet-based | source-port-based);
        }
        vcp-no-hold-time;
    }
```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

**Description** Configure a Virtual Chassis or a Virtual Chassis Fabric (VCF).

The remaining statements are explained separately.

**Default** A standalone EX Series switch is a Virtual Chassis by default. It has a default member ID of 0, a default mastership priority of 128, and a default role as master.

A QFX Series device configured in standalone mode is a Virtual Chassis by default. It has a default member ID of 0, a default mastership priority of 128, and a default role as master.

A standalone XRE200 External Routing Engine or EX8200 switch is not part of an EX8200 Virtual Chassis until a Virtual Chassis configuration is set up.

|                                 |                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration. |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------|

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li><li>• <i>Configuring a QFX Series Virtual Chassis (CLI Procedure)</i></li><li>• <i>Example: Configuring an EX3300 Virtual Chassis with a Master and Backup</i></li><li>• <i>Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet</i></li><li>• <i>Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</i></li><li>• <i>Configuring an EX3300 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX8200 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX9200 Virtual Chassis</i></li></ul> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

# Administration

- [Routine Monitoring on page 5149](#)
- [Operational Commands on page 5160](#)

## Routine Monitoring

---

- [Command Forwarding Usage with an EX Series Virtual Chassis on page 5149](#)
- [Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 5154](#)
- [Verifying That Virtual Chassis Ports Are Operational on page 5155](#)
- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157](#)
- [Verifying That Graceful Routing Engine Switchover Is Working in the Virtual Chassis on page 5159](#)

## Command Forwarding Usage with an EX Series Virtual Chassis

Some CLI commands can be run either on all members or on a specific member of a Virtual Chassis configuration. This functionality is referred to as command forwarding.

You can always specify that these commands be applied to all member switches in the Virtual Chassis by using the **all-members** option, or to a specific member switch by using the **member-member-id** option. If neither option is specified, the default command forwarding behavior, which varies by command, is used. See the **Default** row in [Table 570 on page 5150](#) to learn the command forwarding behavior for a specific command.

For example, to collect information about a particular member switch prior to contacting Juniper Networks Technical Assistance Center (JTAC), use the **request support information member member-id** command to gather data for the specified member switch. If you want to gather this data for all member switches in the Virtual Chassis, you can enter the **request support information** command, which by default uses the **all-members** option, or the **request support information all-members** command.

[Table 570 on page 5150](#) provides a list of commands that can be run either on all members of the Virtual Chassis configuration or on a specific member switch.

Table 570: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration

| Commands Available for Command Forwarding | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | all-members                                                                   | member-member-id                                         | Default            |
|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|----------------------------------------------------------|--------------------|
| <b>request support information</b>        | <p>Use this command when you contact JTAC about your component problem. This command is the equivalent of using the following CLI commands:</p> <ul style="list-style-type: none"> <li>• <b>show version</b></li> <li>• <b>show chassis firmware</b></li> <li>• <b>show chassis hardware</b></li> <li>• <b>show chassis environment</b></li> <li>• <b>show interfaces extensive</b> (for each configured interface)</li> <li>• <b>show configuration</b> (excluding any SECRET-DATA)</li> <li>• <b>show system virtual-memory</b></li> </ul> | Displays information for all members of the Virtual Chassis configuration.    | Displays information for the specified member switch.    | <b>all-members</b> |
| <b>request system partition hard-disk</b> | Set up the hard disk for partitioning. After this command is issued, the hard disk is partitioned the next time the system is rebooted. When the hard disk is partitioned, the contents of <b>/altroot</b> and <b>/altconfig</b> are saved and restored. All other data on the hard disk is at risk of being lost.                                                                                                                                                                                                                           | Partitions the hard disk on all members of the Virtual Chassis configuration. | Partitions the hard disk on the specified member switch. | <b>all-members</b> |
| <b>request system reboot</b>              | Reboot Junos OS for EX Series switches after a software upgrade and occasionally to recover from an error condition.                                                                                                                                                                                                                                                                                                                                                                                                                         | Reboots all members of the Virtual Chassis configuration.                     | Reboots the specified member switch.                     | <b>all-members</b> |

Table 570: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration (*continued*)

| Commands Available for Command Forwarding | Purpose                                                                                                                                      | all-members                                                                    | member-member-id                                         | Default            |
|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|----------------------------------------------------------|--------------------|
| <b>request system snapshot</b>            | Back up the currently running and active file system.                                                                                        | Backs up the file systems on all members of the Virtual Chassis configuration. | Backs up the file system on the specified member switch. | <b>all-members</b> |
| <b>request system storage cleanup</b>     | Free storage space on the switch by rotating log files and proposing a list of files for deletion. User input is required for file deletion. | Runs cleanup on all members of the Virtual Chassis configuration.              | Runs cleanup on the specified member switch.             | <b>all-members</b> |
| <b>show log user</b>                      | Display users who are viewing the system log.                                                                                                | Displays information for all members of the Virtual Chassis configuration.     | Displays information for the specified member switch.    | master switch only |
| <b>show system alarms</b>                 | Display active system alarms.                                                                                                                | Displays information for all members of the Virtual Chassis configuration.     | Displays information for the specified member switch.    | <b>all-members</b> |
| <b>show system audit</b>                  | Display the state and checksum values for file systems.                                                                                      | Displays information for all members of the Virtual Chassis configuration.     | Displays information for the specified member switch.    | <b>all-members</b> |
| <b>show system boot-messages</b>          | Display initial messages generated by the system kernel upon startup. These messages are the contents of <code>/var/run/dmesg.boot</code> .  | Displays information for all members of the Virtual Chassis configuration.     | Displays information for the specified member switch.    | <b>all-members</b> |

Table 570: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration (*continued*)

| Commands Available for Command Forwarding | Purpose                                                                                                                                                                                                                                                                                                                                                           | all-members                                                                | member-member-id                                      | Default            |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|-------------------------------------------------------|--------------------|
| <b>show system buffers</b>                | Display information about the buffer pool that the Routing Engine uses for local traffic. Local traffic is the routing and management traffic that is exchanged between the Routing Engine and the Packet Forwarding Engine within the switch, as well as the routing and management traffic from IP (that is, from OSPF, BGP, SNMP, ping operations, and so on). | Displays information for all members of the Virtual Chassis configuration. | Displays information for the specified member switch. | <b>all-members</b> |
| <b>show system connections</b>            | Display information about the active IP sockets on the Routing Engine. Use this command to verify which servers are active on a system and which connections are currently in progress.                                                                                                                                                                           | Displays information for all members of the Virtual Chassis configuration. | Displays information for the specified member switch. | <b>all-members</b> |
| <b>show system core-dumps</b>             | Display a core file generated by an internal Junos OS process.                                                                                                                                                                                                                                                                                                    | Displays information for all members of the Virtual Chassis configuration. | Displays information for the specified member switch. | <b>all-members</b> |
| <b>show system directory-usage</b>        | Display directory usage information.                                                                                                                                                                                                                                                                                                                              | Displays information for all members of the Virtual Chassis configuration. | Displays information for the specified member switch. | master switch only |
| <b>show system processes</b>              | Display information about software processes that are running on the switch and that have controlling terminals.                                                                                                                                                                                                                                                  | Displays information for all members of the Virtual Chassis configuration. |                                                       | <b>all-members</b> |



**Table 570: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration (*continued*)**

| Commands Available for Command Forwarding | Purpose                                                                                                                                                                                                         | all-members                                                                                      | member-member-id                                                            | Default            |
|-------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------|--------------------|
| <b>show system reboot</b>                 | Display pending system reboots or halts.                                                                                                                                                                        | Displays information for all members of the Virtual Chassis configuration.                       | Displays information for the specified member switch.                       | <b>all-members</b> |
| <b>show system snapshot</b>               | Display information about the backup software that is located in the <code>/altroot</code> and <code>/altconfig</code> file systems. To back up software, use the <code>request system snapshot</code> command. | Displays information for all members of the Virtual Chassis configuration.                       | Displays information for the specified member switch.                       | <b>all-members</b> |
| <b>show system software</b>               | Display the Junos OS extensions loaded on your switch.                                                                                                                                                          | Displays information for all members of the Virtual Chassis configuration.                       | Displays information for the specified member switch.                       | <b>all-members</b> |
| <b>show system statistics</b>             | Display systemwide protocol-related statistics.                                                                                                                                                                 | Displays information for all members of the Virtual Chassis configuration.                       | Displays information for the specified member switch.                       | <b>all-members</b> |
| <b>show system storage</b>                | Display statistics about the amount of free disk space in the switch's file systems.                                                                                                                            | Displays information for all members of the Virtual Chassis configuration.                       | Displays information for the specified member switch.                       | <b>all-members</b> |
| <b>show system uptime</b>                 | Display the current time and information about how long the switch, the switch software, and any existing protocols have been running                                                                           | Displays information for all members of the Virtual Chassis configuration.                       | Displays information for the specified member switch.                       | <b>all-members</b> |
| <b>show system users</b>                  | Show all users who are currently logged in.                                                                                                                                                                     | Shows all users who are currently logged in to any members of the Virtual Chassis configuration. | Shows all users who are currently logged in to the specified member switch. | <b>all-members</b> |

Table 570: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration (*continued*)

| Commands Available for Command Forwarding | Purpose                                                                                                                                                                         | all-members                                                                | member-member-id                                      | Default            |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|-------------------------------------------------------|--------------------|
| <b>show system virtual-memory</b>         | Display the usage of Junos OS kernel memory, listed first by size of allocation and then by type of usage. Use <b>show system virtual-memory</b> for troubleshooting with JTAC. | Displays information for all members of the Virtual Chassis configuration. | Displays information for the specified member switch. | <b>all-members</b> |

**Related Documentation**

- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157](#)
- [Understanding EX Series Virtual Chassis Components on page 5067](#)
- [Junos OS System Basics and Services Command Reference](#)

**Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member**

**Purpose** You can designate the role that a member performs within a Virtual Chassis or you can allow the role to be assigned by default. You can designate the member ID that is assigned to a specific switch by creating a permanent association between the switch's serial number and a member ID, using a preprovisioned configuration. Or you can let the member ID be assigned by the master, based on the sequence in which the member switch is powered on and on which member IDs are currently available.

The role and member ID of the member switch are displayed on the front-panel LCD.

Each member switch can be cabled to one or two other member switches, using either the dedicated Virtual Chassis ports (VCPs) on the rear panel, an uplink port that has been configured as a VCP, or an optical port that has been configured as a VCP. The members that are cabled together are considered neighbor members.

**Action** To display the role and member ID assignments using the CLI:

```
user@switch> show virtual-chassis
```

```
Virtual Chassis ID: 0000.e255.00e0
```

| Member ID | Status | Serial No | Model      | Mastership Priority | Role    | Neighbor List ID, Interface |
|-----------|--------|-----------|------------|---------------------|---------|-----------------------------|
| 0 (FPC 0) | Prsnt  | abc123    | ex4200-48p | 255                 | Master* | 1 vcp-0<br>2 vcp-1          |
| 1 (FPC 1) | Prsnt  | def456    | ex4200-24t | 255                 | Backup  | 2 vcp-0<br>0 vcp-1          |

```

2 (FPC 2)  Prsnt  abd231      ex4200-24p   128  Linecard    0 vcp-0
   1 vcp-1

```

**Meaning** This output verifies that three EX4200 switches have been interconnected as a Virtual Chassis configuration through their dedicated VCPs to create an EX4200 Virtual Chassis. The display shows which of the VCPs is connected to which neighbor. The first port (**vcp-0**) of member **0** is connected to member **1** and the second port of member **0** (**vcp-1**) is connected to member **2**. The FPC slots for the switches are the same as the member IDs.

The **Mastership Priority** values indicate that the master and backup members have been explicitly configured, because they are not using the default value (**128**).



**NOTE:** This example uses output from an EX4200 Virtual Chassis. The output, with the exception of the **Model** column, would be identical on all other Virtual Chassis.

**Related Documentation**

- [Configuring Mastership of a Virtual Chassis \(CLI Procedure\) on page 5107](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a Virtual Chassis on an EX Series Switch \(J-Web Procedure\)](#)
- [Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches \(CLI Procedure\)](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\)](#)
- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157](#)

## Verifying That Virtual Chassis Ports Are Operational

**Purpose** Display the status of Virtual Chassis ports (VCPs) in a Virtual Chassis or Virtual Chassis Fabric (VCF).



**NOTE:** VCPs are not displayed when you issue the **show interfaces** command.

**Action** Display the VCPs:

```
user@switch> show virtual-chassis vc-port all-members
```

```
fpc0:
```

| Interface<br>or<br>PIC / Port | Type | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID Interface |
|-------------------------------|------|-------------|--------|-----------------|--------------------------|
|-------------------------------|------|-------------|--------|-----------------|--------------------------|

|       |            |   |    |       |   |              |
|-------|------------|---|----|-------|---|--------------|
| vcp-0 | Dedicated  | 1 | Up | 32000 | 1 | vcp-0        |
| vcp-1 | Dedicated  | 2 | Up | 32000 | 1 | vcp-1        |
| 1/0   | Configured | 3 | Up | 1000  | 2 | vcp-255/1/0  |
| 1/1   | Configured | 3 | Up | 1000  | 2 | vcp-255/1/1  |
| 1/2   | Configured | 4 | Up | 1000  | 4 | vcp-255/0/20 |
| 1/3   | Configured | 4 | Up | 1000  | 4 | vcp-255/0/21 |

## fpc1:

| Interface<br>or<br>PIC / Port | Type       | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface   |
|-------------------------------|------------|-------------|--------|-----------------|----------------|-------------|
| vcp-0                         | Dedicated  | 1           | Up     | 32000           | 0              | vcp-0       |
| vcp-1                         | Dedicated  | 2           | Up     | 32000           | 0              | vcp-1       |
| 1/0                           | Configured | 3           | Up     | 10000           | 3              | vcp-255/1/0 |
| 1/1                           | Configured | 3           | Up     | 10000           | 3              | vcp-255/1/1 |

## fpc2:

| Interface<br>or<br>PIC / Port | Type       | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface   |
|-------------------------------|------------|-------------|--------|-----------------|----------------|-------------|
| vcp-0                         | Dedicated  | 1           | Up     | 32000           | 3              | vcp-0       |
| vcp-1                         | Dedicated  | 2           | Up     | 32000           | 3              | vcp-1       |
| 1/0                           | Configured | 3           | Up     | 1000            | 0              | vcp-255/1/0 |
| 1/1                           | Configured | 3           | Up     | 1000            | 0              | vcp-255/1/1 |
| 1/2                           |            | -1          | Down   | 1000            |                |             |
| 1/3                           |            | -1          | Down   | 1000            |                |             |

## fpc3:

| Interface<br>or<br>PIC / Port | Type       | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface   |
|-------------------------------|------------|-------------|--------|-----------------|----------------|-------------|
| vcp-0                         | Dedicated  | 1           | Up     | 32000           | 2              | vcp-0       |
| vcp-1                         | Dedicated  | 2           | Up     | 32000           | 2              | vcp-1       |
| 1/0                           | Configured | 3           | Up     | 10000           | 1              | vcp-255/1/0 |
| 1/1                           | Configured | 3           | Up     | 10000           | 1              | vcp-255/1/1 |

## fpc4:

| Interface<br>or<br>PIC / Port | Type       | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface   |
|-------------------------------|------------|-------------|--------|-----------------|----------------|-------------|
| vcp-0                         | Dedicated  | 1           | Down   | 32000           |                |             |
| vcp-1                         | Dedicated  | 2           | Down   | 32000           |                |             |
| 0/20                          | Configured | 3           | Up     | 1000            | 0              | vcp-255/1/2 |
| 0/21                          | Configured | 3           | Up     | 1000            | 0              | vcp-255/1/3 |

**Meaning** The dedicated VCPs are displayed as **vcp-0** and **vcp-1**. The uplink interfaces that have been set as uplink VCPs are displayed as **1/0**, **1/1**, **1/2**, and **1/3**. The network interfaces that have been set as VCPs are displayed as **0/20** and **0/21**. The neighbor interface names of uplink and network VCPs are of the form **vcp-255/pic/port**—for example, **vcp-255/1/0**. In that name, **vcp-255** indicates that the interface is a VCP, **1** is the uplink PIC number, and **0** is the port number. The **fpc** number is the same as the member ID. The trunk ID is a positive number ID assigned to the link aggregation group (LAG) formed by the Virtual Chassis. If no LAG is formed, the value is **-1**.



**NOTE:** This example uses output from an EX4200 Virtual Chassis. The output is similar on all other types of Virtual Chassis or for a VCF.

#### Related Documentation

- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157](#)
- *Configuring an EX3300 Virtual Chassis (CLI Procedure)*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- *Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)*
- *Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)*

## Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis

#### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to view the following information about the switches and the ports on EX2200, EX3300, EX4200, EX4300, EX4500, EX4550, and EX8200 switches that are members of a Virtual Chassis:

- Member details and how members are connected with each other
- Traffic statistics for Virtual Chassis ports (VCPs) of the selected members
- Details of the VCP packet counters

#### Action

To view Virtual Chassis monitoring details in the J-Web interface for a Virtual Chassis, select **Monitor > Virtual Chassis**.



**NOTE:** Virtual Chassis monitoring is supported on J-Web, on all Virtual Chassis platforms except EX3200 and EX6210.

To view member details for all members in the CLI, enter the following command:

```
user@switch> show virtual-chassis
```

To view VCP traffic statistics for a specific member in the CLI, enter the following command:

```
user@switch> show virtual-chassis vc-port statistics member member-id
```

To view the path a packet takes when going from a source interface to a destination interface in a Virtual Chassis configuration using the CLI, enter the following command:

```
user@switch> show virtual-chassis vc-path
```

**Meaning** In the J-Web interface, the top half of the screen displays details of the Virtual Chassis configuration, such as:

- Member
- Role
- Status
- Interface
- Type
- Speed
- Neighboring Member ID
- Link Status
- Error count



**NOTE:** If the member switch in the Virtual Chassis is not provisioned, the member ID will be displayed as -.

---

Click the **Stop** button to stop fetching values from the switch, and click the **Start** button to start plotting data again from the point where it was stopped.

To view a graph of the statistics for the selected VCP of the member, click **Show Graph**.

**Refresh Interval (sec)**—Displays the time interval you have set for page refresh.

Click **Clear Statistics** to clear the monitoring statistics for the selected member switch. You can specify the interval at which the member details and statistics must be refreshed.

The bottom half of the screen displays a chart of the Virtual Chassis statistics and the port packet counters.

For details about the output from CLI commands, see the [show virtual-chassis](#) and [show virtual-chassis vc-port statistics](#) command summaries.

**Related Documentation**

- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a Virtual Chassis on an EX Series Switch \(J-Web Procedure\)](#)
- [Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches \(CLI Procedure\)](#)
- [Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 5154](#)

## Verifying That Graceful Routing Engine Switchover Is Working in the Virtual Chassis

**Purpose** Verify that a Graceful Routing Engine switchover (GRES) between two member switches acting as the master and backup routing engines in a Virtual Chassis has occurred.

**Action** On the master switch, verify the member ID of the backup Routing Engine:

```
{master:0}
user@switch> show virtual-chassis
Virtual Chassis ID: 5efa.4b7a.aae6
```

| Member ID | Status | Serial No    | Model      | Mastership<br>priority | Role    | Neighbor List<br>ID | Interface |
|-----------|--------|--------------|------------|------------------------|---------|---------------------|-----------|
| 0 (FPC 0) | Prsnt  | BM0208105281 | ex4200-24t | 255                    | Master* | 1                   | vcp-0     |
| 1 (FPC 1) | Prsnt  | BP0208192350 | ex4200-48t | 255                    | Backup  | 0                   | vcp-0     |

Member ID for next new member: 2 (FPC 2)

1. Connect to the backup Routing Engine:

```
{master:0}
user@switch> request session member 1
```

```
{backup:1}
user@switch>
```

2. Verify that the backup Routing Engine is ready for switchover on member ID 1:

```
{backup:1}
user@switch> show system switchover
```

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

3. Switch the current backup Routing Engine to master Routing Engine:



**NOTE:** You must wait a minimum of two minutes between Routing Engine failovers for the Routing Engines to synchronize.

```
{backup:1}
user@switch> request chassis routing-engine master acquire
```

4. Verify that the master and backup Routing Engines have switched roles:



**NOTE:** Member ID 1 is now the master and member ID 0 is now the backup.

```
{master:1}
user@switch> show virtual-chassis
```

Virtual Chassis ID: 5efa.4b7a.aae6

|           |        |              |            | Mastership |         | Neighbor List |           |
|-----------|--------|--------------|------------|------------|---------|---------------|-----------|
| Member ID | Status | Serial No    | Model      | priority   | Role    | ID            | Interface |
| 0 (FPC 0) | Prsnt  | BM0208105281 | ex4200-24t | 255        | Backup  | 1             | vcp-0     |
| 1 (FPC 1) | Prsnt  | BP0208192350 | ex4200-48t | 255        | Master* | 0             | vcp-0     |

Member ID for next new member: 2 (FPC 2)

**Meaning** With graceful Routing Engine switchover enabled, when you initiated a switchover from the backup Routing Engine, the backup Routing Engine became the master and the master Routing Engine became the backup.

**Related Documentation**

- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 5116](#)

---

## Operational Commands

- [clear virtual-chassis vc-port statistics](#)
- [request session member](#)
- [request virtual-chassis recycle](#)
- [request virtual-chassis renumber](#)
- [request virtual-chassis vc-port](#)
- [request virtual-chassis vc-port diagnostics optics](#)
- [show system uptime](#)
- [show virtual-chassis active-topology](#)
- [show virtual-chassis device-topology](#)
- [show virtual-chassis login](#)
- [show virtual-chassis protocol adjacency](#)
- [show virtual-chassis protocol database](#)
- [show virtual-chassis protocol interface](#)
- [show virtual-chassis protocol route](#)
- [show virtual-chassis protocol statistics](#)
- [show virtual-chassis](#)
- [show virtual-chassis vc-path](#)
- [show virtual-chassis vc-port](#)
- [show virtual-chassis vc-port diagnostics optics](#)
- [show virtual-chassis vc-port statistics](#)



## clear virtual-chassis vc-port statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear virtual-chassis vc-port statistics<br><all-members><br><interface-name><br><local><br><member member-id>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>The options <b>all-members</b> and <b>local</b> were added in Junos OS Release 9.3 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric.                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Clear—reset to zero (0)—the traffic statistics counters on Virtual Chassis ports (VCPs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Clear traffic statistics for the specified VCP.</p> <p><b>local</b>—(Optional) Clear traffic statistics for VCPs from the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Clear traffic statistics for VCPs from the specified member of a Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis vc-port statistics on page 5227</a></li> <li>• <a href="#">show virtual-chassis vc-port on page 5209</a></li> <li>• <a href="#">Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157</a></li> </ul>                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">clear virtual-chassis vc-port statistics (EX4200 Virtual Chassis) on page 5161</a><br><a href="#">clear virtual-chassis vc-port statistics (EX8200 Virtual Chassis) on page 5162</a><br><a href="#">clear virtual-chassis vc-port statistics member 3 on page 5162</a>                                                                                                                                                                                                                                                                                                           |

### Sample Output

#### clear virtual-chassis vc-port statistics (EX4200 Virtual Chassis)

```
user@switch> clear virtual-chassis vc-port statistics
fpc0:
-----
Statistics cleared
```

### clear virtual-chassis vc-port statistics (EX8200 Virtual Chassis)

```
user@external-routing-engine> clear virtual-chassis vc-port statistics
```

```
member0:
```

```
-----  
Statistics cleared
```

```
member1:
```

```
-----  
Statistics cleared
```

```
member8:
```

```
-----  
Statistics cleared
```

```
member9:
```

```
-----  
Statistics cleared
```

### clear virtual-chassis vc-port statistics member 3

```
user@switch> clear virtual-chassis vc-port statistics member 3
```

```
Cleared statistics on member 3
```

---

## request session member

---

|                                 |                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>request session member <i>member-id</i></code>                                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                          |
| <b>Description</b>              | Start a session with the specified member of a Virtual Chassis or a VCF.                                                                                                                                                                                      |
| <b>Options</b>                  | <i>member-id</i> —Member ID for the specific member of the Virtual Chassis or VCF.                                                                                                                                                                            |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">member on page 5131</a></li><li>• <a href="#">Understanding EX Series Virtual Chassis Components on page 5067</a></li><li>• <a href="#">Understanding QFX Series Virtual Chassis Components</a></li></ul> |

## request virtual-chassis recycle

---

**Syntax**    `request virtual-chassis recycle member-id member-id`

**Release Information**    Command introduced in Junos OS Release 9.0 for EX Series switches.  
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.

**Description**    Make a previously used member ID available for reassignment.

When you remove a member switch from the Virtual Chassis configuration, the master reserves that member ID. To make the member ID available for reassignment, you must use this command.



**NOTE:** You must run this command from the Virtual Chassis member in the master role.

**Options**    `member-id member-id`—Specify the member ID that you want to make available for reassignment to a different member.

**Required Privilege Level**    system-control

**Related Documentation**

- [request virtual-chassis renumber on page 5165](#)
- [Replacing a Member Switch of a Virtual Chassis Configuration \(CLI Procedure\) on page 5104](#)
- [Adding or Replacing a Member Switch or an External Routing Engine in an EX8200 Virtual Chassis \(CLI Procedure\)](#)

**List of Sample Output**    [request virtual-chassis recycle member-id 3 on page 5164](#)  
[request virtual-chassis recycle member-id 1 on page 5164](#)

### Sample Output

[request virtual-chassis recycle member-id 3](#)


```
user@switch> request virtual-chassis recycle member-id 3
```

### Sample Output

[request virtual-chassis recycle member-id 1](#)

```
user@external-routing-engine> request virtual-chassis recycle member-id 1
```

## request virtual-chassis renumber

|                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                             | <code>request virtual-chassis renumber member-id <i>old-member-id</i> new-member-id <i>new-member-id</i></code>                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                                                                                                                                                                | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.                                                                                                                                                                                                                                    |
| <b>Description</b>                                                                                                                                                                        | Renumber a member of a Virtual Chassis configuration.                                                                                                                                                                                                                                                                                                                           |
| <div>  <b>NOTE:</b> You must run this command from the Virtual Chassis member in the master role. </div> |                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                                                                                                                                                                            | <p><code>member-id <i>old-member-id</i></code>—Specify the ID of the member that you wish to renumber.</p> <p><code>new-member-id <i>new-member-id</i></code>—Specify an unassigned member ID.</p>                                                                                                                                                                              |
| <b>Required Privilege Level</b>                                                                                                                                                           | system-control                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                                                                                                                                                              | <ul style="list-style-type: none"> <li>• <a href="#">request virtual-chassis recycle on page 5164</a></li> <li>• <a href="#">Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 5104</a></li> <li>• <a href="#">Adding or Replacing a Member Switch or an External Routing Engine in an EX8200 Virtual Chassis (CLI Procedure)</a></li> </ul> |
| <b>List of Sample Output</b>                                                                                                                                                              | <a href="#">request virtual-chassis renumber member-id 5 new-member-id 4 on page 5165</a><br><a href="#">request virtual-chassis renumber member-id 1 new-member-id 0 on page 5165</a>                                                                                                                                                                                          |

### Sample Output

`request virtual-chassis renumber member-id 5 new-member-id 4`

```
user@switch> request virtual-chassis renumber member-id 5 new-member-id 4
```

`request virtual-chassis renumber member-id 1 new-member-id 0`

```
user@external-routing-engine> request virtual-chassis renumber member-id 1 new-member-id 0
```

## request virtual-chassis vc-port

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>request virtual-chassis vc-port set   delete</b> <fpc-slot <i>fpc-slot</i> > pic-slot <i>pic-slot</i> port <i>port-number</i> <member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Option <b>fpc-slot</b> introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Enable or disable an optical port as a Virtual Chassis port (VCP).</p> <p>If you omit <b>member <i>member-id</i></b>, this command defaults to enabling or disabling the uplink VCP or SFP network port configured as a VCP on the switch where the command is issued.</p> <p>On an EX3300 switch, uplink ports 2 and 3 are configured as VCPs by default. No other uplink ports on any other EX Series switches are configured as VCPs by default.</p> <p>You might experience a temporary traffic disruption immediately after creating or deleting a user-configured VCP in an EX8200 Virtual Chassis.</p>                |
| <b>Options</b>                  | <p><b>pic-slot <i>pic-slot</i></b>—Number of the PIC slot for the port on the switch.</p> <p><b>port <i>port-number</i></b>—Number of the port that is to be enabled or disabled as a VCP.</p> <p><b>member <i>member-id</i></b>—(Optional) Enable or disable the specified VCP on the specified member of the Virtual Chassis or VCF.</p>                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | system-control                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">request virtual-chassis vc-port</a> (dedicated port)</li><li>• <a href="#">show virtual-chassis vc-port on page 5209</a></li><li>• <a href="#">show virtual-chassis vc-port statistics on page 5227</a></li><li>• <a href="#">clear virtual-chassis vc-port statistics on page 5161</a></li><li>• <a href="#">Virtual Chassis Port (VCP) Interface Names in an EX8200 Virtual Chassis</a></li><li>• <a href="#">Understanding EX Series Virtual Chassis Components on page 5067</a></li><li>• <a href="#">Understanding QFX Series Virtual Chassis Components</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">request virtual-chassis vc-port set pic-slot 1 port 0 on page 5167</a><br><a href="#">request virtual-chassis vc-port set pic-slot 1 port 1 member 3 on page 5167</a><br><a href="#">request virtual-chassis vc-port delete pic-slot 1 port 1 member 3 on page 5167</a>                                                                                                                                                                                                                                                                                                                                             |

## Sample Output

**request virtual-chassis vc-port set pic-slot 1 port 0**

user@switch> **request virtual-chassis vc-port set pic-slot 1 port 0**

To check the results of this command, use the **show virtual-chassis vc-port** command.

**request virtual-chassis vc-port set pic-slot 1 port 1 member 3**

user@switch> **request virtual-chassis vc-port set pic-slot 1 port 1 member 3**

To check the results of this command, use the **show virtual-chassis vc-port** command.

**request virtual-chassis vc-port delete pic-slot 1 port 1 member 3**

user@switch> **request virtual-chassis vc-port delete pic-slot 1 port 1 member 3**

To check the results of this command, use the **show virtual-chassis vc-port** command.

## request virtual-chassis vc-port diagnostics optics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>request virtual-chassis vc-port diagnostics optics</b>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Run a digital optical monitoring (DOM) scan on the optical ports configured as Virtual Chassis ports (VCPs).</p> <p>Enter the <b>show virtual-chassis vc-port diagnostics optics</b> command to view the results of the diagnostic scan.</p> <p>On certain EX Series switches, the <b>request virtual-chassis vc-port diagnostics optics</b> command must be entered to run a diagnostic scan before you can gather the <b>show virtual-chassis vc-port diagnostics optics</b> output.</p> |
| <b>Required Privilege Level</b> | system-control                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show virtual-chassis vc-port diagnostics optics on page 5213</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                |

## Sample Output

### request virtual-chassis vc-port diagnostics optics

```
user@switch> request virtual-chassis vc-port diagnostics optics
fpc0:
-----
vc-port Diagnostics Optics Done
```



## show system uptime

|                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                 | <a href="#">Syntax on page 5169</a><br><a href="#">Syntax (EX Series Switches) on page 5169</a><br><a href="#">Syntax (QFX Series) on page 5169</a><br><a href="#">Syntax (TX Matrix Router) on page 5169</a><br><a href="#">Syntax (TX Matrix Plus Router) on page 5169</a><br><a href="#">Syntax (MX Series Router) on page 5169</a>                                                                                                                                                                                                   |
| <b>Syntax</b>                         | show system uptime                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switches)</b>    | show system uptime<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (QFX Series)</b>            | show system uptime<br><director-group <i>name</i> ><br><infrastructure <i>name</i> ><br><interconnect-device <i>name</i> ><br><node-group <i>name</i> >                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (TX Matrix Router)</b>      | show system uptime<br><all-chassis   all-lcc   lcc <i>number</i>   scc>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (TX Matrix Plus Router)</b> | show system uptime<br><detail><br><all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (MX Series Router)</b>      | show system uptime<br><all-members><br><invoke-on><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>            | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>sfc</b> option introduced for the TX Matrix Plus router in JUNOS Release 9.6.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>                    | Display the current time and information about how long the router or switch, router or switch software, and routing protocols have been running.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                        | <b>none</b> —Show time since the system rebooted and processes started.<br><br><b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started on all the routers in the chassis.<br><br><b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus |

router, show time since the system rebooted and processes started for all connected T1600 or T4000 LCCs.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on all members of the Virtual Chassis configuration.

**director-group *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Director group.

**infrastructure *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the fabric control Routing Engine and fabric manager Routing Engine.

**interconnect-device *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Interconnect device.

**invoke-on**—(MX Series routers only) (Optional) Display the time since the system rebooted and processes started on the master Routing Engine, backup Routing Engine, or both, on a router with two Routing Engines.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show time since the system rebooted and processes started for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Node group.

**scc**—(TX Matrix routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system uptime** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation**

- [Monitoring System Process Information](#)
- [Monitoring System Properties](#)
- [10-Gigabit Ethernet LAN/WAN PIC with XFP \(T640 Router\)](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

- [show system uptime on page 5172](#)
- [show system uptime all-lcc \(TX Matrix Router\) on page 5172](#)
- [show system uptime all-lcc \(TX Matrix Plus Router\) on page 5172](#)
- [show system uptime \(EX Series\) on page 5173](#)
- [show system uptime \(QFX Series\) on page 5173](#)

**Output Fields** [Table 19 on page 223](#) describes the output fields for the **show system uptime** command. Output fields are listed in the approximate order in which they appear.

**Table 571: show system uptime Output Fields**

| Field Name               | Field Description                                                                                                                 |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <b>Current time</b>      | Current system time in UTC.                                                                                                       |
| <b>System booted</b>     | Date and time when the Routing Engine on the router or switch was last booted and how long it has been running.                   |
| <b>Protocols started</b> | Date and time when the routing protocols were last started and how long they have been running.                                   |
| <b>Last configured</b>   | Date and time when a configuration was last committed. Also shows the name of the user who issued the last <b>commit</b> command. |
| <b>time and up</b>       | Current time, in the local time zone, and how long the router or switch has been operational.                                     |
| <b>users</b>             | Number of users logged in to the router or switch.                                                                                |
| <b>load averages</b>     | Load averages for the last 1 minute, 5 minutes, and 15 minutes.                                                                   |

## Sample Output

### show system uptime

```
user@host> show system uptime
Current time:      1998-10-13 19:45:47 UTC
System booted:     1998-10-12 20:51:41 UTC (22:54:06 ago)
Protocols started: 1998-10-13 19:33:45 UTC (00:12:02 ago)
Last configured:   1998-10-13 19:33:45 UTC (00:12:02 ago) by abc
12:45PM up 22:54, 2 users, load averages: 0.07, 0.02, 0.01
```

### show system uptime all-lcc (TX Matrix Router)

```
user@host> show system uptime all-lcc
lcc0-re0:
-----
Current time: 2004-09-13 09:55:35 PDT
System booted: 2004-09-13 03:13:55 PDT (06:41:40 ago)
Last configured: 2004-09-13 03:17:48 PDT (06:37:47 ago) by root
9:55AM PDT up 6:42, 1 user, load averages: 0.02, 0.03, 0.00
lcc2-re0:
-----
Current time: 2004-09-13 09:55:35 PDT
System booted: 2004-09-12 03:23:43 PDT (1d 06:31 ago)
Last configured: 2004-09-13 03:05:36 PDT (06:49:59 ago) by root
9:55AM PDT up 1 day, 6:32, 1 user, load averages: 0.02, 0.01, 0.00
```

### show system uptime all-lcc (TX Matrix Plus Router)

```
user@host> show system uptime all-lcc
sfc0-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:33 PDT (17:44:57 ago)
Protocols started: 2009-05-24 06:40:30 PDT (17:44:00 ago)
Last configured: 2009-05-24 06:33:27 PDT (17:51:03 ago) by gregdo
12:24AM up 17:45, 2 users, load averages: 0.07, 0.05, 0.01

lcc0-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:46 PDT (17:44:44 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:47 PDT (17:43:43 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

lcc1-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:38 PDT (17:44:52 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:18 PDT (17:44:12 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

lcc2-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:48 PDT (17:44:42 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:44 PDT (17:43:46 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00
```

lcc3-re0:

-----  
Current time: 2009-05-25 00:24:30 PDT  
System booted: 2009-05-24 06:39:44 PDT (17:44:46 ago)  
error: the routing subsystem is not running  
Last configured: 2009-05-24 06:40:08 PDT (17:44:22 ago) by root  
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

#### show system uptime (EX Series)

```
user@switch> show system uptime
Current time: 2014-03-12 16:39:56 UTC
System booted: 2014-03-12 14:58:05 UTC (01:41:51 ago)
Protocols started: 2014-03-12 14:59:48 UTC (01:40:08 ago)
Last configured: 2014-03-12 14:58:58 UTC (01:40:58 ago) by root
4:39PM up 1:42, 4 users, load averages: 0.02, 0.02, 0.00
```

#### show system uptime (QFX Series)

```
user@switch> show system uptime
Current time: 2010-08-27 03:12:30 PDT
System booted: 2010-08-13 17:11:54 PDT (1w6d 10:00 ago)
Protocols started: 2010-08-13 17:13:56 PDT (1w6d 09:58 ago)
Last configured: 2010-08-26 05:54:00 PDT (21:18:30 ago) by regress
3:12AM up 13 days, 10:01, 3 users, load averages: 0.00, 0.00, 0.00
```

## show virtual-chassis active-topology

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis active-topology<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display the active topology of the Virtual Chassis or VCF with next-hop reachability information.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>none</b>—Display the active topology of the member switch where the command is issued.</p> <p><b>all-members</b>—(Optional) Display the active topology of all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the active topology of the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the active topology of the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157</a></li> <li>• <a href="#">Understanding EX Series Virtual Chassis Configuration on page 5086</a></li> </ul>                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis active-topology (EX4200 Virtual Chassis) on page 5175</a><br><a href="#">show virtual-chassis active-topology (EX8200 Virtual Chassis) on page 5175</a><br><a href="#">show virtual-chassis active-topology (Virtual Chassis Fabric) on page 5176</a>                                                                                                                                                                                                    |
| <b>Output Fields</b>            | Table 572 on page 5174 lists the output fields for the <b>show virtual-chassis active-topology</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                       |

**Table 572: show virtual-chassis active-topology Output Fields**

| Field Name            | Field Description                                                                                                                                                                              |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Destination ID</b> | Specifies the member ID of the destination.                                                                                                                                                    |
| <b>Next-hop</b>       | <p>Specifies the member ID and Virtual Chassis port (VCP) of the next hop to which packets for the destination ID are forwarded.</p> <p>The next hop can be more than one device in a VCF.</p> |

## Sample Output

### show virtual-chassis active-topology (EX4200 Virtual Chassis)

```

user@switch> show virtual-chassis active-topology
 1                      1(vcp-1)

 2                      1(vcp-1)

 3                      1(vcp-1)

 4                      1(vcp-1)

 5                      8(vcp-0) 1(vcp-1)

 6                      8(vcp-0)

 7                      8(vcp-0)

 8                      8(vcp-0)

```

### show virtual-chassis active-topology (EX8200 Virtual Chassis)

```

user@external-routing-engine> show virtual-chassis active-topology
member0:

```

| Destination ID | Next-hop           |
|----------------|--------------------|
| 1              | 1(vcp-4/0/4.32768) |
| 8              | 8(vcp-0/0.32768)   |
| 9              | 8(vcp-0/0.32768)   |

```
member1:
```

| Destination ID | Next-hop           |
|----------------|--------------------|
| 0              | 0(vcp-3/0/4.32768) |
| 8              | 8(vcp-0/0.32768)   |
| 9              | 8(vcp-0/0.32768)   |

```
member8:
```

| Destination ID | Next-hop         |
|----------------|------------------|
| 0              | 0(vcp-1/1.32768) |
| 1              | 1(vcp-1/2.32768) |
| 9              | 9(vcp-2/1.32768) |

member9:

| Destination ID | Next-hop         |
|----------------|------------------|
| 0              | 8(vcp-1/2.32768) |
| 1              | 8(vcp-1/2.32768) |
| 8              | 8(vcp-1/2.32768) |

### show virtual-chassis active-topology (Virtual Chassis Fabric)

user@device> show virtual-chassis active-topology  
fpc0:

| Destination ID            | Next-hop                                  |
|---------------------------|-------------------------------------------|
| 1<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 2<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 3<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 4                         | 4(vcp-255/0/2.32768)                      |
| 5                         | 5(vcp-255/0/3.32768)                      |
| 6                         | 6(vcp-255/0/1.32768)                      |

fpc1:

| Destination ID            | Next-hop                                  |
|---------------------------|-------------------------------------------|
| 0<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 2<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 3<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 4                         | 4(vcp-255/0/2.32768)                      |
| 5                         | 5(vcp-255/0/3.32768)                      |
| 6                         | 6(vcp-255/0/1.32768)                      |

fpc2:



| Destination ID            | Next-hop                                  |
|---------------------------|-------------------------------------------|
| 0<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 1<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 3<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 4                         | 4(vcp-255/0/2.32768)                      |
| 5                         | 5(vcp-255/0/3.32768)                      |
| 6                         | 6(vcp-255/0/1.32768)                      |

fpc3:

| Destination ID            | Next-hop                                  |
|---------------------------|-------------------------------------------|
| 0<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 1<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 2<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 4                         | 4(vcp-255/0/2.32768)                      |
| 5                         | 5(vcp-255/0/3.32768)                      |
| 6                         | 6(vcp-255/0/1.32768)                      |

fpc4:

| Destination ID             | Next-hop                                                             |
|----------------------------|----------------------------------------------------------------------|
| 0                          | 0(vcp-255/0/48.32768)                                                |
| 1                          | 1(vcp-255/0/49.32768)                                                |
| 2                          | 2(vcp-255/0/50.32768)                                                |
| 3                          | 3(vcp-255/0/51.32768)                                                |
| 5<br>0(vcp-255/0/48.32768) | 3(vcp-255/0/51.32768) 2(vcp-255/0/50.32768)<br>1(vcp-255/0/49.32768) |
| 6<br>0(vcp-255/0/48.32768) | 3(vcp-255/0/51.32768) 2(vcp-255/0/50.32768)<br>1(vcp-255/0/49.32768) |

fpc5:

| Destination ID | Next-hop              |
|----------------|-----------------------|
| 0              | 0(vcp-255/0/48.32768) |

|                       |                       |                       |
|-----------------------|-----------------------|-----------------------|
| 1                     | 1(vcp-255/0/49.32768) |                       |
| 2                     | 2(vcp-255/0/50.32768) |                       |
| 3                     | 3(vcp-255/0/51.32768) |                       |
| 4                     | 3(vcp-255/0/51.32768) | 2(vcp-255/0/50.32768) |
| 0(vcp-255/0/48.32768) | 1(vcp-255/0/49.32768) |                       |
| 6                     | 3(vcp-255/0/51.32768) | 2(vcp-255/0/50.32768) |
| 0(vcp-255/0/48.32768) | 1(vcp-255/0/49.32768) |                       |

fpc6:

| Destination ID       | Next-hop                                  |
|----------------------|-------------------------------------------|
| 0                    | 0(vcp-255/0/0.32768)                      |
| 1                    | 1(vcp-255/0/1.32768)                      |
| 2                    | 2(vcp-255/0/2.32768)                      |
| 3                    | 3(vcp-255/0/3.32768)                      |
| 4                    | 3(vcp-255/0/3.32768) 2(vcp-255/0/2.32768) |
| 0(vcp-255/0/0.32768) | 1(vcp-255/0/1.32768)                      |
| 5                    | 3(vcp-255/0/3.32768) 2(vcp-255/0/2.32768) |
| 0(vcp-255/0/0.32768) | 1(vcp-255/0/1.32768)                      |

## show virtual-chassis device-topology

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis device-topology<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Display the device topology—the member and system IDs, the VCP numbers, and device status—for all hardware devices in the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>none</b>—Display the device topology for all members of the Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Display the device topology for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the device topology for the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the device topology for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding EX Series Virtual Chassis Port Link Aggregation on page 5083</a></li> <li>• <a href="#">Understanding EX8200 Virtual Chassis Topologies</a></li> </ul>                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>            | <a href="#">Table 573 on page 5179</a> lists the output fields for the <b>show virtual-chassis device-topology</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                       |

**Table 573: show virtual-chassis device-topology Output Fields**

| Field Name    | Field Description                                                                                                                                                                                                                                                                                             |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Member</b> | Assigned member ID.                                                                                                                                                                                                                                                                                           |
| <b>Device</b> | Assigned device ID.<br><br>For an EX8200 Virtual Chassis, the member ID and the device ID are always identical.                                                                                                                                                                                               |
| <b>Status</b> | The status of the device within the Virtual Chassis or VCF. Outputs include: <ul style="list-style-type: none"> <li>• <b>Prsnt</b>—Device is currently connected to and participating in the Virtual Chassis or VCF.</li> <li>• <b>NotPrsnt</b>—Device is assigned but is not currently connected.</li> </ul> |

Table 573: show virtual-chassis device-topology Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                      |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <b>System ID</b>                 | System ID of the device.<br><br>The system ID of the device is the device's MAC address.                                               |
| <b>Member (Neighbor List)</b>    | Assigned member ID of the neighbor device.                                                                                             |
| <b>Device (Neighbor List)</b>    | Assigned device ID of the neighbor device.<br><br>For an EX8200 Virtual Chassis, the member ID and the device ID are always identical. |
| <b>Interface (Neighbor List)</b> | The interface connecting the device to the neighbor.                                                                                   |

## Sample Output

### show virtual-chassis device-topology

```
user@switch> show virtual-chassis device-topology
```

```
member0:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
0         0      Prsnt   0021.59f7.d000  8         8      vcp-0/0
                                1         1      vcp-4/0/1
1         1      Prsnt   0026.888d.6800  8         8      vcp-0/0
                                9         9      vcp-0/1
                                0         0      vcp-3/0/4
8         8      Prsnt   0000.4a75.9b7c  9         9      vcp-1/0
                                0         0      vcp-1/1
                                1         1      vcp-1/2
9         9      Prsnt   0000.73e9.9a57  8         8      vcp-1/0
                                1         1      vcp-1/1
```

```
member1:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
0         0      Prsnt   0021.59f7.d000  8         8      vcp-0/0
                                1         1      vcp-4/0/1
1         1      Prsnt   0026.888d.6800  8         8      vcp-0/0
                                9         9      vcp-0/1
                                0         0      vcp-3/0/4
8         8      Prsnt   0000.4a75.9b7c  9         9      vcp-1/0
                                0         0      vcp-1/1
                                1         1      vcp-1/2
9         9      Prsnt   0000.73e9.9a57  8         8      vcp-1/0
                                1         1      vcp-1/1
```

```
member8:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
```

|   |   |       |                |   |   |           |
|---|---|-------|----------------|---|---|-----------|
| 0 | 0 | Prsnt | 0021.59f7.d000 | 8 | 8 | vcp-0/0   |
|   |   |       |                | 1 | 1 | vcp-4/0/1 |
| 1 | 1 | Prsnt | 0026.888d.6800 | 8 | 8 | vcp-0/0   |
|   |   |       |                | 9 | 9 | vcp-0/1   |
|   |   |       |                | 0 | 0 | vcp-3/0/4 |
| 8 | 8 | Prsnt | 0000.4a75.9b7c | 9 | 9 | vcp-1/0   |
|   |   |       |                | 0 | 0 | vcp-1/1   |
|   |   |       |                | 1 | 1 | vcp-1/2   |
| 9 | 9 | Prsnt | 0000.73e9.9a57 | 8 | 8 | vcp-1/0   |
|   |   |       |                | 1 | 1 | vcp-1/1   |

member9:

|        |        |        |                | Neighbor List |        |           |
|--------|--------|--------|----------------|---------------|--------|-----------|
| Member | Device | Status | System ID      | Member        | Device | Interface |
| 0      | 0      | Prsnt  | 0021.59f7.d000 | 8             | 8      | vcp-0/0   |
|        |        |        |                | 1             | 1      | vcp-4/0/1 |
| 1      | 1      | Prsnt  | 0026.888d.6800 | 8             | 8      | vcp-0/0   |
|        |        |        |                | 9             | 9      | vcp-0/1   |
|        |        |        |                | 0             | 0      | vcp-3/0/4 |
| 8      | 8      | Prsnt  | 0000.4a75.9b7c | 9             | 9      | vcp-1/0   |
|        |        |        |                | 0             | 0      | vcp-1/1   |
|        |        |        |                | 1             | 1      | vcp-1/2   |
| 9      | 9      | Prsnt  | 0000.73e9.9a57 | 8             | 8      | vcp-1/0   |
|        |        |        |                | 1             | 1      | vcp-1/1   |

#### show virtual-chassis device-topology (Virtual Chassis Fabric)

user@device> show virtual-chassis device-topology  
fpc0:

|        |        |        |                | Neighbor List |        |              |
|--------|--------|--------|----------------|---------------|--------|--------------|
| Member | Device | Status | System ID      | Member        | Device | Interface    |
| 0      | 0      | Prsnt  | 100e.7eb6.a900 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 1      | 1      | Prsnt  | 100e.7eb8.3a40 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 2      | 2      | Prsnt  | 100e.7eb5.d700 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 3      | 3      | Prsnt  | 100e.7eb5.c440 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 4      | 4      | Prsnt  | 100e.7eb5.7e40 | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
| 5      | 5      | Prsnt  | 100e.7eb5.80c0 | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
|        |        |        |                | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
| 6      | 6      | Prsnt  | 100e.7eb6.3b00 | 3             | 3      | vcp-255/0/3  |
|        |        |        |                | 2             | 2      | vcp-255/0/2  |
|        |        |        |                | 0             | 0      | vcp-255/0/0  |
|        |        |        |                | 1             | 1      | vcp-255/0/1  |

fpc1:

Neighbor List

| Member | Device | Status | System ID      | Member | Device | Interface    |
|--------|--------|--------|----------------|--------|--------|--------------|
| 0      | 0      | Prsnt  | 100e.7eb6.a900 | 4      | 4      | vcp-255/0/2  |
|        |        |        |                | 5      | 5      | vcp-255/0/3  |
|        |        |        |                | 6      | 6      | vcp-255/0/1  |
| 1      | 1      | Prsnt  | 100e.7eb8.3a40 | 4      | 4      | vcp-255/0/2  |
|        |        |        |                | 5      | 5      | vcp-255/0/3  |
|        |        |        |                | 6      | 6      | vcp-255/0/1  |
| 2      | 2      | Prsnt  | 100e.7eb5.d700 | 4      | 4      | vcp-255/0/2  |
|        |        |        |                | 5      | 5      | vcp-255/0/3  |
|        |        |        |                | 6      | 6      | vcp-255/0/1  |
| 3      | 3      | Prsnt  | 100e.7eb5.c440 | 4      | 4      | vcp-255/0/2  |
|        |        |        |                | 5      | 5      | vcp-255/0/3  |
|        |        |        |                | 6      | 6      | vcp-255/0/1  |
| 4      | 4      | Prsnt  | 100e.7eb5.7e40 | 3      | 3      | vcp-255/0/51 |
|        |        |        |                | 2      | 2      | vcp-255/0/50 |
|        |        |        |                | 0      | 0      | vcp-255/0/48 |
|        |        |        |                | 1      | 1      | vcp-255/0/49 |
| 5      | 5      | Prsnt  | 100e.7eb5.80c0 | 3      | 3      | vcp-255/0/51 |
|        |        |        |                | 2      | 2      | vcp-255/0/50 |
|        |        |        |                | 1      | 1      | vcp-255/0/49 |
|        |        |        |                | 0      | 0      | vcp-255/0/48 |
| 6      | 6      | Prsnt  | 100e.7eb6.3b00 | 3      | 3      | vcp-255/0/3  |
|        |        |        |                | 2      | 2      | vcp-255/0/2  |
|        |        |        |                | 0      | 0      | vcp-255/0/0  |
|        |        |        |                | 1      | 1      | vcp-255/0/1  |

fpc2:

| Neighbor List |        |        |                |        |        |              |
|---------------|--------|--------|----------------|--------|--------|--------------|
| Member        | Device | Status | System ID      | Member | Device | Interface    |
| 0             | 0      | Prsnt  | 100e.7eb6.a900 | 4      | 4      | vcp-255/0/2  |
|               |        |        |                | 5      | 5      | vcp-255/0/3  |
|               |        |        |                | 6      | 6      | vcp-255/0/1  |
| 1             | 1      | Prsnt  | 100e.7eb8.3a40 | 4      | 4      | vcp-255/0/2  |
|               |        |        |                | 5      | 5      | vcp-255/0/3  |
|               |        |        |                | 6      | 6      | vcp-255/0/1  |
| 2             | 2      | Prsnt  | 100e.7eb5.d700 | 4      | 4      | vcp-255/0/2  |
|               |        |        |                | 5      | 5      | vcp-255/0/3  |
|               |        |        |                | 6      | 6      | vcp-255/0/1  |
| 3             | 3      | Prsnt  | 100e.7eb5.c440 | 4      | 4      | vcp-255/0/2  |
|               |        |        |                | 5      | 5      | vcp-255/0/3  |
|               |        |        |                | 6      | 6      | vcp-255/0/1  |
| 4             | 4      | Prsnt  | 100e.7eb5.7e40 | 3      | 3      | vcp-255/0/51 |
|               |        |        |                | 2      | 2      | vcp-255/0/50 |
|               |        |        |                | 0      | 0      | vcp-255/0/48 |
|               |        |        |                | 1      | 1      | vcp-255/0/49 |
| 5             | 5      | Prsnt  | 100e.7eb5.80c0 | 3      | 3      | vcp-255/0/51 |
|               |        |        |                | 2      | 2      | vcp-255/0/50 |
|               |        |        |                | 1      | 1      | vcp-255/0/49 |
|               |        |        |                | 0      | 0      | vcp-255/0/48 |
| 6             | 6      | Prsnt  | 100e.7eb6.3b00 | 3      | 3      | vcp-255/0/3  |
|               |        |        |                | 2      | 2      | vcp-255/0/2  |
|               |        |        |                | 0      | 0      | vcp-255/0/0  |
|               |        |        |                | 1      | 1      | vcp-255/0/1  |

fpc3:

| Neighbor List |        |        |                |        |        |             |
|---------------|--------|--------|----------------|--------|--------|-------------|
| Member        | Device | Status | System ID      | Member | Device | Interface   |
| 0             | 0      | Prsnt  | 100e.7eb6.a900 | 4      | 4      | vcp-255/0/2 |
|               |        |        |                | 5      | 5      | vcp-255/0/3 |

|   |   |       |                |   |   |              |
|---|---|-------|----------------|---|---|--------------|
| 1 | 1 | Prsnt | 100e.7eb8.3a40 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 4 | 4 | vcp-255/0/2  |
|   |   |       |                | 5 | 5 | vcp-255/0/3  |
| 2 | 2 | Prsnt | 100e.7eb5.d700 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 4 | 4 | vcp-255/0/2  |
|   |   |       |                | 5 | 5 | vcp-255/0/3  |
| 3 | 3 | Prsnt | 100e.7eb5.c440 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 4 | 4 | vcp-255/0/2  |
|   |   |       |                | 5 | 5 | vcp-255/0/3  |
| 4 | 4 | Prsnt | 100e.7eb5.7e40 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 3 | 3 | vcp-255/0/51 |
|   |   |       |                | 2 | 2 | vcp-255/0/50 |
| 5 | 5 | Prsnt | 100e.7eb5.80c0 | 0 | 0 | vcp-255/0/48 |
|   |   |       |                | 1 | 1 | vcp-255/0/49 |
|   |   |       |                | 3 | 3 | vcp-255/0/51 |
|   |   |       |                | 2 | 2 | vcp-255/0/50 |
|   |   |       |                | 1 | 1 | vcp-255/0/49 |
| 6 | 6 | Prsnt | 100e.7eb6.3b00 | 0 | 0 | vcp-255/0/48 |
|   |   |       |                | 3 | 3 | vcp-255/0/3  |
|   |   |       |                | 2 | 2 | vcp-255/0/2  |
|   |   |       |                | 0 | 0 | vcp-255/0/0  |
|   |   |       |                | 1 | 1 | vcp-255/0/1  |

fpc4:

|        |        |        |                | Neighbor List |        |              |
|--------|--------|--------|----------------|---------------|--------|--------------|
| Member | Device | Status | System ID      | Member        | Device | Interface    |
| 0      | 0      | Prsnt  | 100e.7eb6.a900 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 1      | 1      | Prsnt  | 100e.7eb8.3a40 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 2      | 2      | Prsnt  | 100e.7eb5.d700 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 3      | 3      | Prsnt  | 100e.7eb5.c440 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 4      | 4      | Prsnt  | 100e.7eb5.7e40 | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
| 5      | 5      | Prsnt  | 100e.7eb5.80c0 | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
|        |        |        |                | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
| 6      | 6      | Prsnt  | 100e.7eb6.3b00 | 3             | 3      | vcp-255/0/3  |
|        |        |        |                | 2             | 2      | vcp-255/0/2  |
|        |        |        |                | 0             | 0      | vcp-255/0/0  |
|        |        |        |                | 1             | 1      | vcp-255/0/1  |

fpc5:

|        |        |        |                | Neighbor List |        |             |
|--------|--------|--------|----------------|---------------|--------|-------------|
| Member | Device | Status | System ID      | Member        | Device | Interface   |
| 0      | 0      | Prsnt  | 100e.7eb6.a900 | 4             | 4      | vcp-255/0/2 |
|        |        |        |                | 5             | 5      | vcp-255/0/3 |
|        |        |        |                | 6             | 6      | vcp-255/0/1 |
| 1      | 1      | Prsnt  | 100e.7eb8.3a40 | 4             | 4      | vcp-255/0/2 |
|        |        |        |                | 5             | 5      | vcp-255/0/3 |

|   |   |       |                |   |   |              |
|---|---|-------|----------------|---|---|--------------|
| 2 | 2 | Prsnt | 100e.7eb5.d700 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 4 | 4 | vcp-255/0/2  |
|   |   |       |                | 5 | 5 | vcp-255/0/3  |
| 3 | 3 | Prsnt | 100e.7eb5.c440 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 4 | 4 | vcp-255/0/2  |
|   |   |       |                | 5 | 5 | vcp-255/0/3  |
| 4 | 4 | Prsnt | 100e.7eb5.7e40 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 3 | 3 | vcp-255/0/51 |
|   |   |       |                | 2 | 2 | vcp-255/0/50 |
| 5 | 5 | Prsnt | 100e.7eb5.80c0 | 0 | 0 | vcp-255/0/48 |
|   |   |       |                | 1 | 1 | vcp-255/0/49 |
|   |   |       |                | 3 | 3 | vcp-255/0/51 |
| 6 | 6 | Prsnt | 100e.7eb6.3b00 | 2 | 2 | vcp-255/0/50 |
|   |   |       |                | 1 | 1 | vcp-255/0/49 |
|   |   |       |                | 0 | 0 | vcp-255/0/48 |
|   |   |       |                | 3 | 3 | vcp-255/0/3  |
|   |   |       |                | 2 | 2 | vcp-255/0/2  |
|   |   |       |                | 0 | 0 | vcp-255/0/0  |
|   |   |       |                | 1 | 1 | vcp-255/0/1  |

fpc6:

|        |        |        |                | Neighbor List |        |              |
|--------|--------|--------|----------------|---------------|--------|--------------|
| Member | Device | Status | System ID      | Member        | Device | Interface    |
| 0      | 0      | Prsnt  | 100e.7eb6.a900 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 1      | 1      | Prsnt  | 100e.7eb8.3a40 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 2      | 2      | Prsnt  | 100e.7eb5.d700 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 3      | 3      | Prsnt  | 100e.7eb5.c440 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 4      | 4      | Prsnt  | 100e.7eb5.7e40 | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
| 5      | 5      | Prsnt  | 100e.7eb5.80c0 | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
| 6      | 6      | Prsnt  | 100e.7eb6.3b00 | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
|        |        |        |                | 3             | 3      | vcp-255/0/3  |
|        |        |        |                | 2             | 2      | vcp-255/0/2  |
|        |        |        |                | 0             | 0      | vcp-255/0/0  |
|        |        |        |                | 1             | 1      | vcp-255/0/1  |



## show virtual-chassis login

|                                 |                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show virtual-chassis login</b>                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                |
| <b>Description</b>              | <p>Supply the address of the host that logged into the Virtual Chassis or VCF, or identify the location of the member switch that redirected the current session to a different member switch.</p> <p>You might need this information for tracing or troubleshooting purposes.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">request session member on page 5163</a></li> <li>• <a href="#">Understanding Global Management of a Virtual Chassis on page 5080</a></li> </ul>                                                                               |
| <b>List of Sample Output</b>    | <p><a href="#">show virtual-chassis login (Direct Login to the Master Console Port) on page 5185</a></p> <p><a href="#">show virtual-chassis login (Backup Console Session Redirected to the Master Console Port) on page 5185</a></p>                                             |

### Sample Output

#### show virtual-chassis login (Direct Login to the Master Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host 248.1.2.3
```

#### show virtual-chassis login (Backup Console Session Redirected to the Master Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host backup
```

## show virtual-chassis protocol adjacency

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show virtual-chassis protocol adjacency</code><br><code>&lt;brief   detail   extensive&gt;</code><br><code>&lt;all-members&gt;</code><br><code>&lt;local&gt;</code><br><code>&lt;member <i>member-id</i>&gt;</code><br><code>&lt;system-id&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display the Virtual Chassis Control Protocol (VCCP) adjacency statistics in the Virtual Chassis or VCF for all hardware devices.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>none</b>—Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> and <b>extensive</b> options provide identical displays.</p> <p><b>all-members</b>—(Optional) Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display VCCP adjacency statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.</p> <p><b>system-id</b>—(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding EX Series Virtual Chassis Port Link Aggregation on page 5083</a></li><li>• <a href="#">Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis protocol adjacency on page 5187</a><br><a href="#">show virtual-chassis protocol adjacency detail on page 5188</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>            | <a href="#">Table 574 on page 5187</a> lists the output fields for the <b>show virtual-chassis protocol adjacency</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 574: show virtual-chassis protocol adjacency Output Fields

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                        | Level of Output |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>           | Name of the Virtual Chassis port (VCP) interface.                                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>System</b>              | The MAC address of the device on the receiving side of the VCP link.                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>State</b>               | State of the link. Outputs include: <ul style="list-style-type: none"> <li>• <b>Up</b>—The link is up.</li> <li>• <b>Down</b>—The link is down.</li> <li>• <b>New</b>—The link is new.</li> <li>• <b>One-way</b>—The link is transmitting traffic in one direction.</li> <li>• <b>Initializing</b>—The link is initializing.</li> <li>• <b>Rejected</b>—The link is rejected.</li> </ul> | All levels      |
| <b>Hold, Expires in</b>    | Remaining holdtime of the adjacency.                                                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>Priority</b>            | Priority to become the designated intermediary system.                                                                                                                                                                                                                                                                                                                                   | detail          |
| <b>Up/Down Transitions</b> | Count of adjacency status transition changes from up to down or down to up.                                                                                                                                                                                                                                                                                                              | detail          |
| <b>Last transition</b>     | Time of the last up/down transition.                                                                                                                                                                                                                                                                                                                                                     | detail          |

## Sample Output

### show virtual-chassis protocol adjacency

```
user@switch> show virtual-chassis protocol adjacency
```

```
member0:
```

```
-----
Interface      System      State      Hold (secs)
vcp-0/0.32768  0000.4a75.9b7c Up          57
vcp-0/1.32768  0000.4a75.9b7c Up          59
vcp-4/0/1.32768 0026.888d.6800 Up          57
```

```
member1:
```

```
-----
Interface      System      State      Hold (secs)
vcp-0/0.32768  0000.4a75.9b7c Up          58
vcp-0/1.32768  0000.73e9.9a57 Up          59
vcp-3/0/4.32768 0021.59f7.d000 Up          58
```

```
member8:
```

```
-----
Interface      System      State      Hold (secs)
vcp-1/0.32768  0000.73e9.9a57 Up          58
vcp-1/1.32768  0021.59f7.d000 Up          58
vcp-1/2.32768  0026.888d.6800 Up          59
vcp-2/0.32768  0021.59f7.d000 Up          59
```

```
member9:
```

```
-----
Interface      System      State      Hold (secs)
```

|               |                   |    |
|---------------|-------------------|----|
| vcp-1/0.32768 | 0000.4a75.9b7c Up | 58 |
| vcp-1/1.32768 | 0026.888d.6800 Up | 59 |

### show virtual-chassis protocol adjacency detail

```
user@switch> show virtual-chassis protocol adjacency detail
```

```
member0:
```

-----

```
0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago
```

```
0000.4a75.9b7c
  interface-name: vcp-0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago
```

```
0026.888d.6800
  interface-name: vcp-4/0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago
```

```
member1:
```

-----

```
0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0000.73e9.9a57
  interface-name: vcp-0/1.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago
```

```
0021.59f7.d000
  interface-name: vcp-3/0/4.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago
```

```
member8:
```

-----

```
0000.73e9.9a57
  interface-name: vcp-1/0.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0021.59f7.d000
  interface-name: vcp-1/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0026.888d.6800
  interface-name: vcp-1/2.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0021.59f7.d000
  interface-name: vcp-2/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
member9:
```

-----

```
0000.4a75.9b7c
  interface-name: vcp-1/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0026.888d.6800
  interface-name: vcp-1/1.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago
```

## show virtual-chassis protocol database

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis protocol database<br><brief   detail   extensive><br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display the Virtual Chassis Control Protocol (VCCP) database statistics for all hardware devices within the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>none</b>—Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> option provides more output than the <b>brief</b> option. The <b>extensive</b> option provides all output and is most useful for customer support personnel.</p> <p><b>all-members</b>—(Optional) Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display VCCP database statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display VCCP database statistics for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</a></li> <li>• <a href="#">Understanding EX Series Virtual Chassis Components on page 5067</a></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis Components</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis protocol database on page 5191</a><br><a href="#">show virtual-chassis protocol database detail on page 5192</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | <a href="#">Table 575 on page 5190</a> lists the output fields for the <b>show virtual-chassis protocol database</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Table 575: show virtual-chassis protocol database Output Fields**

| Field Name | Field Description                               | Level of Output |
|------------|-------------------------------------------------|-----------------|
| LSP ID     | Link-state protocol (LSP) data unit identifier. | All levels      |

Table 575: show virtual-chassis protocol database Output Fields (*continued*)

| Field Name       | Field Description                                      | Level of Output |
|------------------|--------------------------------------------------------|-----------------|
| <b>Sequence</b>  | Sequence number of the LSP.                            | All levels      |
| <b>Checksum</b>  | Checksum value of the LSP.                             | All levels      |
| <b>Lifetime</b>  | Remaining lifetime of the LSP, in seconds.             | All levels      |
| <b>Neighbor</b>  | MAC address of the neighbor on the advertising system. | detail          |
| <b>Interface</b> | Virtual Chassis port (VCP) interface name.             | detail          |
| <b>Metric</b>    | Metric of the prefix or neighbor.                      | detail          |

The **extensive** output was omitted from this list. The **extensive** output is useful for customer support personnel only.

## Sample Output

### show virtual-chassis protocol database

```
user@switch> show virtual-chassis protocol database
```

```
member0:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361   0x27e8   113
0021.59f7.d000.00-00  0x16882  0x3993   118
0026.888d.6800.00-00  0x1691f  0x82b7   116
  4 LSPs
```

```
member1:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361   0x27e8   114
0021.59f7.d000.00-00  0x16883  0x289    116
0026.888d.6800.00-00  0x1691f  0x82b7   118
  4 LSPs
```

```
member8:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   118
0000.73e9.9a57.00-00  0xf361   0x27e8   114
0021.59f7.d000.00-00  0x16883  0x289    116
0026.888d.6800.00-00  0x16920  0xa335   116
  4 LSPs
```

```
member9:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361   0x27e8   116
0021.59f7.d000.00-00  0x16883  0x289    114
```

```
0026.888d.6800.00-00      0x16920   0xa335      116
4 LSPs
```

### show virtual-chassis protocol database detail

```
user@switch> show virtual-chassis protocol database detail
member0:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 114 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 118 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15
```

```
0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15
```

```
member1:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15
```

```
0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 117 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15
```

```
member8:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 118 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
```



```

Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15

0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15

```

member9:

```

-----

0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 116 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150

0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 117 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 113 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15

0026.888d.6800.00-00 Sequence: 0x1694f, Checksum: 0xa61a, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15

```

## show virtual-chassis protocol interface

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show virtual-chassis protocol interface</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;all-members&gt;</code><br><code>&lt;interface-name&gt;</code><br><code>&lt;local&gt;</code><br><code>&lt;member member-id&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Display information about Virtual Chassis Control Protocol (VCCP) statistics for VCCP-enabled interfaces within the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>none</b>—Display the VCCP interface statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail</b> —(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> option provides more output than the <b>brief</b> option.</p> <p><b>all-members</b>—(Optional) Display VCCP interface statistics for all members of the Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display VCCP interface statistics for the specified interface.</p> <p><b>local</b>—(Optional) Display VCCP interface statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display VCCP interface statistics for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>EX Series Virtual Chassis Overview</i></li><li>• <i>Understanding QFX Series Virtual Chassis</i></li><li>• <i>Understanding Virtual Chassis Ports in an EX8200 Virtual Chassis</i></li><li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis protocol interface on page 5195</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | <a href="#">Table 576 on page 5195</a> lists the output fields for the <b>show virtual-chassis protocol interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 576: show virtual-chassis protocol interface Output Fields

| Field Name       | Field Description                                                                                                                                          | Level of Output |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b> | Name of the VCP.                                                                                                                                           | All levels      |
| <b>State</b>     | State of the link. Outputs include: <ul style="list-style-type: none"> <li>• <b>Up</b>—The link is up.</li> <li>• <b>Down</b>—The link is down.</li> </ul> | All levels      |
| <b>Metric</b>    | Metric of the prefix or neighbor.                                                                                                                          | All levels      |

## Sample Output

### show virtual-chassis protocol interface

```
user@switch> show virtual-chassis protocol interface
```

```
member0:
```

```
-----
IS-IS interface database:
```

| Interface       | State | Metric |
|-----------------|-------|--------|
| vcp-0/0.32768   | Up    | 150    |
| vcp-0/1.32768   | Up    | 150    |
| vcp-4/0/1.32768 | Up    | 15     |
| vcp-4/0/7.32768 | Down  | 15     |

```
member1:
```

```
-----
IS-IS interface database:
```

| Interface       | State | Metric |
|-----------------|-------|--------|
| vcp-0/0.32768   | Up    | 150    |
| vcp-0/1.32768   | Up    | 150    |
| vcp-3/0/4.32768 | Up    | 15     |

```
member8:
```

```
-----
IS-IS interface database:
```

| Interface     | State | Metric |
|---------------|-------|--------|
| vcp-0/0.32768 | Down  | 150    |
| vcp-1/0.32768 | Up    | 150    |
| vcp-1/1.32768 | Up    | 150    |
| vcp-1/2.32768 | Up    | 150    |
| vcp-1/3.32768 | Down  | 150    |
| vcp-2/0.32768 | Up    | 150    |
| vcp-2/1.32768 | Down  | 150    |
| vcp-2/2.32768 | Down  | 150    |
| vcp-2/3.32768 | Down  | 150    |

```
member9:
```

```
-----
IS-IS interface database:
```

| Interface     | State | Metric |
|---------------|-------|--------|
| vcp-0/0.32768 | Down  | 150    |
| vcp-1/0.32768 | Up    | 150    |
| vcp-1/1.32768 | Up    | 150    |
| vcp-1/2.32768 | Down  | 150    |
| vcp-1/3.32768 | Down  | 150    |



## show virtual-chassis protocol route

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis protocol route<br><all-members><br><destination-id><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display the unicast and multicast Virtual Chassis Control Protocol (VCCP) routing tables within the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>none</b>—Display the unicast and multicast routing tables for all members of the Virtual Chassis.</p> <p><b>all-members</b>—(Optional) Display the unicast and multicast routing tables for all members of the Virtual Chassis or VCF.</p> <p><b>destination-id</b>—(Optional) Display the unicast and multicast routing tables to the specified destination member ID for each member of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the unicast and multicast routing tables on the device where this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the unicast and multicast routing tables for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>EX Series Virtual Chassis Overview</i></li> <li>• <i>Understanding QFX Series Virtual Chassis</i></li> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis protocol route on page 5198</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 577 on page 5197</a> lists the output fields for the <b>show virtual-chassis protocol route</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Table 577: show virtual-chassis protocol route Output Fields**

| Field Name     | Field Description                                                              |
|----------------|--------------------------------------------------------------------------------|
| <b>Dev</b>     | MAC address of the member storing the VCCP routing table.                      |
| <b>Version</b> | Version of the shortest-path-first algorithm that generated the routing table. |

Table 577: show virtual-chassis protocol route Output Fields (*continued*)

| Field Name       | Field Description                                                            |
|------------------|------------------------------------------------------------------------------|
| <b>System ID</b> | MAC address of the device.                                                   |
| <b>Version</b>   | Version of the shortest-path-first (SPF) algorithm that generated the route. |
| <b>Metric</b>    | The metric number to get to that device.                                     |
| <b>Interface</b> | Name of the Virtual Chassis port (VCP) interface connecting the devices.     |
| <b>Via</b>       | MAC address of the next-hop device, if applicable.                           |

## Sample Output

### show virtual-chassis protocol route

```

user@switch> show virtual-chassis protocol route
member0:
-----
Dev 0021.59f7.d000 ucast routing table           Current version: 21
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    21      150 vcp-0/1.32768 0000.4a75.9b7c
0000.73e9.9a57    21      165 vcp-4/0/1.32768 0026.888d.6800
0021.59f7.d000    21         0
0026.888d.6800    21      15 vcp-4/0/1.32768 0026.888d.6800

Dev 0021.59f7.d000 mcast routing table           Current version: 21
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    21
0000.73e9.9a57    21
0021.59f7.d000    21          vcp-4/0/1.32768
                   vcp-0/1.32768
0026.888d.6800    21

member1:
-----
Dev 0026.888d.6800 ucast routing table           Current version: 25
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    25      150 vcp-0/0.32768 0000.4a75.9b7c
0000.73e9.9a57    25      150 vcp-0/1.32768 0000.73e9.9a57
0021.59f7.d000    25         15 vcp-3/0/4.32768 0021.59f7.d000
0026.888d.6800    25         0

Dev 0026.888d.6800 mcast routing table           Current version: 25
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    25
0000.73e9.9a57    25          vcp-3/0/4.32768
0021.59f7.d000    25          vcp-0/1.32768
0026.888d.6800    25          vcp-3/0/4.32768
                   vcp-0/0.32768

```

vcp-0/1.32768

member8:

```
-----
Dev 0000.4a75.9b7c ucast routing table          Current version: 39
-----
```

| System ID      | Version | Metric | Interface     | Via            |
|----------------|---------|--------|---------------|----------------|
| 0000.4a75.9b7c | 39      | 0      |               |                |
| 0000.73e9.9a57 | 39      | 150    | vcp-1/0.32768 | 0000.73e9.9a57 |
| 0021.59f7.d000 | 39      | 150    | vcp-2/0.32768 | 0021.59f7.d000 |
| 0026.888d.6800 | 39      | 150    | vcp-1/2.32768 | 0026.888d.6800 |

```
Dev 0000.4a75.9b7c mcast routing table          Current version: 39
-----
```

| System ID      | Version | Metric | Interface     | Via |
|----------------|---------|--------|---------------|-----|
| 0000.4a75.9b7c | 39      |        | vcp-1/0.32768 |     |
|                |         |        | vcp-2/0.32768 |     |
|                |         |        | vcp-1/2.32768 |     |
| 0000.73e9.9a57 | 39      |        |               |     |
| 0021.59f7.d000 | 39      |        |               |     |
| 0026.888d.6800 | 39      |        |               |     |

member9:

```
-----
Dev 0000.73e9.9a57 ucast routing table          Current version: 31
-----
```

| System ID      | Version | Metric | Interface     | Via            |
|----------------|---------|--------|---------------|----------------|
| 0000.4a75.9b7c | 31      | 150    | vcp-1/0.32768 | 0000.4a75.9b7c |
| 0000.73e9.9a57 | 31      | 0      |               |                |
| 0021.59f7.d000 | 31      | 165    | vcp-1/1.32768 | 0026.888d.6800 |
| 0026.888d.6800 | 31      | 150    | vcp-1/1.32768 | 0026.888d.6800 |

```
Dev 0000.73e9.9a57 mcast routing table          Current version: 31
-----
```

| System ID      | Version | Metric | Interface     | Via |
|----------------|---------|--------|---------------|-----|
| 0000.4a75.9b7c | 31      |        |               |     |
| 0000.73e9.9a57 | 31      |        | vcp-1/0.32768 |     |
|                |         |        | vcp-1/1.32768 |     |
| 0021.59f7.d000 | 31      |        |               |     |
| 0026.888d.6800 | 31      |        |               |     |

## show virtual-chassis protocol statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis protocol statistics<br><all-members><br><interface-name><br><local><br><member member-id>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display the Virtual Chassis Control Protocol (VCCP) statistics for all hardware devices within the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>none</b>—Display VCCP statistics for all members of the Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Display VCCP statistics for all members of the Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display VCCP statistics for the specified interface.</p> <p><b>local</b>—(Optional) Display VCCP statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display VCCP statistics for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>EX Series Virtual Chassis Overview</i></li> <li>• <i>Understanding QFX Series Virtual Chassis</i></li> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> </ul>                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis protocol statistics on page 5201</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>            | <a href="#">Table 578 on page 5200</a> lists the output fields for the <b>show virtual-chassis protocol interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                          |

**Table 578: show virtual-chassis protocol statistics Output Fields**

| Field Name       | Field Description                                                                    |
|------------------|--------------------------------------------------------------------------------------|
| <b>PDU type</b>  | Protocol data unit type.                                                             |
| <b>Received</b>  | Number of PDUs received since VCCP started or since the statistics were set to zero. |
| <b>Processed</b> | Number of PDUs received minus the number of PDUs dropped.                            |



Table 578: show virtual-chassis protocol statistics Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                        |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Drops</b>                  | Number of PDUs dropped.                                                                                                                                  |
| <b>Sent</b>                   | Number of PDUs transmitted since VCCP started or since the statistics were set to zero.                                                                  |
| <b>Rexmit</b>                 | Number of PDUs retransmitted since VCCP started or since the statistics were set to zero.                                                                |
| <b>Total Packets Received</b> | Number of PDUs received since VCCP started or since the statistics were set to zero.                                                                     |
| <b>Total Packets Sent</b>     | Number of PDUs sent since VCCP started or since the statistics were set to zero.                                                                         |
| <b>LSP queue length</b>       | Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.                                                            |
| <b>SPF runs</b>               | Number of shortest-path-first (SPF) calculations that have been performed.                                                                               |
| <b>Fragments Rebuilt</b>      | Number of link-state PDU fragments that the local system has computed.                                                                                   |
| <b>LSP Regenerations</b>      | Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed. |
| <b>Purges initiated</b>       | Number of purges that the system initiated. A purge is initiated if the software determines that a link-state PDU must be removed from the network.      |

## Sample Output

### show virtual-chassis protocol statistics

```

user@switch> show virtual-chassis protocol statistics
member0:
-----
IS-IS statistics for 0021.59f7.d000:
PDU type      Received    Processed      Drops      Sent      Rexmit
LSP            8166        8166           0         4551         0
HELLO          1659        1659           0         1693         0
CSNP             2            2             0            3         0
PSNP           1909        1909           0         2293         0
Unknown         0            0             0            0         0
Totals        11736       11736           0         8540         0

Total packets received: 11736 Sent: 8540

LSP queue length: 0 Drops: 0
SPF runs: 9
Fragments rebuilt: 1640
LSP regenerations: 1
Purges initiated: 0

member1:
-----
IS-IS statistics for 0026.888d.6800:

```

| PDU type | Received | Processed | Drops | Sent  | Rexmit |
|----------|----------|-----------|-------|-------|--------|
| LSP      | 10909    | 10909     | 0     | 12088 | 0      |
| HELLO    | 1877     | 1877      | 0     | 2251  | 0      |
| CSNP     | 3        | 3         | 0     | 3     | 0      |
| PSNP     | 3846     | 3846      | 0     | 3732  | 0      |
| Unknown  | 0        | 0         | 0     | 0     | 0      |
| Totals   | 16635    | 16635     | 0     | 18074 | 0      |

Total packets received: 16635 Sent: 18074

LSP queue length: 0 Drops: 0  
SPF runs: 13  
Fragments rebuilt: 1871  
LSP regenerations: 2  
Purges initiated: 0

member8:

-----  
IS-IS statistics for 0000.4a75.9b7c:

| PDU type | Received | Processed | Drops | Sent  | Rexmit |
|----------|----------|-----------|-------|-------|--------|
| LSP      | 7935     | 7935      | 0     | 14865 | 0      |
| HELLO    | 2695     | 2695      | 0     | 7124  | 0      |
| CSNP     | 4        | 4         | 0     | 4     | 0      |
| PSNP     | 4398     | 4398      | 0     | 3666  | 0      |
| Unknown  | 0        | 0         | 0     | 0     | 0      |
| Totals   | 15032    | 15032     | 0     | 25659 | 0      |

Total packets received: 15032 Sent: 25659

LSP queue length: 0 Drops: 0  
SPF runs: 26  
Fragments rebuilt: 2666  
LSP regenerations: 4  
Purges initiated: 0

member9:

-----  
IS-IS statistics for 0000.73e9.9a57:

| PDU type | Received | Processed | Drops | Sent  | Rexmit |
|----------|----------|-----------|-------|-------|--------|
| LSP      | 10800    | 10800     | 0     | 6327  | 0      |
| HELLO    | 1492     | 1492      | 0     | 2356  | 0      |
| CSNP     | 2        | 2         | 0     | 2     | 0      |
| PSNP     | 2683     | 2683      | 0     | 3149  | 0      |
| Unknown  | 0        | 0         | 0     | 0     | 0      |
| Totals   | 14977    | 14977     | 0     | 11834 | 0      |

Total packets received: 14977 Sent: 11834

LSP queue length: 0 Drops: 0  
SPF runs: 19  
Fragments rebuilt: 1510  
LSP regenerations: 6  
Purges initiated: 0

## show virtual-chassis

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show virtual-chassis</b><br><b>&lt;status&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF). <b>Fabric ID</b>, <b>Fabric Mode</b>, and <b>Route Mode</b> output fields introduced in Junos OS Release 13.2X51-D20.</p> <p><b>Alias-Name</b> output field introduced in Junos OS Release 14.1X53-D10.</p> |
| <b>Description</b>              | Display information about all members of the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>none</b>—Display information about all Virtual Chassis or VCF member devices.</p> <p><b>status</b>—Same output as for <b>show virtual-chassis</b>.</p>                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis active-topology on page 5174</a></li> <li>• <a href="#">show virtual-chassis protocol adjacency on page 5186</a></li> <li>• <a href="#">show virtual-chassis vc-path on page 5207</a></li> <li>• <a href="#">Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157</a></li> </ul>                                                        |
| <b>List of Sample Output</b>    | <p><a href="#">show virtual-chassis (EX4200 Virtual Chassis) on page 5205</a></p> <p><a href="#">show virtual-chassis (EX8200 Virtual Chassis) on page 5205</a></p> <p><a href="#">show virtual-chassis (Virtual Chassis Fabric) on page 5206</a></p>                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <p><a href="#">Table 579 on page 5203</a> lists the output fields for the <b>show virtual-chassis</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                  |

**Table 579: show virtual-chassis Output Fields**

| Field Name                  | Field Description                                                |
|-----------------------------|------------------------------------------------------------------|
| <b>Fabric ID</b>            | Assigned ID used to identify the VCF.                            |
| <b>Fabric Mode</b>          | Mode of the VCF: Enabled, Disabled, or Mixed.                    |
| <b>Virtual Chassis ID</b>   | Assigned ID that applies to the entire Virtual Chassis or VCF.   |
| <b>Virtual Chassis Mode</b> | Mode of the Virtual Chassis or VCF: Enabled, Disabled, or Mixed. |

Table 579: show virtual-chassis Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Member ID</b>           | Assigned member ID and FPC: <ul style="list-style-type: none"> <li>On all EX Series Virtual Chassis except EX8200 Virtual Chassis, and on a VCF, the FPC number refers to the member ID assigned to the switch.</li> <li>On EX8200 Virtual Chassis, member IDs are numbered 0 through 9. The FPC number indicates the slot number of the line card within the Virtual Chassis. The FPC number on member 0 is always 0 through 15. The FPC number on member 1 is always 16 through 31. The FPC number on member 2 is always 32 through 47; and so on for the members.</li> </ul>                                                                                                                         |
| <b>Status</b>              | For a nonprovisioned configuration: <ul style="list-style-type: none"> <li><b>Prsnt</b> for a member that is currently connected to the Virtual Chassis or VCF configuration.</li> <li><b>NotPrsnt</b> for a member ID that has been assigned but is not currently connected.</li> </ul> For a preprovisioned configuration: <ul style="list-style-type: none"> <li><b>Prsnt</b> for a member that is specified in the preprovisioned configuration file and is currently connected to the Virtual Chassis or VCF.</li> <li><b>Unprvsnd</b> for a member that is interconnected with the Virtual Chassis or VCF configuration but is not specified in the preprovisioned configuration file.</li> </ul> |
| <b>Serial No</b>           | Serial number of the member device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Alias-Name</b>          | The user-configured alias of the member device.<br><br>The <b>Alias-Name</b> field appears only if an alias has been configured for at least one device in the Virtual Chassis or VCF. Aliases are configured using the <b>alias-name</b> statement in the <code>[edit virtual-chassis aliases serial-number serial-number]</code> hierarchy.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Model</b>               | Model number of the member device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Mastership Priority</b> | Mastership priority value of the member device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Role</b>                | Role of the member device: master, backup, or linecard.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Mixed Mode</b>          | Mixed mode configuration status: <ul style="list-style-type: none"> <li><b>Y</b> for a member device configured in mixed mode.</li> <li><b>N</b> for a member device not configured in mixed mode.</li> <li><b>NA</b> for a member device that cannot be configured in mixed mode.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Route Mode</b>          | The route mode of the member device: fabric (F) or Virtual Chassis (V).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Location</b>            | Location of the member device.<br><br>If this field is empty, the location field was not set for the device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Neighbor List</b>       | Member ID of the neighbor member to which this member's Virtual Chassis port (VCP) is connected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## Sample Output

### show virtual-chassis (EX4200 Virtual Chassis)

```

user@switch> show virtual-chassis
Virtual Chassis ID: 0019.e250.47a0
Virtual Chassis Mode: Enabled

```

| Member ID | Status | Serial No    | Model      | Mastership<br>priority | Role     | Mixed<br>Mode | Neighbor List<br>ID | Interface |
|-----------|--------|--------------|------------|------------------------|----------|---------------|---------------------|-----------|
| 0 (FPC 0) | Prsnt  | AK0207360276 | ex4200-24t | 249                    | Master*  | N             | 8                   | vcp-0     |
|           |        |              |            |                        |          |               | 1                   | vcp-1     |
| 1 (FPC 1) | Prsnt  | AK0207360281 | ex4200-24t | 248                    | Backup   | N             | 0                   | vcp-0     |
|           |        |              |            |                        |          |               | 2                   | vcp-1     |
| 2 (FPC 2) | Prsnt  | AJ0207391130 | ex4200-48p | 247                    | Linecard | N             | 1                   | vcp-0     |
|           |        |              |            |                        |          |               | 3                   | vcp-1     |
| 3 (FPC 3) | Prsnt  | AK0207360280 | ex4200-24t | 246                    | Linecard | N             | 2                   | vcp-0     |
|           |        |              |            |                        |          |               | 4                   | vcp-1     |
| 4 (FPC 4) | Prsnt  | AJ0207391113 | ex4200-48p | 245                    | Linecard | N             | 3                   | vcp-0     |
|           |        |              |            |                        |          |               | 5                   | vcp-1     |
| 5 (FPC 5) | Prsnt  | BP0207452204 | ex4200-48t | 244                    | Linecard | N             | 4                   | vcp-0     |
|           |        |              |            |                        |          |               | 6                   | vcp-1     |
| 6 (FPC 6) | Prsnt  | BP0207452222 | ex4200-48t | 243                    | Linecard | N             | 5                   | vcp-0     |
|           |        |              |            |                        |          |               | 7                   | vcp-1     |
| 7 (FPC 7) | Prsnt  | BR0207432028 | ex4200-24f | 242                    | Linecard | N             | 6                   | vcp-0     |
|           |        |              |            |                        |          |               | 8                   | vcp-1     |
| 8 (FPC 8) | Prsnt  | BR0207431996 | ex4200-24f | 241                    | Linecard | N             | 7                   | vcp-0     |
|           |        |              |            |                        |          |               | 0                   | vcp-1     |

Member ID for next new member: 9 (FPC 9)

### show virtual-chassis (EX8200 Virtual Chassis)

```

user@external-routing-engine> show virtual-chassis
Virtual Chassis ID: c806.0842.de51
Virtual Chassis Mode: Enabled

```

| Member ID       | Status | Serial No    | Model  | Mastership<br>priority | Role     | Neighbor List<br>ID | Interface |
|-----------------|--------|--------------|--------|------------------------|----------|---------------------|-----------|
| 0 (FPC 0-15)    | Prsnt  | BA0908380001 | ex8216 | 0                      | Linecard | 8                   | vcp-0/0   |
|                 |        |              |        |                        |          | 8                   | vcp-0/1   |
|                 |        |              |        |                        |          | 1                   | vcp-4/0/4 |
| 1 (FPC 16-31)   | Prsnt  | BT0909411634 | ex8208 | 0                      | Linecard | 8                   | vcp-0/0   |
|                 |        |              |        |                        |          | 0                   | vcp-3/0/4 |
| 8 (FPC 128-143) | Prsnt  | 062009000021 | ex-xre | 128                    | Master   | 9                   | vcp-1/0   |
|                 |        |              |        |                        |          | 1                   | vcp-1/2   |

```

9 (FPC 144-159) Prsnt 062009000022 ex-xre 128 Backup*
9 vcp-1/3
0 vcp-2/0
9 vcp-2/1
0 vcp-1/1
8 vcp-1/0
8 vcp-1/2
8 vcp-1/3
8 vcp-1/3

```

### show virtual-chassis (Virtual Chassis Fabric)

```

user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 0282.5fa0.3f08
Fabric Mode: Enabled

```

| List        | Member ID | Status       | Serial No    | Model       | Mstr<br>prio | Role    | Mixed | Route | Neighbor<br>Mode ID |
|-------------|-----------|--------------|--------------|-------------|--------------|---------|-------|-------|---------------------|
| Interface   | 0 (FPC 0) | Prsnt        | AB3112430001 | qfx5100-48s | 129          | Master* | N     | F     | 3                   |
| vcp-255/1/0 |           |              |              |             |              |         |       |       | 2                   |
| vcp-255/1/1 |           |              |              |             |              |         |       |       | 4                   |
| vcp-255/1/2 |           |              |              |             |              |         |       |       | 4                   |
| vcp-255/1/3 |           |              |              |             |              |         |       |       | 4                   |
| 1 (FPC 1)   | Prsnt     | AB3112230001 | qfx5100-48s  | 129         | Backup       | N       | F     | 3     |                     |
| vcp-255/1/0 |           |              |              |             |              |         |       | 2     |                     |
| vcp-255/1/1 |           |              |              |             |              |         |       | 4     |                     |
| vcp-255/1/2 |           |              |              |             |              |         |       | 4     |                     |
| vcp-255/1/3 |           |              |              |             |              |         |       | 4     |                     |
| 2 (FPC 2)   | Prsnt     | AB3112460011 | qfx5100-48s  | 0           | Linecard     | N       | F     | 1     |                     |
| vcp-255/1/0 |           |              |              |             |              |         |       | 0     |                     |
| vcp-255/1/1 |           |              |              |             |              |         |       | 0     |                     |
| 3 (FPC 3)   | Prsnt     | AB3112460011 | qfx5100-48s  | 0           | Linecard     | N       | F     | 1     |                     |
| vcp-255/1/0 |           |              |              |             |              |         |       | 0     |                     |
| vcp-255/1/1 |           |              |              |             |              |         |       | 0     |                     |
| 4 (FPC 4)   | Prsnt     | AB3112430011 | qfx5100-48s  | 0           | Linecard     | N       | F     | 1     |                     |
| vcp-255/1/0 |           |              |              |             |              |         |       | 0     |                     |
| vcp-255/1/1 |           |              |              |             |              |         |       | 0     |                     |

## show virtual-chassis vc-path

|                                 |                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show virtual-chassis vc-path source-interface <i>interface-name</i> destination-interface <i>interface-name</i></b>                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.6 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.                                                                                                                                                                                |
| <b>Description</b>              | Show the path a packet takes when going from a source interface to a destination interface in a Virtual Chassis configuration.                                                                                                                                                                                              |
| <b>Options</b>                  | <b>source-interface <i>interface-name</i></b> —Name of the interface from which the packet originates<br><b>destination-interface <i>interface-name</i></b> —Name of the interface to which the packet is delivered                                                                                                         |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157</a></li> <li>• <a href="#">Understanding EX Series Virtual Chassis Configuration on page 5086</a></li> <li>• <a href="#">EX8200 Virtual Chassis Overview</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis vc-path source-interface destination-interface on page 5208</a>                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | <a href="#">Table 580 on page 5207</a> lists the output fields for the <b>show virtual-chassis vc-path</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                 |

**Table 580: show virtual-chassis vc-path Output Fields**

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                               |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hop</b>        | The number of hops between the source and destination interfaces.                                                                                                                                                                                                                                                                                                               |
| <b>Member</b>     | The Virtual Chassis ID of the member switch that contains the Packet Forwarding Engine for each intermediate hop.                                                                                                                                                                                                                                                               |
| <b>PFE-Device</b> | The number of the Packet Forwarding Engine in each Virtual Chassis member through which a packet passes. Each Packet Forwarding Engine is the next hop of the preceding Packet Forwarding Engine.                                                                                                                                                                               |
| <b>Interface</b>  | The name of the interface through which the Packet Forwarding Engines are connected. The interface for the first hop is always the source interface and the interface for the last hop is always the destination interface. For intermediate hops, the <b>Interface</b> field denotes the Packet Forwarding Engines through which the packet passes on its way to the next hop. |

## Sample Output

show virtual-chassis vc-path source-interface destination-interface

```
user@switch> show virtual-chassis vc-path source-interface ge-0/0/0 destination-interface
ge-1/0/1
vc-path from ge-0/0/0 to ge-1/0/1
Hop      Member    PFE-Device    Interface
0         0          1              ge-0/0/0
1         0          0              internal-1/24
2         1          3              vcp-0
3         1          4              ge-1/0/1
```



## show virtual-chassis vc-port

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis vc-port<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Display the status of the Virtual Chassis ports (VCPs), including both the dedicated VCPs and the uplink ports configured as VCPs.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>none</b>—Display the operational status of all VCPs of the member switch where the command is issued.</p> <p><b>all-members</b>—(Optional) Display the operational status of all VCPs on all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the operational status of the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the operational status of all VCPs for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis vc-port statistics on page 5227</a></li> <li>• <a href="#">Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157</a></li> <li>• <a href="#">Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis</a></li> </ul>                                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis vc-port (EX4200 Virtual Chassis) on page 5211</a><br><a href="#">show virtual-chassis vc-port (EX8200 Virtual Chassis) on page 5211</a><br><a href="#">show virtual-chassis vc-port all-members on page 5212</a>                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | Table 581 on page 5209 lists the output fields for the <b>show virtual-chassis vc-port</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                |

Table 581: show virtual-chassis vc-port Output Fields

| Field Name       | Field Description                            |
|------------------|----------------------------------------------|
| <i>fpcnumber</i> | The FPC number is the same as the member ID. |

Table 581: show virtual-chassis vc-port Output Fields (*continued*)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface or PIC/Port | <p>VCP name.</p> <ul style="list-style-type: none"> <li>The dedicated VCPs in an EX4200 or EX4500 Virtual Chassis are <b>vcp-0</b> and <b>vcp-1</b>. The dedicated VCPs in an EX4550 Virtual Chassis are <b>VCP-1/0</b>, <b>VCP-1/1</b>, <b>VCP-2/0</b>, and <b>VCP-2/1</b>.</li> <li>Optical ports set as VCPs are named <b>1/0</b> and <b>1/1</b>, representing the PIC number and the port number.</li> <li>The native VCP (port 0) on an XRE200 External Routing Engine in an EX8200 Virtual Chassis is named <b>vcp-0</b>.</li> <li>The VCPs on each Virtual Chassis Control Interface (VCCI) module in an XRE200 External Routing Engine are named using the <b>vcp-slot-number/port-number</b> convention; for instance, <b>vcp-1/0</b>.</li> <li>The VCPs on EX8200 member switches are named using the <b>vcp-slot-number/pic-number/interface-number</b> convention; for instance, <b>vcp-3/0/2</b>.</li> <li>A <b>255</b> as the first number in your port number indicates that your VCP is part of a Link Aggregation group (LAG) bundle. For instance, a display of <b>vcp-255/1/0</b> indicates that the dedicated VCP named <b>vcp-1/0</b> is part of a LAG bundle. A display of <b>vcp-255/1/0</b> indicates that an uplink port that was previously named <b>xe-0/1/0</b> is now part of a VCP LAG bundle.</li> </ul> |
| Type                  | <p>Type of VCP:</p> <ul style="list-style-type: none"> <li><b>Dedicated</b>—The rear panel VCP on an EX4200, EX4500, or EX4550 switch, or any VCP link connected to an XRE200 External Routing Engine in an EX8200 Virtual Chassis.</li> <li><b>Configured</b>—Optical port configured as a VCP.</li> <li><b>Auto-Configured</b>—Optical port autoconfigured as a VCP.</li> </ul> <p>See <a href="#">“Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)” on page 5109</a> or <a href="#">Setting a 10-Gigabit Ethernet Port as a Virtual Chassis Port in an EX8200 Virtual Chassis (CLI Procedure)</a> for information about configuring VCPs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Trunk ID              | <p>A positive-number ID assigned to a link aggregation group (LAG) formed by the Virtual Chassis. The trunk ID value is –1 if no trunk is formed. A LAG between uplink VCPs requires that the link speed be the same on connected interfaces and that at least two VCPs on one member be connected to at least two VCPs on the other member in an EX4200 or EX4500 Virtual Chassis.</p> <p>Dedicated VCP LAGs are assigned trunk IDs 1 and 2. Trunk IDs for LAGs formed with uplink VCPs therefore have values of 3 or greater.</p> <p>The trunk ID value changes if the link-adjacency state between LAG members changes; trunk membership is then allocated or deallocated.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Status                | <p>Interface status:</p> <ul style="list-style-type: none"> <li><b>absent</b>—Interface is not a VCP link.</li> <li><b>down</b>—VCP link is down.</li> <li><b>up</b>—VCP link is up.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Speed (mbps)          | Speed of the interface in megabits per second.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Neighbor ID/Interface | The Virtual Chassis member ID and interface of a VCP on a member that is connected to the interface or PIC/Port field in the same row as this interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## Sample Output

### show virtual-chassis vc-port (EX4200 Virtual Chassis)

```
user@switch> show virtual-chassis vc-port
```

```
fpc0:
```

| Interface<br>or<br>PIC / Port | Type            | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface   |
|-------------------------------|-----------------|-------------|--------|-----------------|----------------|-------------|
| vcp-0                         | Dedicated       | 1           | Up     | 32000           | 1              | vcp-1       |
| vcp-1                         | Dedicated       | 2           | Up     | 32000           | 0              | vcp-0       |
| 1/0                           | Auto-Configured | 3           | Up     | 1000            | 2              | vcp-255/1/0 |
| 1/0                           | Auto-Configured | 3           | Up     | 1000            | 2              | vcp-255/1/1 |

### show virtual-chassis vc-port (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis vc-port
```

```
member0:
```

| Interface<br>or<br>Slot/PIC/Port | Type       | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface |
|----------------------------------|------------|-------------|--------|-----------------|----------------|-----------|
| vcp-0/0                          | Dedicated  | -1          | Up     | 1000            | 8              | vcp-1/1   |
| vcp-0/1                          | Dedicated  | -1          | Up     | 1000            | 8              | vcp-2/0   |
| 4/0/4                            | Configured | -1          | Up     | 10000           | 1              | vcp-3/0/4 |
| 4/0/7                            | Configured | -1          | Down   | 10000           |                |           |
| 4/0/3                            | Configured |             | Absent |                 |                |           |
| 4/0/2                            | Configured |             | Absent |                 |                |           |
| 4/0/5                            | Configured |             | Absent |                 |                |           |
| 4/0/6                            | Configured |             | Absent |                 |                |           |
| 4/0/1                            | Configured |             | Absent |                 |                |           |
| 4/0/0                            | Configured |             | Absent |                 |                |           |

```
member1:
```

| Interface<br>or<br>Slot/PIC/Port | Type       | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface |
|----------------------------------|------------|-------------|--------|-----------------|----------------|-----------|
| vcp-0/0                          | Dedicated  | -1          | Up     | 1000            | 8              | vcp-1/2   |
| 3/0/0                            | Configured | -1          | Down   | 10000           |                |           |
| 3/0/1                            | Configured | -1          | Down   | 10000           |                |           |
| 3/0/4                            | Configured | -1          | Up     | 10000           | 0              | vcp-4/0/4 |
| 3/0/5                            | Configured |             | Absent |                 |                |           |
| 4/0/5                            | Configured |             | Absent |                 |                |           |
| 4/0/4                            | Configured |             | Absent |                 |                |           |

```
member8:
```

| Interface<br>or<br>Slot/PIC/Port | Type      | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface |
|----------------------------------|-----------|-------------|--------|-----------------|----------------|-----------|
| vcp-0/0                          | Dedicated | -1          | Down   | 1000            |                |           |
| vcp-1/0                          | Dedicated | -1          | Up     | 1000            | 9              | vcp-1/0   |
| vcp-1/1                          | Dedicated | -1          | Up     | 1000            | 0              | vcp-0/0   |
| vcp-1/2                          | Dedicated | -1          | Up     | 1000            | 1              | vcp-0/0   |
| vcp-1/3                          | Dedicated | -1          | Up     | 1000            | 9              | vcp-1/3   |
| vcp-2/0                          | Dedicated | -1          | Up     | 1000            | 0              | vcp-0/1   |
| vcp-2/1                          | Dedicated | -1          | Up     | 1000            | 9              | vcp-1/2   |
| vcp-2/2                          | Dedicated | -1          | Down   | 1000            |                |           |

```
vcp-2/3      Dedicated      -1   Down      1000
```

```
member9:
```

| Interface<br>or<br>Slot/PIC/Port | Type      | Trunk<br>ID | Status   | Speed<br>(mbps) | Neighbor<br>ID | Interface |
|----------------------------------|-----------|-------------|----------|-----------------|----------------|-----------|
| vcp-0/0                          | Dedicated | -1          | Disabled | 1000            |                |           |
| vcp-1/0                          | Dedicated | -1          | Up       | 1000            | 8              | vcp-1/0   |
| vcp-1/1                          | Dedicated | -1          | Down     | 1000            |                |           |
| vcp-1/2                          | Dedicated | -1          | Up       | 1000            | 8              | vcp-2/1   |
| vcp-1/3                          | Dedicated | -1          | Up       | 1000            | 8              | vcp-1/3   |

### show virtual-chassis vc-port all-members

```
user@switch> show virtual-chassis vc-port all-members
```

```
fpc0:
```

| Interface<br>or<br>PIC / Port | Type            | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface   |
|-------------------------------|-----------------|-------------|--------|-----------------|----------------|-------------|
| vcp-0                         | Dedicated       | 1           | Up     | 32000           | 1              | vcp-1       |
| vcp-1                         | Dedicated       | 2           | Up     | 32000           | 0              | vcp-0       |
| 1/0                           | Auto-Configured | 3           | Up     | 1000            | 2              | vcp-255/1/0 |
| 1/1                           | Auto-Configured | 3           | Up     | 1000            | 2              | vcp-255/1/1 |

```
fpc1:
```

| Interface<br>or<br>PIC / Port | Type            | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface   |
|-------------------------------|-----------------|-------------|--------|-----------------|----------------|-------------|
| vcp-0                         | Dedicated       | 1           | Up     | 32000           | 0              | vcp-1       |
| vcp-1                         | Dedicated       | 2           | Up     | 32000           | 0              | vcp-0       |
| 1/0                           | Auto-Configured | -1          | Up     | 1000            | 3              | vcp-255/1/0 |

```
fpc2:
```

| Interface<br>or<br>PIC / Port | Type            | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface   |
|-------------------------------|-----------------|-------------|--------|-----------------|----------------|-------------|
| vcp-0                         | Dedicated       | 1           | Up     | 32000           | 3              | vcp-1       |
| vcp-1                         | Dedicated       | 2           | Up     | 32000           | 3              | vcp-0       |
| 1/0                           | Auto-Configured | 3           | Up     | 1000            | 0              | vcp-255/1/0 |
| 1/1                           | Auto-Configured | 3           | Up     | 1000            | 0              | vcp-255/1/1 |

```
fpc3:
```

| Interface<br>or<br>PIC / Port | Type            | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface   |
|-------------------------------|-----------------|-------------|--------|-----------------|----------------|-------------|
| vcp-0                         | Dedicated       | 1           | Up     | 32000           | 2              | vcp-0       |
| vcp-1                         | Dedicated       | 2           | Up     | 32000           | 2              | vcp-1       |
| 1/0                           | Auto-Configured | -1          | Up     | 1000            | 1              | vcp-255/1/0 |

## show virtual-chassis vc-port diagnostics optics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show virtual-chassis vc-port diagnostics optics &lt;all-members&gt; &lt;interface-name&gt; &lt;local&gt; &lt;member member-id&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Display diagnostics data and alarms for Ethernet optical transceivers installed in ports configured as Virtual Chassis Ports (VCPs) in an EX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that a transceiver is not operating properly. DOM information can be used to diagnose why a transceiver is not working.</p> <p>On some EX Series switches, the <b>request virtual-chassis vc-port diagnostics optics</b> command must be entered to run a diagnostic scan before you can gather the <b>show virtual-chassis vc-port diagnostics optics</b> output.</p> |
| <b>Options</b>                  | <p><b>none</b>—Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display diagnostics information for the transceiver installed in a specified VCP.</p> <p><b>local</b>—(Optional) Display diagnostics information for transceivers installed in VCPs on the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display diagnostics information for transceivers installed in VCPs on a specified member of a Virtual Chassis or VCF.</p>                                                       |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis vc-port on page 5209</a></li> <li>• <i>Installing a Transceiver in an EX Series Switch</i></li> <li>• <i>Removing a Transceiver from an EX Series Switch</i></li> <li>• <a href="#">Junos OS Ethernet Interfaces Configuration Guide</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <p><a href="#">show virtual-chassis vc-port diagnostics optics on page 5216</a></p> <p><a href="#">show virtual-chassis vc-port diagnostics optics (interface-name) on page 5221</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

[show virtual-chassis vc-port diagnostics optics local](#) on page 5223

[show virtual-chassis vc-port diagnostics optics \(member member-id\)](#) on page 5225

**Output Fields** [Table 582 on page 5214](#) lists the output fields for the **show virtual-chassis vc-port diagnostics optics** command. Output fields are listed in the approximate order in which they appear.

**Table 582: show virtual-chassis vc-port diagnostics optics Output Fields**

| Field Name                            | Field Description                                                                                                                                                       |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FPC                                   | Displays the FPC slot number.                                                                                                                                           |
| Virtual chassis port                  | Displays the name of the VCP.                                                                                                                                           |
| Laser bias current                    | Displays the magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents. |
| Laser output power                    | Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).                                                                              |
| Module temperature                    | Displays the temperature, in Celsius and Fahrenheit.                                                                                                                    |
| Module voltage                        | Displays the voltage, in Volts.                                                                                                                                         |
| Receiver signal average optical power | Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).                                                           |
| Laser bias current high alarm         | Displays whether the laser bias power setting high alarm is <i>On</i> or <i>Off</i> .                                                                                   |
| Laser bias current low alarm          | Displays whether the laser bias power setting low alarm is <i>On</i> or <i>Off</i> .                                                                                    |
| Laser bias current high warning       | Displays whether the laser bias power setting high warning is <i>On</i> or <i>Off</i> .                                                                                 |
| Laser bias current low warning        | Displays whether the laser bias power setting low warning is <i>On</i> or <i>Off</i> .                                                                                  |
| Laser output power high alarm         | Displays whether the laser output power high alarm is <i>On</i> or <i>Off</i> .                                                                                         |
| Laser output power low alarm          | Displays whether the laser output power low alarm is <i>On</i> or <i>Off</i> .                                                                                          |
| Laser output power high warning       | Displays whether the laser output power high warning is <i>On</i> or <i>Off</i> .                                                                                       |
| Laser output power low warning        | Displays whether the laser output power low warning is <i>On</i> or <i>Off</i> .                                                                                        |
| Module temperature high alarm         | Displays whether the module temperature high alarm is <i>On</i> or <i>Off</i> .                                                                                         |
| Module temperature low alarm          | Displays whether the module temperature low alarm is <i>On</i> or <i>Off</i> .                                                                                          |
| Module temperature high warning       | Displays whether the module temperature high warning is <i>On</i> or <i>Off</i> .                                                                                       |
| Module temperature low warning        | Displays whether the module temperature low warning is <i>On</i> or <i>Off</i> .                                                                                        |

Table 582: show virtual-chassis vc-port diagnostics optics Output Fields (*continued*)

| Field Name                                | Field Description                                                                  |
|-------------------------------------------|------------------------------------------------------------------------------------|
| Module voltage high alarm                 | Displays whether the module voltage high alarm is <i>On</i> or <i>Off</i> .        |
| Module voltage low alarm                  | Displays whether the module voltage low alarm is <i>On</i> or <i>Off</i> .         |
| Module voltage high warning               | Displays whether the module voltage high warning is <i>On</i> or <i>Off</i> .      |
| Module voltage low warning                | Displays whether the module voltage low warning is <i>On</i> or <i>Off</i> .       |
| Laser rx power high alarm                 | Displays whether the receive laser power high alarm is <i>On</i> or <i>Off</i> .   |
| Laser rx power low alarm                  | Displays whether the receive laser power low alarm is <i>On</i> or <i>Off</i> .    |
| Laser rx power high warning               | Displays whether the receive laser power high warning is <i>On</i> or <i>Off</i> . |
| Laser rx power low warning                | Displays whether the receive laser power low warning is <i>On</i> or <i>Off</i> .  |
| Laser bias current high alarm threshold   | Displays the vendor-specified threshold for the laser bias current high alarm.     |
| Laser bias current low alarm threshold    | Displays the vendor-specified threshold for the laser bias current low alarm.      |
| Laser bias current high warning threshold | Displays the vendor-specified threshold for the laser bias current high warning.   |
| Laser bias current low warning threshold  | Displays the vendor-specified threshold for the laser bias current low warning.    |
| Laser output power high alarm threshold   | Displays the vendor-specified threshold for the laser output power high alarm.     |
| Laser output power low alarm threshold    | Displays the vendor-specified threshold for the laser output power low alarm.      |
| Laser output power high warning threshold | Displays the vendor-specified threshold for the laser output power high warning.   |
| Laser output power low warning threshold  | Displays the vendor-specified threshold for the laser output power low warning.    |
| Module temperature high alarm threshold   | Displays the vendor-specified threshold for the module temperature high alarm.     |
| Module temperature low alarm threshold    | Displays the vendor-specified threshold for the module temperature low alarm.      |
| Module temperature high warning threshold | Displays the vendor-specified threshold for the module temperature high warning.   |
| Module temperature low warning threshold  | Displays the vendor-specified threshold for the module temperature low warning.    |

Table 582: show virtual-chassis vc-port diagnostics optics Output Fields (*continued*)

| Field Name                            | Field Description                                                            |
|---------------------------------------|------------------------------------------------------------------------------|
| Module voltage high alarm threshold   | Displays the vendor-specified threshold for the module voltage high alarm.   |
| Module voltage low alarm threshold    | Displays the vendor-specified threshold for the module voltage low alarm.    |
| Module voltage high warning threshold | Displays the vendor-specified threshold for the module voltage high warning. |
| Module voltage low warning threshold  | Displays the vendor-specified threshold for the module voltage low warning.  |
| Laser rx power high alarm threshold   | Displays the vendor-specified threshold for the laser rx power high alarm.   |
| Laser rx power low alarm threshold    | Displays the vendor-specified threshold for the laser rx power low alarm.    |
| Laser rx power high warning threshold | Displays the vendor-specified threshold for the laser rx power high warning. |
| Laser rx power low warning threshold  | Displays the vendor-specified threshold for the laser rx power low warning.  |

## Sample Output

### show virtual-chassis vc-port diagnostics optics

```

user@switch> show virtual-chassis vc-port diagnostics optics
fpc0:
-----
Virtual chassis port: vcp-0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-1
  Optical diagnostics                : N/A

fpc1:
-----
Virtual chassis port: vcp-0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-1
  Optical diagnostics                : N/A

fpc2:
-----
Virtual chassis port: vcp-2/0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current                 : 4.130 mA
  Laser output power                  : 0.2450 mW / -6.11 dBm
  Module temperature                 : 32 degrees C / 90 degrees F
  Module voltage                     : 3.3530 V
  Receiver signal average optical power : 0.0971 mW / -10.13 dBm
  Laser bias current high alarm      : Off
  Laser bias current low alarm       : Off

```



```

Laser bias current high warning      : Off
Laser bias current low warning       : Off
Laser output power high alarm        : Off
Laser output power low alarm         : Off
Laser output power high warning      : Off
Laser output power low warning       : Off
Module temperature high alarm        : Off
Module temperature low alarm         : Off
Module temperature high warning      : Off
Module temperature low warning       : Off
Module voltage high alarm            : Off
Module voltage low alarm             : Off
Module voltage high warning          : Off
Module voltage low warning           : Off
Laser rx power high alarm            : Off
Laser rx power low alarm             : Off
Laser rx power high warning          : Off
Laser rx power low warning           : Off
Laser bias current high alarm threshold : 14.998 mA
Laser bias current low alarm threshold : 0.998 mA
Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold : 1.198 mA
Laser output power high alarm threshold : 0.7940 mW / -1.00 dBm
Laser output power low alarm threshold : 0.0790 mW / -11.02 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0990 mW / -10.04 dBm
Module temperature high alarm threshold : 85 degrees C / 185 degrees F
Module temperature low alarm threshold : -10 degrees C / 14 degrees F
Module temperature high warning threshold : 80 degrees C / 176 degrees F
Module temperature low warning threshold : -5 degrees C / 23 degrees F
Module voltage high alarm threshold : 3.600 V
Module voltage low alarm threshold : 3.000 V
Module voltage high warning threshold : 3.499 V
Module voltage low warning threshold : 3.099 V
Laser rx power high alarm threshold : 1.5848 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                  : 5.428 mA
Laser output power                  : 0.4760 mW / -3.22 dBm
Module temperature                  : 28 degrees C / 83 degrees F
Module voltage                      : 3.3440 V
Receiver signal average optical power : 0.4002 mW / -3.98 dBm
Laser bias current high alarm       : Off
Laser bias current low alarm        : Off
Laser bias current high warning     : Off
Laser bias current low warning      : Off
Laser output power high alarm       : Off
Laser output power low alarm        : Off
Laser output power high warning     : Off
Laser output power low warning      : Off
Module temperature high alarm       : Off
Module temperature low alarm        : Off
Module temperature high warning     : Off
Module temperature low warning      : Off
Module voltage high alarm           : Off
Module voltage low alarm            : Off
Module voltage high warning         : Off
Module voltage low warning          : Off
Laser rx power high alarm           : Off

```

```

Laser rx power low alarm           : Off
Laser rx power high warning        : Off
Laser rx power low warning         : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

fpc3:

-----  
Virtual chassis port: vcp-255/0/2

```

Laser bias current           : 7.876 mA
Laser output power           : 0.5330 mW / -2.73 dBm
Module temperature           : 26 degrees C / 78 degrees F
Module voltage               : 3.3060 V
Receiver signal average optical power : 0.4885 mW / -3.11 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm  : Off
Laser output power low alarm   : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm  : Off
Module temperature low alarm   : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm      : Off
Module voltage low alarm       : Off
Module voltage high warning    : Off
Module voltage low warning     : Off
Laser rx power high alarm      : Off
Laser rx power low alarm       : Off
Laser rx power high warning    : Off
Laser rx power low warning     : Off
Laser bias current high alarm threshold : 14.500 mA
Laser bias current low alarm threshold : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F

```

```

Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current : 5.052 mA
Laser output power : 0.5030 mW / -2.98 dBm
Module temperature : 24 degrees C / 75 degrees F
Module voltage : 3.2890 V
Receiver signal average optical power : 0.5028 mW / -2.99 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/4
Laser bias current : 7.978 mA
Laser output power : 0.5460 mW / -2.63 dBm
Module temperature : 24 degrees C / 76 degrees F

```

```

Module voltage                               : 3.3060 V
Receiver signal average optical power       : 0.6305 mW / -2.00 dBm
Laser bias current high alarm               : Off
Laser bias current low alarm               : Off
Laser bias current high warning            : Off
Laser bias current low warning            : Off
Laser output power high alarm              : Off
Laser output power low alarm              : Off
Laser output power high warning           : Off
Laser output power low warning            : Off
Module temperature high alarm              : Off
Module temperature low alarm              : Off
Module temperature high warning           : Off
Module temperature low warning            : Off
Module voltage high alarm                 : Off
Module voltage low alarm                 : Off
Module voltage high warning               : Off
Module voltage low warning               : Off
Laser rx power high alarm                 : Off
Laser rx power low alarm                 : Off
Laser rx power high warning              : Off
Laser rx power low warning              : Off
Laser bias current high alarm threshold   : 14.500 mA
Laser bias current low alarm threshold   : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold   : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold   : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold   : 75 degrees C / 167 degrees F
Module temperature low alarm threshold   : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold       : 3.630 V
Module voltage low alarm threshold       : 2.970 V
Module voltage high warning threshold     : 3.465 V
Module voltage low warning threshold     : 3.135 V
Laser rx power high alarm threshold       : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold       : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold     : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold     : 0.1023 mW / -9.90 dBm

```

fpc4:

```

-----
Virtual chassis port: vcp-0
  Optical diagnostics                       : N/A
Virtual chassis port: vcp-1
  Optical diagnostics                       : N/A
Virtual chassis port: vcp-255/0/4
  Laser bias current                       : 7.860 mA
  Laser output power                       : 0.5370 mW / -2.70 dBm
  Module temperature                       : 24 degrees C / 75 degrees F
  Module voltage                           : 3.2920 V
  Receiver signal average optical power    : 0.6271 mW / -2.03 dBm
  Laser bias current high alarm            : Off
  Laser bias current low alarm            : Off
  Laser bias current high warning          : Off
  Laser bias current low warning          : Off
  Laser output power high alarm            : Off
  Laser output power low alarm            : Off
  Laser output power high warning          : Off

```

```

Laser output power low warning      : Off
Module temperature high alarm       : Off
Module temperature low alarm        : Off
Module temperature high warning     : Off
Module temperature low warning      : Off
Module voltage high alarm           : Off
Module voltage low alarm            : Off
Module voltage high warning         : Off
Module voltage low warning          : Off
Laser rx power high alarm           : Off
Laser rx power low alarm            : Off
Laser rx power high warning         : Off
Laser rx power low warning          : Off
Laser bias current high alarm threshold : 14.500 mA
Laser bias current low alarm threshold : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

#### show virtual-chassis vc-port diagnostics optics (interface-name)

```

user@external-routing-engine> show virtual-chassis vc-port diagnostics optics vcp-255/0/3
fpc0:
-----

```

```

fpc1:
-----

```

```

fpc2:
-----

```

```

Virtual chassis port: vcp-255/0/3
Laser bias current      : 5.448 mA
Laser output power      : 0.4770 mW / -3.21 dBm
Module temperature      : 28 degrees C / 82 degrees F
Module voltage          : 3.3450 V
Receiver signal average optical power : 0.3973 mW / -4.01 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm  : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm  : Off

```

```

Module temperature high warning      : Off
Module temperature low warning       : Off
Module voltage high alarm            : Off
Module voltage low alarm             : Off
Module voltage high warning          : Off
Module voltage low warning           : Off
Laser rx power high alarm            : Off
Laser rx power low alarm             : Off
Laser rx power high warning          : Off
Laser rx power low warning           : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

fpc3:

-----  
Virtual chassis port: vcp-255/0/3

```

Laser bias current      : 5.040 mA
Laser output power      : 0.5020 mW / -2.99 dBm
Module temperature      : 24 degrees C / 74 degrees F
Module voltage          : 3.2870 V
Receiver signal average optical power : 0.5073 mW / -2.95 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm  : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm  : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm  : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm  : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA

```

```

Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold  : 2.500 mA
Laser output power high alarm threshold   : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold    : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold  : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold   : 75 degrees C / 167 degrees F
Module temperature low alarm threshold    : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold  : 0 degrees C / 32 degrees F
Module voltage high alarm threshold       : 3.630 V
Module voltage low alarm threshold        : 2.970 V
Module voltage high warning threshold     : 3.465 V
Module voltage low warning threshold      : 3.135 V
Laser rx power high alarm threshold       : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold        : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold     : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold      : 0.1023 mW / -9.90 dBm

```

fpc4:

-----

#### show virtual-chassis vc-port diagnostics optics local

```

user@switch> show virtual-chassis vc-port diagnostics optics local
Virtual chassis port: vcp-2/0
  Optical diagnostics : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current : 4.130 mA
  Laser output power : 0.2450 mW / -6.11 dBm
  Module temperature : 32 degrees C / 90 degrees F
  Module voltage     : 3.3530 V
  Receiver signal average optical power : 0.0961 mW / -10.17 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm  : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm : Off
  Module temperature low alarm  : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Module voltage high alarm : Off
  Module voltage low alarm  : Off
  Module voltage high warning : Off
  Module voltage low warning : Off
  Laser rx power high alarm : Off
  Laser rx power low alarm  : Off
  Laser rx power high warning : Off
  Laser rx power low warning : Off
  Laser bias current high alarm threshold : 14.998 mA
  Laser bias current low alarm threshold  : 0.998 mA

```

```
Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold  : 1.198 mA
Laser output power high alarm threshold   : 0.7940 mW / -1.00 dBm
Laser output power low alarm threshold    : 0.0790 mW / -11.02 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold  : 0.0990 mW / -10.04 dBm
Module temperature high alarm threshold    : 85 degrees C / 185 degrees F
Module temperature low alarm threshold     : -10 degrees C / 14 degrees F
Module temperature high warning threshold  : 80 degrees C / 176 degrees F
Module temperature low warning threshold   : -5 degrees C / 23 degrees F
Module voltage high alarm threshold        : 3.600 V
Module voltage low alarm threshold         : 3.000 V
Module voltage high warning threshold      : 3.499 V
Module voltage low warning threshold       : 3.099 V
Laser rx power high alarm threshold        : 1.5848 mW / 2.00 dBm
Laser rx power low alarm threshold         : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold      : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold       : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                        : 5.426 mA
Laser output power                        : 0.4760 mW / -3.22 dBm
Module temperature                        : 28 degrees C / 83 degrees F
Module voltage                            : 3.3450 V
Receiver signal average optical power     : 0.3955 mW / -4.03 dBm
Laser bias current high alarm             : Off
Laser bias current low alarm              : Off
Laser bias current high warning           : Off
Laser bias current low warning            : Off
Laser output power high alarm              : Off
Laser output power low alarm              : Off
Laser output power high warning           : Off
Laser output power low warning            : Off
Module temperature high alarm              : Off
Module temperature low alarm              : Off
Module temperature high warning           : Off
Module temperature low warning            : Off
Module voltage high alarm                  : Off
Module voltage low alarm                   : Off
Module voltage high warning                : Off
Module voltage low warning                 : Off
Laser rx power high alarm                  : Off
Laser rx power low alarm                   : Off
Laser rx power high warning                : Off
Laser rx power low warning                 : Off
Laser bias current high alarm threshold   : 10.500 mA
Laser bias current low alarm threshold    : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold  : 2.500 mA
Laser output power high alarm threshold   : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold    : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold  : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold   : 75 degrees C / 167 degrees F
Module temperature low alarm threshold    : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold  : 0 degrees C / 32 degrees F
Module voltage high alarm threshold        : 3.630 V
Module voltage low alarm threshold         : 2.970 V
Module voltage high warning threshold      : 3.465 V
Module voltage low warning threshold       : 3.135 V
Laser rx power high alarm threshold        : 1.5849 mW / 2.00 dBm
```



```

Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold    : 0.1023 mW / -9.90 dBm

```

#### show virtual-chassis vc-port diagnostics optics (member member-id)

```

user@switch> show virtual-chassis vc-port diagnostics optics member 2
fpc2:

```

```

-----
Virtual chassis port: vcp-2/0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current                  : 4.130 mA
  Laser output power                  : 0.2450 mW / -6.11 dBm
  Module temperature                  : 31 degrees C / 88 degrees F
  Module voltage                      : 3.3530 V
  Receiver signal average optical power : 0.0961 mW / -10.17 dBm
  Laser bias current high alarm       : Off
  Laser bias current low alarm        : Off
  Laser bias current high warning     : Off
  Laser bias current low warning      : Off
  Laser output power high alarm       : Off
  Laser output power low alarm        : Off
  Laser output power high warning     : Off
  Laser output power low warning      : Off
  Module temperature high alarm       : Off
  Module temperature low alarm        : Off
  Module temperature high warning     : Off
  Module temperature low warning      : Off
  Module voltage high alarm           : Off
  Module voltage low alarm            : Off
  Module voltage high warning         : Off
  Module voltage low warning          : Off
  Laser rx power high alarm           : Off
  Laser rx power low alarm            : Off
  Laser rx power high warning         : Off
  Laser rx power low warning          : Off
  Laser bias current high alarm threshold : 14.998 mA
  Laser bias current low alarm threshold : 0.998 mA
  Laser bias current high warning threshold : 14.000 mA
  Laser bias current low warning threshold : 1.198 mA
  Laser output power high alarm threshold : 0.7940 mW / -1.00 dBm
  Laser output power low alarm threshold : 0.0790 mW / -11.02 dBm
  Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
  Laser output power low warning threshold : 0.0990 mW / -10.04 dBm
  Module temperature high alarm threshold : 85 degrees C / 185 degrees F
  Module temperature low alarm threshold : -10 degrees C / 14 degrees F
  Module temperature high warning threshold : 80 degrees C / 176 degrees F
  Module temperature low warning threshold : -5 degrees C / 23 degrees F
  Module voltage high alarm threshold : 3.600 V
  Module voltage low alarm threshold : 3.000 V
  Module voltage high warning threshold : 3.499 V
  Module voltage low warning threshold : 3.099 V
  Laser rx power high alarm threshold : 1.5848 mW / 2.00 dBm

```

```
Laser rx power low alarm threshold      : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold   : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold    : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                      : 5.418 mA
Laser output power                      : 0.4770 mW / -3.21 dBm
Module temperature                      : 28 degrees C / 83 degrees F
Module voltage                          : 3.3450 V
Receiver signal average optical power   : 0.3964 mW / -4.02 dBm
Laser bias current high alarm           : Off
Laser bias current low alarm            : Off
Laser bias current high warning         : Off
Laser bias current low warning          : Off
Laser output power high alarm           : Off
Laser output power low alarm            : Off
Laser output power high warning         : Off
Laser output power low warning          : Off
Module temperature high alarm           : Off
Module temperature low alarm            : Off
Module temperature high warning         : Off
Module temperature low warning          : Off
Module voltage high alarm               : Off
Module voltage low alarm                : Off
Module voltage high warning             : Off
Module voltage low warning              : Off
Laser rx power high alarm               : Off
Laser rx power low alarm                : Off
Laser rx power high warning             : Off
Laser rx power low warning              : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold  : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold  : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold     : 3.630 V
Module voltage low alarm threshold      : 2.970 V
Module voltage high warning threshold   : 3.465 V
Module voltage low warning threshold    : 3.135 V
Laser rx power high alarm threshold     : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold    : 0.1023 mW / -9.90 dBm
```

## show virtual-chassis vc-port statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show virtual-chassis vc-port statistics &lt;all-members&gt; &lt;brief   detail   extensive &gt; &lt;interface-name&gt; &lt;local&gt; &lt;member member-id&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The options <b>all-members</b>, <b>brief</b>, <b>detail</b>, <b>extensive</b>, and <b>local</b> were added in Junos OS Release 9.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Display the traffic statistics collected on Virtual Chassis ports (VCPs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>none</b>—Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> and <b>extensive</b> options provide identical displays.</p> <p><b>all-members</b>—(Optional) Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display traffic statistics for the specified VCP.</p> <p><b>local</b>—(Optional) Display traffic statistics for VCPs on the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display traffic statistics for VCPs on the specified member of a Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear virtual-chassis vc-port statistics on page 5161</a></li> <li>• <a href="#">show virtual-chassis vc-port on page 5209</a></li> <li>• <a href="#">Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157</a></li> <li>• <a href="#">Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <p><a href="#">show virtual-chassis vc-port statistics on page 5230</a></p> <p><a href="#">show virtual-chassis vc-port statistics (EX8200 Virtual Chassis) on page 5231</a></p> <p><a href="#">show virtual-chassis vc-port statistics brief on page 5231</a></p> <p><a href="#">show virtual-chassis vc-port statistics extensive on page 5231</a></p> <p><a href="#">show virtual-chassis vc-port statistics member 0 on page 5233</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Output Fields** Table 582 on page 5214 lists the output fields for the **show virtual-chassis vc-port statistics** command. Output fields are listed in the approximate order in which they appear.

**Table 583: show virtual-chassis vc-port statistics Output Fields**

| Field Name                   | Field Description                                                                                                                                                                                  | Level of Output            |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|
| <b>fpcnumber</b>             | (All Virtual Chassis except EX8200 Virtual Chassis. VCF) ID of the Virtual Chassis member. The FPC number is the same as the member ID.                                                            | All levels                 |
| <b>member number</b>         | (EX8200 Virtual Chassis only) Member ID of the Virtual Chassis member.                                                                                                                             | All levels                 |
| <b>Interface</b>             | VCP name.                                                                                                                                                                                          | <b>brief</b>               |
| <b>Input Octets/Packets</b>  | Number of octets and packets received on the VCP.                                                                                                                                                  | <b>brief, member, none</b> |
| <b>Output Octets/Packets</b> | Number of octets and packets transmitted on the VCP.                                                                                                                                               | <b>brief, member, none</b> |
| <b>master: number</b>        | Member ID of the master Routing Engine.                                                                                                                                                            | All levels                 |
| <b>Port</b>                  | VCP for which <b>RX</b> (Receive) statistics, <b>TX</b> (Transmit) statistics, or both are reported by the VCP subsystem during a sampling interval—since the statistics counter was last cleared. | <b>detail, extensive</b>   |
| <b>Total octets</b>          | Total number of octets received and transmitted on the VCP.                                                                                                                                        | <b>detail, extensive</b>   |
| <b>Total packets</b>         | Total number of packets received and transmitted on the VCP.                                                                                                                                       | <b>detail, extensive</b>   |
| <b>Unicast packets</b>       | Number of unicast packets received and transmitted on the VCP.                                                                                                                                     | <b>detail, extensive</b>   |
| <b>Broadcast packets</b>     | Number of broadcast packets received and transmitted on the VCP.                                                                                                                                   | <b>detail, extensive</b>   |
| <b>Multicast packets</b>     | Number of multicast packets received and transmitted on the VCP.                                                                                                                                   | <b>detail, extensive</b>   |
| <b>MAC control frames</b>    | Number of media access control (MAC) control frames received and transmitted on the VCP.                                                                                                           | <b>detail, extensive</b>   |

Table 583: show virtual-chassis vc-port statistics Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output          |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>CRC alignment errors</b> | <p>Number of packets received on the VCP that had a length—excluding framing bits, but including frame check sequence (FCS) octets—of between 64 and 1518 octets, inclusive, and had one of the following errors:</p> <ul style="list-style-type: none"> <li>Invalid FCS with an integral number of octets (FCS error)</li> <li>Invalid FCS with a nonintegral number of octets (alignment error)</li> </ul>                                                                                                   | <b>detail, extensive</b> |
| <b>Oversize packets</b>     | Number of packets received on the VCP that were longer than 1518 octets (excluding framing bits, but including FCS octets) but were otherwise well formed.                                                                                                                                                                                                                                                                                                                                                     | <b>detail, extensive</b> |
| <b>Undersize packets</b>    | Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits but including FCS octets) and were otherwise well formed..                                                                                                                                                                                                                                                                                                                                                      | <b>detail, extensive</b> |
| <b>Jabber packets</b>       | <p>Number of packets received on the VCP that were longer than 1518 octets—excluding framing bits, but including FCS octets—and that had either an FCS error or an alignment error.</p> <p><b>NOTE:</b> This definition of <i>jabber</i> is different from the definition in IEEE-802.3 section 8.2.1.5 (10Base5) and section 10.3.1.4 (10Base2). These documents define <i>jabber</i> as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.</p> | <b>detail, extensive</b> |
| <b>Fragments received</b>   | <p>Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error.</p> <p>Fragment frames normally increment because both runs (which are normal occurrences caused by collisions) and noise hits are counted.</p>                                                                                                                                                                               | <b>detail, extensive</b> |
| <b>Ifout errors</b>         | Number of outbound packets received on the VCP that could not be transmitted because of errors.                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail, extensive</b> |
| <b>Packet drop events</b>   | Number of outbound packets received on the VCP that were dropped, rather than being encapsulated and sent out of the switch as fragments. The packet drop counter is incremented if a temporary shortage of packet memory causes packet fragmentation to fail.                                                                                                                                                                                                                                                 | <b>detail, extensive</b> |
| <b>64 octets frames</b>     | Number of packets received on the VCP (including invalid packets) that were 64 octets in length (excluding framing bits, but including FCS octets).                                                                                                                                                                                                                                                                                                                                                            | <b>detail, extensive</b> |

Table 583: show virtual-chassis vc-port statistics Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                 | Level of Output          |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>65–127 octets frames</b>    | Number of packets received on the VCP (including invalid packets) that were between 65 and 127 octets in length, inclusive (excluding framing bits, but including FCS octets).    | <b>detail, extensive</b> |
| <b>128–255 octets frames</b>   | Number of packets received on the VCP (including invalid packets) that were between 128 and 255 octets in length, inclusive (excluding framing bits, but including FCS octets).   | <b>detail, extensive</b> |
| <b>256–511 octets frames</b>   | Number of packets received on the VCP (including invalid packets) that were between 256 and 511 octets in length, inclusive (excluding framing bits, but including FCS octets).   | <b>detail, extensive</b> |
| <b>512–1023 octets frames</b>  | Number of packets received on the VCP (including invalid packets) that were between 512 and 1023 octets in length, inclusive (excluding framing bits, but including FCS octets).  | <b>detail, extensive</b> |
| <b>1024–1518 octets frames</b> | Number of packets received on the VCP (including invalid packets) that were between 1024 and 1518 octets in length, inclusive (excluding framing bits, but including FCS octets). | <b>detail, extensive</b> |
| <b>Rate packets per second</b> | Number of packets per second received and transmitted on the VCP.                                                                                                                 | <b>detail, extensive</b> |
| <b>Rate bytes per second</b>   | Number of bytes per second received and transmitted on the VCP.                                                                                                                   | <b>detail, extensive</b> |

## Sample Output

### show virtual-chassis vc-port statistics

```
user@switch> show virtual-chassis vc-port statistics
fpc0:
```

```
-----
Interface          Input  Octets/Packets      Output  Octets/Packets
internal-0/24       0      / 0                0      / 0
internal-0/25       0      / 0                0      / 0
internal-1/26       0      / 0                0      / 0
internal-1/27       0      / 0                0      / 0
vcp-0               0      / 0                0      / 0
vcp-1               0      / 0                0      / 0
internal-0/26       0      / 0                0      / 0
internal-0/27       0      / 0                0      / 0
internal-1/24       0      / 0                0      / 0
internal-1/25       0      / 0                0      / 0
```

```
{master:0}
```

### show virtual-chassis vc-port statistics (EX8200 Virtual Chassis)

```

user@external-routing-engine> show virtual-chassis vc-port statistics
member0:
-----
Interface          Input Octets/Packets      Output Octets/Packets
vcp-4/0/4          43171238 / 48152          47687133 / 51891
vcp-4/0/7          0 / 0                     0 / 0

member1:
-----
Interface          Input Octets/Packets      Output Octets/Packets
vcp-3/0/0          0 / 0                     0 / 0
vcp-3/0/1          0 / 0                     0 / 0
vcp-3/0/4          47695376 / 51899          43180556 / 48160

member8:
-----

member9:
-----

```

### show virtual-chassis vc-port statistics brief

```

user@switch> show virtual-chassis vc-port statistics brief
fpc0:
-----
Interface          Input Octets/Packets      Output Octets/Packets
internal-0/24       0 / 0                     0 / 0
internal-0/25       0 / 0                     0 / 0
internal-1/26       0 / 0                     0 / 0
internal-1/27       0 / 0                     0 / 0
vcp-0               0 / 0                     0 / 0
vcp-1               0 / 0                     0 / 0
internal-0/26       0 / 0                     0 / 0
internal-0/27       0 / 0                     0 / 0
internal-1/24       0 / 0                     0 / 0
internal-1/25       0 / 0                     0 / 0

{master:0}

```

### show virtual-chassis vc-port statistics extensive

```

user@switch> show virtual-chassis vc-port statistics extensive
fpc0:
-----

```

|                       | RX | TX |
|-----------------------|----|----|
| Port: internal-0/24   |    |    |
| Total octets:         | 0  | 0  |
| Total packets:        | 0  | 0  |
| Unicast packets:      | 0  | 0  |
| Broadcast packets:    | 0  | 0  |
| Multicast packets:    | 0  | 0  |
| MAC control frames:   | 0  | 0  |
| CRC alignment errors: | 0  |    |
| Oversize packets:     | 0  |    |
| Undersize packets:    | 0  |    |
| Jabber packets:       | 0  |    |
| Fragments received:   | 0  |    |

```
Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127  octets frames: 0
128-255 octets frames: 0
256-511 octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second: 0          0

...

Port: vcp-0
Total octets:          0          0
Total packets:         0          0
Unicast packets:       0          0
Broadcast packets:     0          0
Multicast packets:     0          0
MAC control frames:    0          0
CRC alignment errors:  0
Oversize packets:     0
Undersize packets:     0
Jabber packets:        0
Fragments received:    0
Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127  octets frames: 0
128-255 octets frames: 0
256-511 octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second: 0          0

Port: vcp-1
Total octets:          0          0
Total packets:         0          0
Unicast packets:       0          0
Broadcast packets:     0          0
Multicast packets:     0          0
MAC control frames:    0          0
CRC alignment errors:  0
Oversize packets:     0
Undersize packets:     0
Jabber packets:        0
Fragments received:    0
Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127  octets frames: 0
128-255 octets frames: 0
256-511 octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second: 0          0

...
```



```
{master:0}
```

#### show virtual-chassis vc-port statistics member 0

```
user@switch>show virtual-chassis vc-port statistics member 0  
fpc0:
```

```
-----  
Interface          Input  Octets/Packets      Output  Octets/Packets  
internal-0/24       0      / 0              0      / 0  
internal-0/25       0      / 0              0      / 0  
internal-1/26       0      / 0              0      / 0  
internal-1/27       0      / 0              0      / 0  
vcp-0               0      / 0              0      / 0  
vcp-1               0      / 0              0      / 0  
internal-0/26       0      / 0              0      / 0  
internal-0/27       0      / 0              0      / 0  
internal-1/24       0      / 0              0      / 0  
internal-1/25       0      / 0              0      / 0
```

```
{master:0}
```



# Troubleshooting

- [Troubleshooting Procedures on page 5235](#)

## Troubleshooting Procedures

---

- [Troubleshooting an EX Series Virtual Chassis on page 5235](#)

### Troubleshooting an EX Series Virtual Chassis

This topic describes the following troubleshooting issues for a Virtual Chassis:

- [A Disconnected Member Switch's ID Is Not Available for Reassignment on page 5235](#)
- [Load Factory Default Does Not Commit on a Multimember Virtual Chassis on page 5235](#)
- [The Member ID Persists When a Member Switch Is Disconnected From a Virtual Chassis on page 5236](#)
- [A Member Switch Is Not Participating in a Mixed Virtual Chassis on page 5236](#)

#### [A Disconnected Member Switch's ID Is Not Available for Reassignment](#)

---

**Problem** **Description:** You disconnected a switch from the Virtual Chassis, but the disconnected switch's member ID is still displayed in the status output. You cannot reassign that member ID to another switch.

**Solution** When you disconnect a member of a Virtual Chassis configuration, the master retains the member ID and member configuration in its configuration database. Output from the [show virtual-chassis](#) command continues to display the member ID of the disconnected member with a status of **NotPrsnt**.

If want to permanently disconnect the member switch, you can free up the member ID by using the [request virtual-chassis recycle](#) command. This will also clear the status of that member.

#### [Load Factory Default Does Not Commit on a Multimember Virtual Chassis](#)

---

**Problem** **Description:** The **load factory-default** command fails on a multimember Virtual Chassis.

**Solution** The **load factory-default** command is not supported on a multimember Virtual Chassis configuration. For information on how to revert the switches in the Virtual Chassis to factory default settings, see [“Reverting to the Default Factory Configuration for the EX Series Switch” on page 620](#).

### The Member ID Persists When a Member Switch Is Disconnected From a Virtual Chassis

---

**Problem** **Description:** Gigabit Ethernet interfaces retain their previous slot numbers when a member switch is disconnected from the Virtual Chassis.

**Solution** If a switch had been previously connected as a member of a Virtual Chassis configuration, it retains the member ID that it was assigned as a member of that configuration even after it is disconnected and operating as a standalone switch. The interfaces that were configured while the switch was a member of the Virtual Chassis configuration retain the old member ID as the first digit of the interface name.

For example, if the switch was previously member 1, its interfaces are named **ge-1/0/0** and so on.

To change the switch's member ID, so that its member ID is **0**, and to rename the switch's interfaces accordingly:

1. To change the member ID to 0:  

```
user@switch> request virtual-chassis renumber member-id 1 new-member-id 0
```
2. To rename the interfaces to match the new member ID:  

```
[edit virtual-chassis]  
user@switch# replace pattern ge-1/ with ge-0/
```

### A Member Switch Is Not Participating in a Mixed Virtual Chassis

---

**Problem** **Description:** A member switch in a mixed Virtual Chassis is not participating in the Virtual Chassis. The **show virtual-chassis** output indicates the member switch status is **Inactive** or **NotPrsnt**.

This issue is most likely to occur immediately after you have cabled a mixed Virtual Chassis.

**Solution** The Virtual Chassis mode on the switch might not be set to **mixed** mode. If the member switch is an EX4500 switch and is cabled into the Virtual Chassis through the dedicated Virtual Chassis port (VCP), the PIC mode might also be set to **Intraconnect** instead of **virtual-chassis**.

To verify the Virtual Chassis mode:

```
user@switch> show virtual-chassis mode  
fpc0:  
-----  
Mixed Mode: Enabled  
fpc1:  
-----
```

```
Mixed Mode: Enabled
fpc2:
```

```
-----
Mixed Mode: Enabled
fpc3:
```

```
-----
Mixed Mode: Enabled
fpc4:
```

```
-----
Mixed Mode: Disabled
fpc5:
```

```
-----
Mixed Mode: Enabled
```

To change the Virtual Chassis mode on a member switch (in this case, member ID 4) to **mixed** mode:

```
user@switch> request virtual-chassis mode mixed member 4
```

(EX4500 switch only) To verify the PIC mode:

```
user@switch> show chassis pic-mode
```

```
fpc0:
```

```
-----
Pic Mode: Not-Applicable
```

```
fpc1:
```

```
-----
Pic Mode: Not-Applicable
```

```
fpc2:
```

```
-----
Pic Mode: Not-Applicable
```

```
fpc3:
```

```
-----
Pic Mode: Not-Applicable
```

```
fpc4:
```

```
-----
Pic Mode: PIC 3: Intraconnect
```

```
fpc5:
```

```
-----
Pic Mode: PIC 3: virtual-chassis
```

To change the PIC mode on an EX4500 switch to **virtual-chassis** mode (in this case, member ID 4):

```
user@switch> request chassis pic-mode virtual-chassis member 4
```

The member switch must be rebooted for the Virtual Chassis mode or PIC mode setting change to take effect. To reboot the member switch (in this case, member ID 4):

```
user@switch> request system reboot member 4
```

#### Related Documentation

- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157](#)
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- *Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)*
- *Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)*



## PART 27

# Virtual Chassis Fabric

- [Overview on page 5241](#)
- [Configuration on page 5261](#)
- [Administration on page 5313](#)
- [Troubleshooting Procedures on page 5391](#)





## CHAPTER 88

# Overview

- [Virtual Chassis Fabric Overview on page 5241](#)

### Virtual Chassis Fabric Overview

---

- [Virtual Chassis Fabric Overview on page 5241](#)
- [Understanding Virtual Chassis Fabric Components on page 5243](#)
- [Understanding Virtual Chassis Fabric Configuration on page 5251](#)
- [Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric on page 5254](#)
- [Understanding Traffic Flow Through a Virtual Chassis Fabric on page 5258](#)
- [Understanding Software Upgrades in a Virtual Chassis Fabric on page 5258](#)

### Virtual Chassis Fabric Overview


The Juniper Networks Virtual Chassis Fabric (VCF) provides a low-latency, high-performance fabric architecture that can be managed as a single device. VCF is an evolution of the Virtual Chassis feature, which enables you to interconnect multiple devices into a single logical device, inside of a fabric architecture. The VCF architecture is optimized to support small and medium-sized data centers that contain a mix of 1-Gbps, 10-Gbps, and 40-Gbps Ethernet interfaces.



Video: [What is Virtual Chassis Fabric?](#)

A VCF is constructed using a spine-and-leaf architecture. In the spine-and-leaf architecture, each spine device is interconnected to each leaf device. A VCF supports up to thirty-two total devices, and up to four devices can be configured as spine devices. See [Figure 81 on page 5242](#) for an illustration of the VCF spine-and-leaf architecture.

Figure 81: VCF Spine-and-Leaf Architecture



Each spine device must be a QFX5100 device. In an optimal VCF configuration, the leaf devices are also QFX5100 devices. You can, however, also create a mixed VCF by configuring QFX3600, QFX3500, and EX4300 switches as leaf devices. See [“Understanding Virtual Chassis Fabric Components” on page 5243](#) for more information about the spine-and-leaf architecture.

A VCF provides the following benefits:

- **Latency**—VCF provides predictable low latency because it uses a fabric architecture that ensures each device is one or two hops away from every other device in the fabric. The weighted algorithm that makes traffic-forwarding decisions in a VCF is designed to avoid congestion and ensures low latency by intelligently forwarding traffic over all paths within the VCF to any destination device., ensuring predictable low latency for all traffic traversing the VCF.
- **Resiliency**—The VCF architecture provides a resilient framework because traffic has multiple paths across the fabric. Traffic is, therefore, easily diverted within the fabric when a device or link fails.
- **Flexibility**—You can easily expand the size of your VCF by adding devices to the fabric as your networking needs grow.
- **Investment protection**—In environments that need to expand because the capabilities of a traditional QFX5100, QFX3600, QFX3500, or EX4300 Virtual Chassis are maximized, a VCF is often a logical upgrade option because it enables the system to evolve without having to remove the existing, previously purchased devices from the network.
- **Manageability**—VCF provides multiple features that simplify configuration and management. VCF, for instance, has an autoprovisioning feature that enables you to plug and play devices into the fabric after minimal initial configuration. VCF leverages many of the existing configuration procedures from a Virtual Chassis, so that you can configure and maintain a VCF easily if you are already familiar with the procedures for configuring and maintaining a Virtual Chassis.

#### Related Documentation

- [Network Configuration Example: MetaFabric™ Architecture 1.1: Configuring Virtual Chassis Fabric and Network Director 1.6](#)
- [Understanding Virtual Chassis Fabric Components on page 5243](#)

- [Understanding Virtual Chassis Fabric Configuration on page 5251](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)

## Understanding Virtual Chassis Fabric Components

This topic describes the components of a Virtual Chassis Fabric (VCF).

This topic covers:

- [Spine-and-Leaf Topology on page 5243](#)
- [Spine Devices on page 5244](#)
- [Leaf Devices on page 5244](#)
- [Routing Engine Role on page 5245](#)
- [Linecard Role on page 5246](#)
- [Master Routing Engine Election Process on page 5246](#)
- [Virtual Chassis Ports \(VCPs\) on page 5247](#)
- [Automatic Virtual Chassis Port \(VCP\) Conversion on page 5247](#)
- [VCF Configuration Options on page 5248](#)
- [Fabric Mode on page 5248](#)
- [Mixed Mode on page 5249](#)
- [Virtual Management Ethernet Interface on page 5249](#)
- [Virtual Chassis Fabric Port Link Aggregation Group Bundles on page 5249](#)
- [Virtual Chassis Fabric License Requirements on page 5250](#)
- [Hardware Requirements for a Virtual Chassis Fabric on page 5250](#)
- [Software Requirements in a Virtual Chassis Fabric on page 5250](#)

---

### Spine-and-Leaf Topology


The VCF uses a spine-and-leaf architecture where each device in the fabric is either a spine device or a leaf device.

A VCF can have up to four spine devices, and up to thirty-two total devices. Each spine device has at least one direct Virtual Chassis port (VCP) connection to each leaf device in the VCF.

All traffic entering a leaf device can, therefore, be forwarded to any directly connected spine device and is always two hops away from any other leaf device—leaf device to leaf device traffic travels from the source leaf device to a spine device to the destination leaf device—within the VCF.

See [Figure 82 on page 5244](#) for an illustration of the VCF spine-and-leaf architecture:

Figure 82: VCF Spine-and-Leaf Architecture



Traffic is forwarded through a VCF using a weighted algorithm designed to avoid congestion. Traffic travelling across the VCF from one leaf device to another leaf device is forwarded using the best path available at the time, so any connection to a spine device can be used to transport traffic from one leaf device to another leaf device.

### Spine Devices

A spine device:

- Must be a QFX5100 device.
- Is configured into the Routing Engine role, but can operate in the Routing Engine role or the linecard role..



**NOTE:** Only two devices can simultaneously operate in the Routing Engine role within a VCF. A VCF, however, supports up to four spine devices. In scenarios where a VCF has three or more spine devices, the devices that are not operating in the Routing Engine role operate in the linecard role.

A spine device that is configured into the Routing Engine role but is operating in the linecard role assumes the Routing Engine role when an active Routing Engine fails.

- Has a direct connection to each leaf device.
- Typically connects a router, firewall, or other data center networking device to the VCF.

A VCF should always have at least two active spine devices. A VCF supports up to four spine devices.

You can configure any QFX5100 device as a spine device. In the most common VCF configurations, QFX5100-24Q devices are used as spine devices.

### Leaf Devices

A leaf device:

- Is optimally a QFX5100 device, but can also be a QFX3500, QFX3600, or EX4300 device.
- Has a direct connection to each spine device.
- Always operates in the linecard role.
- Typically connects an endpoint device—for instance, a server or other storage device in a data center—to the VCF.

A VCF can have up to thirty-two total devices and up to four devices can be configuring into spine devices. The devices that are not spine devices in a VCF operate as leaf devices.

In the most common VCF configurations, QFX5100-48S devices are used as leaf devices.

### Routing Engine Role

A VCF has two devices operating in the Routing Engine role—a master Routing Engine and a backup Routing Engine.

The device that functions as the master Routing Engine:

- Is a spine device.
- Manages the member devices.
- Runs the chassis management processes and control protocols.
- Represents all the member devices interconnected within the VCF configuration. (The hostname and other parameters that you assign to this device during setup apply to all members of the VCF.)

The device that functions as the backup Routing Engine:

- Is a spine device.
- Maintains a state of readiness to take over the master role if the master fails.
- Synchronizes with the master in terms of protocol states, forwarding tables, and so forth, so that it preserves routing information and maintains network connectivity without disruption when the master is unavailable.

All spine devices in a VCF are configured into the Routing Engine role.



**NOTE:** Only two devices can simultaneously operate in the Routing Engine role within a VCF. A VCF, however, supports up to four spine devices. In scenarios where a VCF has three or more spine devices, the devices that are not operating in the Routing Engine role operate in the linecard role.

A spine device that is configured into the Routing Engine role but is operating in the linecard role assumes the Routing Engine role when an active Routing Engine fails.

The master and backup Routing Engines are selected by the master election algorithm. If a VCF has more than two spine devices, the spine devices that are not selected as the master or backup Routing Engine operate in the linecard role even though the devices are configured into the Routing Engine role.

A spine device operating in the linecard role can complete all spine-related functions with no limitations within a VCF. A spine device operating in the linecard role assumes the backup Routing Engine role when the master or backup Routing Engine fails.

### Linecard Role

---

All leaf devices in a VCF operate in the linecard role. In autoprovisioned configurations, leaf devices are assigned the linecard role when they are cabled into the VCF. In preprovisioned configurations, leaf devices are manually configured into the linecard role. In nonprovisioned configurations, leaf devices are assigned the linecard role according to the master election algorithm, which uses the mastership priority values to set the roles of each device in the VCF.

All spine devices in a VCF are configured into the Routing Engine role, but operate in the linecard role when they are not selected as the master or backup Routing Engine by the master election algorithm. A spine device operating in the linecard role can complete all spine-related functions with no limitations within a VCF. A spine device operating in the linecard role becomes the new backup Routing Engine when the master or backup Routing Engine fails.

A member that functions in the linecard role in a VCF:

- Runs only a subset of Junos OS.
- Detects certain error conditions (such as an unplugged cable) on any interfaces that have been configured on it through the device functioning as the master Routing Engine.

### Master Routing Engine Election Process

---

The device in the master Routing Engine role in a VCF is always a spine device.

In a preprovisioned or autoprovisioned VCF, up to four spine devices are assigned the Routing Engine role during the configuration process. The spine device that has been powered on the longest assumes the master Routing Engine role; the spine device that has been powered on the second longest assumes the backup Routing Engine role. The remaining spine devices assume the linecard role.

In a nonprovisioned VCF, the master and backup Routing Engines are selected using the following algorithm:

1. Choose the QFX5100 device with the highest user-configured mastership priority (255 is the highest possible value) as the master Routing Engine, and the QFX5100 switch with the second highest mastership priority value as the backup Routing Engine.  
  
A QFX5100 switch with a mastership priority of 0 will always stay in the linecard role.
2. Choose the QFX5100 device that was master the last time the VCF booted.

3. Choose the QFX5100 device that has been included in the VCF configuration for the longest period of time.
4. Choose the QFX5100 device with the lowest MAC address.

QFX3500, QFX3600, and EX4300 devices never assume the master or backup Routing Engine role in a VCF.

We strongly recommend that you configure the mastership priority of the QFX5100 devices in your VCF to ensure that the correct devices assume their intended roles when you configure your VCF using a nonprovisioned configuration.

### Virtual Chassis Ports (VCPs)

Virtual Chassis ports (VCPs) are used in a VCF to interconnect leaf devices to spine devices. All control and data traffic in a VCF is transported over VCPs.

VCPs in a VCF are either SFP+ connections that support 10-Gbps or QSFP+ connections that support 40-Gbps.

10-Gbps SFP+ and 40-Gbps QSFP+ links are automatically converted into VCPs in most scenarios when a device is added to an autoprovisioned or preprovisioned VCF. Automatic VCP conversion is discussed in more detail in the following section.

You can manually configure a 10-Gbps SFP+ and 40-Gbps QSFP+ link into a VCP.

Channelized interfaces cannot be configured into VCPs.

### Automatic Virtual Chassis Port (VCP) Conversion

10-Gbps SFP+ and 40-Gbps QSFP+ links are not configured into VCPs, by default.

10-Gbps SFP+ and 40-Gbps QSFP+ links are automatically converted into VCPs when:

- Link Layer Discovery Protocol (LLDP) is enabled on the interfaces on both ends of the link. LLDP is enabled by default.
- the device being added to the VCF is configured into fabric mode.
- The interfaces on both ends of the link are not configured as VCPs. The following interfaces are configured as VCPs:
  - The 40-Gbps QSFP+ port on an EX4300 switch, by default.
  - Any interface in the VCF that has been a VCP. If a device is removed from a VCF, the interface that was interconnected to the removed device remains configured as a VCP until it is configured into a network port using the **request virtual-chassis vc-port delete** command.
  - Any interface that has been configured into a VCP using the **request virtual-chassis vc-port set** command.

To change any of the above interfaces into a network interface so that the interface can become eligible for automatic VCP conversion, use the **request virtual-chassis vc-port delete** command.

- one of the devices is already part of a VCF that was autoprovisioned or preprovisioned.

Automatic VCP conversion does not work in nonprovisioned VCFs.

Automatic VCP conversion does not convert a VCP interface into a network interface when a device is removed from a VCF. If automatic VCP conversion has converted an interface into a VCP and you want the interface to function as a network interface, you must manually disable the VCP interface.

---

### VCF Configuration Options

You can configure a VCF using autoprovisioned, preprovisioned, or nonprovisioned configuration.

Autoprovisioned configuration allows you to *plug and play* leaf devices into a VCF after completing a minimal initial configuration procedure.

In a preprovisioned configuration, you deterministically control the devices in your VCF by associating each device's serial number to a member ID and role.

Nonprovisioned configuration is possible, but not recommended for most VCF installations. Nonprovisioned configuration is a highly manual procedure that should only be performed by expert users.

See [“Understanding Virtual Chassis Fabric Configuration” on page 5251](#) for additional information on the VCF configuration options.

---

### Fabric Mode

A device must be configured into fabric mode in order for it to join a VCF. You should always configure a device into fabric mode before interconnecting it into a VCF.

In preprovisioned and nonprovisioned configurations, a device is not participating as a VCF member until it is configured into fabric mode.

In autoprovisioned configurations, a spine device is not participating as a VCF member until it is configured into fabric mode. A spine device that is not configured into fabric mode is configured into fabric mode when it is interconnected into the VCF. The final step of the process of configuring the device into fabric mode is a device reboot. We strongly recommend configuring the spine device into fabric mode before interconnecting it into the VCF to eliminate this reboot.

A leaf device in an autoprovisioned configuration is also rebooted to complete the fabric mode configuration when it is interconnected into a VCF without being set into fabric mode. You can avoid the downtime that accompanies the reboot by setting the device into fabric mode before interconnecting it into the VCF.

A standalone device that is not part of a VCF should never be configured into fabric mode. A device is not in fabric mode, by default.



---

### Mixed Mode

---

The optimal method of configuring a VCF is to use QFX5100 devices only. A VCF composed entirely of QFX5100 devices supports the largest breadth of features at the highest scalability while also supporting the highest number of high-speed interfaces.

You can, however, configure other devices as leaf devices in your VCF. QFX5100, QFX3600, QFX3500, or EX4300 devices can be used as leaf devices in a VCF.

If you use QFX3600, QFX3500, or EX4300 devices as leaf devices in your VCF, you must configure all devices in your VCF into mixed mode.

A device that is not part of a Virtual Chassis or a VCF with other devices should never be configured into mixed mode. A device is not configured into mixed mode, by default.

---

### Virtual Management Ethernet Interface

---

VCF configuration can be managed remotely using a global management interface called the virtual management Ethernet (VME) interface. The VME interface is a logical interface representing all of the out-of-band management ports on the member devices. When you connect to the VCF using the VME interface's IP address, the connection is always redirected to the device acting in the master Routing Engine role.

A VME interface should always be used to configure a VCF. The VME interface is not tied to a device, so it can always be used to log in to the VCF even after the master Routing Engine changes.

We strongly recommend cabling the management port on all spine devices to the network to ensure that you always have a direct connection to the master Routing Engine through the VME interface, regardless of which spine device assumes the master Routing Engine role. The management ports on leaf devices can also be used by the VME interface to access the VCF, so you can also cable leaf device management ports to the network, if desired.

---

### Virtual Chassis Fabric Port Link Aggregation Group Bundles

---

You can increase the bandwidth on links configured as VCPs within a VCF between two devices by configuring multiple same-speed links between two devices into VCPs. If, for instance, you configure two 40-Gbps QSFP+ links that are connecting the same devices in a VCF into VCPs, the two VCP links form one LAG bundle with two member links and 80-Gbps of total available bandwidth.

A VCP LAG bundle provides more bandwidth than a single VCP link can provide. A VCP LAG bundle also improves performance by load-sharing traffic across links within the bundle, and provides redundancy because traffic can be forwarded across another member link in the VCP LAG bundle when one VCP link fails.

VCP LAG bundling occurs automatically when same-speed VCP links are configured between two devices. No user configuration is required. VCP LAG bundling works only on same-speed VCP links; 10-Gbps and 40-Gbps links cannot be in the same VCP LAG bundle.

### Virtual Chassis Fabric License Requirements

---

A feature license is required to configure a VCF. The VCF feature license is an independent feature license; the enhanced feature licenses (EFLs) or advanced feature licenses (AFLs) that must be purchased to enable some features on some Juniper switches cannot be purchased to enable VCF.

For a VCF deployment, two license keys are recommended for redundancy—one for the device in the master Routing Engine role and the other for the device in the backup Routing Engine role.

To purchase a feature license for VCF, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with the feature license files and license keys. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show virtual-chassis** command.

### Hardware Requirements for a Virtual Chassis Fabric

---

A VCF can contain up to four devices configured as spines and up to thirty-two total devices.

All spine devices must be QFX5100 devices. We recommend optimizing the performance of your VCF by also configuring QFX5100 devices as your leaf devices. A non-mixed VCF has the highest port density and feature support for a VCF in addition to supporting more spine devices. Nevertheless, you can configure any combination of QFX5100, QFX3600, QFX3500, or EX4300 devices into leaf devices within your VCF.

You can configure any QFX5100 device as a spine device. In the most common VCF configurations, the QFX5100-24Q devices are used as spine devices and QFX5100-48S devices are used as leaf devices.

### Software Requirements in a Virtual Chassis Fabric

---

All devices in a VCF must be running the same version of Junos OS software that supports VCF. VCF is initially supported for twenty total member devices in Junos OS Release 13.2X51-D20 for QFX5100, QFX3600, QFX3500, and EX4300 devices.

In Junos OS Release 14.1X53-D15, support for thirty-two member devices in a VCF was introduced.

The devices in the VCF must be using the version of software for standalone switches.

The flex software bundle is supported on non-mixed VCFs using QFX5100 member switches only. You cannot use the flex software bundle in mixed VCFs. The flex software bundle is the software that includes "jinstall-qfx-5-flex" text in the filename when it is downloaded from the Software Center.

We recommend configuring a device to the Junos OS release running on the VCF before interconnecting it into the VCF. For additional information on VCF software upgrades, see "[Understanding Software Upgrades in a Virtual Chassis Fabric](#)" on page 5258.

For information on software upgrade options for an operational VCF, see [“Understanding Software Upgrades in a Virtual Chassis Fabric” on page 5258](#).

**Related  
Documentation**

- [Network Configuration Example: MetaFabric™ Architecture 1.1: Configuring Virtual Chassis Fabric and Network Director 1.6](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- [Understanding Virtual Chassis Fabric Configuration on page 5251](#)
- [Understanding Software Upgrades in a Virtual Chassis Fabric on page 5258](#)
- [Virtual Chassis Fabric Overview on page 5241](#)

## Understanding Virtual Chassis Fabric Configuration

This topic describes the configuration options available for your Virtual Chassis Fabric (VCF).

This topic covers:

- [Virtual Chassis Fabric Setup on page 5251](#)
- [Configuration File Management in a VCF on page 5253](#)
- [Logging into a Virtual Chassis Fabric on page 5253](#)
- [Understanding Interface Numbering on page 5253](#)

### Virtual Chassis Fabric Setup

---

You must setup your VCF using one of the following options:

- [Autoprovisioned Virtual Chassis Fabric Configuration on page 5251](#)
- [Preprovisioned Virtual Chassis Fabric Configuration on page 5252](#)
- [Nonprovisioned Virtual Chassis Fabric Configuration on page 5252](#)

#### ***Autoprovisioned Virtual Chassis Fabric Configuration***

Autoprovisioned configuration allows you to “plug and play” leaf devices into a VCF after minimal initial configuration.

The minimal configuration requirements for autop provisioning a VCF include setting the configuration mode to autoprovisioned and explicitly identifying the spine devices in your VCF by serial number. After this minimal configuration is complete, all supported devices—supported devices are either devices that have been zeroized or devices in factory default mode that have never been configured into a Virtual Chassis or VCF—are automatically added to the VCF as leaf devices when they are cabled to spine devices using supported 10-Gbps SFP+ ports or 40-Gbps QSFP+ ports. The Virtual Chassis ports (VCPs) are created automatically. Other parameters such as fabric and mixed mode are automatically detected and set.

A spine device in an autoprovisioned configuration should be configured into fabric mode before being interconnected into a VCF. A spine device in an autoprovisioned VCF must

also have the same mixed mode setting as other member devices in the VCF. You should configure your spine device into fabric mode and, if necessary, mixed mode before interconnecting it into the VCF.

A leaf device in an autoprovisioned configuration is rebooted to complete the fabric mode configuration when it is interconnected into a VCF without being set into fabric mode. The leaf device is also rebooted if the device needs to be configured into or out of mixed mode to participate in the VCF. You can avoid the downtime that accompanies the reboot of the leaf device by setting the leaf device into fabric mode and into or out of mixed mode before interconnecting it into the VCF.

### ***Preprovisioned Virtual Chassis Fabric Configuration***

In a preprovisioned configuration, you deterministically control the devices in your VCF by associating each device's serial number to a member ID and role.

The advantage of configuring a VCF using a preprovisioned configuration is that you can explicitly control which devices are added to your VCF, and in what roles. VCF configuration, notably, occurs automatically when two devices that have been configured into fabric mode (and mixed mode, if applicable) are interconnected by a supported 10-Gbps SFP+ port or a 40-Gbps QSFP+ port after the preprovisioned configuration is defined.

The disadvantage of using a preprovisioned configuration is that the configuration process is more manual than the autoprovisioned configuration process.

### ***Nonprovisioned Virtual Chassis Fabric Configuration***



**CAUTION:** We discourage nonprovisioned VCF configuration. You can configure all aspects of a VCF using autoprovisioned or preprovisioned configuration. Nonprovisioned VCF configuration should only be used by VCF experts in specialized scenarios.

A nonprovisioned VCF is the default method for creating a VCF; it is the configuration mode used when a VCF has not been configured into autoprovisioned or preprovisioned mode.

In a nonprovisioned VCF, member roles are determined by a mastership election algorithm. The first value checked by the mastership election algorithm is the mastership priority value. The devices with the highest mastership priority values assume the Routing Engine role, which is used by the spine devices in a VCF. All other devices assume the linecard role.

If two or more devices have the same mastership priority value and are candidates for the Routing Engine role, the mastership election algorithm uses other parameters to determine which device is elected as the Routing Engine. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 5078](#).

The default mastership priority value for all devices is 128. You should always configure your spine devices with the highest mastership priority to ensure a spine device assumes the Routing Engine role.

In a nonprovisioned VCF, you must manually configure every VCP.

### Configuration File Management in a VCF

You configure a VCF by logging onto the master Routing Engine and making configuration changes. See the next section for information on logging into a VCF.

The configuration file that is modified when you are on the master Routing Engine is automatically shared with all other devices in the VCF when it is committed. Each device stores its own copy of the configuration file.

### Logging into a Virtual Chassis Fabric

The recommended method of logging into a VCF is through the use of a Virtual Management Ethernet (VME) interface. The VME interface is a logical interface representing all of the out-of-band management ports on the member devices. When you connect to the VCF configuration using the VME interface's IP address, the connection is always redirected to the management port on device in the master Routing Engine role. The VME interface is not tied to a device, so it can always be used to log in to the VCF even after the master Routing Engine changes. We recommend cabling the management ports—an *me* or *em* interface—on each spine device in your VCF to support the VME interface.

If you log in to the console port of any member device in a VCF, your session is automatically redirected to the device acting in the master Routing Engine role.

### Understanding Interface Numbering

Interfaces in Junos OS are specified as follows:

*type-fpc/pic/port*

A VCF applies this convention as follows:

- *type*—The interface type.
- *fpc*—Flexible PIC Concentrator. In a VCF, the *fpc* is the member ID of the switch. For instance, the *fpc* of member 16 in the VCF is 16.
- *pic*—the number of the PIC (Physical Interface Card) on the member device.
- *port*—the port number.

For more detailed information on interface numbering, see *Understanding Interface Naming Conventions*.

#### Related Documentation

- [Network Configuration Example: MetaFabric™ Architecture 1.1: Configuring Virtual Chassis Fabric and Network Director 1.6](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- [Configuring a Nonprovisioned Virtual Chassis Fabric on page 5267](#)

## Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric

This topic describes the requirements for a mixed Virtual Chassis or a mixed Virtual Chassis Fabric (VCF).

A mixed Virtual Chassis includes two or more types of EX Series switches, two or more types of QFX Series switches, or a mix of EX and QFX Series switches.

A mixed VCF is any VCF that includes two or more types of member switches. Because a VCF must use a QFX5100 switch as a spine device, a mixed VCF is any VCF that includes EX4300, QFX3500, or QFX3600 member switches in addition to the required QFX5100 switches.



**NOTE:** The optimal VCF topology is to use QFX5100 devices only. A VCF composed entirely of QFX5100 devices supports the largest breadth of features at the highest scalability while also supporting the highest number of high-speed interfaces.

This topic covers:

- [Virtual Chassis Fabric Summary on page 5254](#)
- [Understanding Mixed Virtual Chassis Fabric on page 5255](#)
- [Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches on page 5255](#)
- [Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches on page 5256](#)
- [Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis on page 5257](#)
- [Understanding Mixed EX4300 and EX4600 Virtual Chassis on page 5257](#)
- [Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis on page 5257](#)

### Virtual Chassis Fabric Summary

---

[Table 565 on page 5074](#) provides a high-level overview of the permitted hardware allowed in the routing engine and line card roles of a mixed and a non-mixed VCF. The table also includes license requirements and supported configuration methods.

Table 584: Virtual Chassis Fabric Summary

| Category  | Allowed Routing Engines | Allowed Line Cards                      | License Requirement                                                               | Configuration Methods                                                    |
|-----------|-------------------------|-----------------------------------------|-----------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| Non-mixed | QFX5100                 | QFX5100                                 | Yes (on two QFX5100 switches operating in master and backup Routing Engine roles) | Autoprovisioning<br>Preprovisioning<br>Nonprovisioning (not recommended) |
| Mixed     | QFX5100                 | QFX5100<br>QFX3600<br>QFX3500<br>EX4300 | Yes (on two QFX5100 switches operating in master and backup Routing Engine roles) | Autoprovisioning<br>Preprovisioning<br>Nonprovisioning (not recommended) |

### Understanding Mixed Virtual Chassis Fabric

A VCF must use a QFX5100 switch in the spine role. A mixed VCF is, therefore, any VCF that includes EX4300, QFX3500, or QFX3600 member switches in addition to the required QFX5100 switch.

The optimal method of configuring a VCF is to use QFX5100 devices only. A non-mixed VCF composed entirely of QFX5100 devices supports the largest breadth of features at the highest scalability while also supporting the highest number of high-speed interfaces. You can, however, also configure a mixed VCF.

If you use QFX3600, QFX3500, or EX4300 devices as leaf devices in your VCF, you must configure all devices in your VCF into mixed mode. If you are turning a non-mixed VCF into a mixed VCF, you have to reboot the VCF to change the mixed mode setting.

### Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches

[Table 566 on page 5076](#) provides a high-level overview of the permitted hardware allowed in the routing engine and line card roles of a mixed and a non-mixed Virtual Chassis for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 switches. The table also includes license requirements and supported configuration methods.

Table 585: Virtual Chassis Summary

| Category  | Allowed Routing Engines | Allowed Line Cards                      | License Requirement | Configuration Methods              |
|-----------|-------------------------|-----------------------------------------|---------------------|------------------------------------|
| Non-mixed | QFX5100                 | QFX5100                                 | No                  | Nonprovisioning<br>Preprovisioning |
|           | QFX3600<br>QFX3500      | QFX3600<br>QFX3500                      | No                  | Nonprovisioning<br>Preprovisioning |
|           | EX4600                  | EX4600                                  | No                  | Nonprovisioning<br>Preprovisioning |
|           | EX4300                  | EX4300                                  | No                  | Nonprovisioning<br>Preprovisioning |
| Mixed     | QFX5100                 | QFX5100<br>QFX3600<br>QFX3500<br>EX4300 | No                  | Nonprovisioning<br>Preprovisioning |
|           | QFX3600<br>QFX3500      | QFX3600<br>QFX3500<br>EX4300            | No                  | Nonprovisioning<br>Preprovisioning |
|           | EX4600                  | EX4600<br>EX4300                        | No                  | Nonprovisioning<br>Preprovisioning |

#### Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches

In a mixed Virtual Chassis, the switch in the master Routing Engine role determines which switches are supported in the line card role of the mixed Virtual Chassis.

When a mixed Virtual Chassis is using a QFX5100 switch in the master Routing Engine role, you can use QFX5100, QFX3600, QFX3500, or EX4300 switches in the line card role.

When a mixed Virtual Chassis is using a QFX3600 or QFX3500 switch in the master Routing Engine role, you can use QFX3600, QFX3500, or EX4300 switches in the line card role.

In a mixed EX4300 and EX4600 Virtual Chassis, an EX4600 switch automatically assumes the Routing Engine role.

EX4600 switches can only be in a mixed Virtual Chassis with EX4300 switches. EX4600 switches cannot be in a mixed Virtual Chassis with QFX5100, QFX3600, or QFX3500 switches.

We recommend always configuring the same type of switch into the master and backup Routing Engine role, to ensure that the switch operating in the master role remains the same type of switch in the event of a switchover.



In most mixed Virtual Chassis, you must configure your Virtual Chassis to ensure a switch that supports the master Routing Engine assumes the master Routing Engine role. Without user configuration, any switch—with the exception of the EX4300 switch, which can never assume the master or backup Routing Engine role in a mixed Virtual Chassis or VCF—can assume the master or backup Routing Engine role.

### **Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis**

Up to ten EX4300 switches, QFX3500 switches, QFX3600 switches, and QFX5100 switches can be interconnected using Virtual Chassis ports (VCPs) to form a mixed or non-mixed Virtual Chassis. The mixed Virtual Chassis supports up to ten member switches regardless of the switches that compose the mixed Virtual Chassis.

EX4300 switches can also be interconnected into a mixed Virtual Chassis with EX4600 switches. See the following section for information on mixed EX4300 and EX4600 Virtual Chassis.

### **Understanding Mixed EX4300 and EX4600 Virtual Chassis**

EX4300 switches and EX4600 switches can be interconnected into the same Virtual Chassis. An EX4600 switch automatically assumes the master Routing Engine role in a mixed EX4300 and EX4600 Virtual Chassis, since EX4300 switches cannot assume the Routing Engine role in a mixed Virtual Chassis. EX4600 switches cannot be in a mixed Virtual Chassis with any other type of switch.

The mixed Virtual Chassis supports up to ten member switches.

### **Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis**

EX4200 switches, EX4500 switches, and EX4550 switches can be interconnected into the same Virtual Chassis to form a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, mixed EX4500 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, and EX4550 Virtual Chassis. The mixed Virtual Chassis supports up to 10 member switches regardless of whether the switches are EX4200 switches, EX4500 switches, or EX4550 switches. Any model of EX4200, EX4500, or EX4550 switch can be interconnected into the same mixed Virtual Chassis. The master election process that decides member switch roles in a mixed Virtual Chassis is identical to the master election process in a non-mixed Virtual Chassis, so any member switch in a mixed Virtual Chassis can assume the master, backup, or linecard role.

EX4200 switches, EX4500 switches, and EX4550 switches cannot be interconnected into a Virtual Chassis with any other switches.

#### **Related Documentation**

- [Virtual Chassis Fabric Overview on page 5241](#)
- [Understanding QFX Series Virtual Chassis](#)
- [EX Series Virtual Chassis Overview](#)
- [Understanding Virtual Chassis Fabric Components on page 5243](#)
- [Understanding QFX Series Virtual Chassis Components](#)
- [Understanding EX Series Virtual Chassis Components on page 5067](#)

## Understanding Traffic Flow Through a Virtual Chassis Fabric

A Virtual Chassis Fabric (VCF) forwards unicast traffic using a smart trunking algorithm that sends all traffic across multiple paths based on end-to-end available bandwidth. The smart trunking algorithm avoids unnecessary congestion due to improper traffic allocation while optimizing fabric bandwidth utilization because traffic is forwarded through the VCF relative to available bandwidth.

The smart trunking algorithm works by considering the overall available path bandwidth of each path in the VCF when making traffic-forwarding decisions, and then forwarding traffic across the paths relative to available path bandwidth. If a VCF with two spine devices, for instance, has one path from leaf device 1 to leaf device 4 that contains two 40-Gbps QSFP+ links and a second path from leaf device 1 to leaf device 4 that contains two 10-Gbps SFP+ links, the algorithm tries to balance traffic sent on the paths so that four times more packets are sent on the first path with 40 Gbps of available bandwidth across the entire path than are sent on the second path with 10 Gbps of total bandwidth.

You can optimize how traffic is forwarded through the VCF by adding spine devices to maximize the number of available paths between all leaf devices, and by using as many 40-Gbps QSFP+ interfaces as Virtual Chassis ports (VCPs) as possible.

VCF also supports adaptive load balancing (ALB). ALB enables the VCF trunking algorithm to use dynamic load information on interfaces and traffic queues to make forwarding decisions within the VCF. When ALB is implemented using flowlets, traffic flows that enter the VCF are spliced into smaller flows—flowlets—and individually forwarded across the VCF to the same destination device over different paths when the inactivity time between packet bursts on the sending interface exceeds the user-configurable inactivity interval. When ALB is implemented using per-packet mode, the sending interface actively monitors all paths available between two member devices and forwards traffic through the VCF using the best available path at the moment.

Implementing ALB using flowlets is effective in environments that periodically experience extremely large traffic flows—*elephant flows*—that are substantially larger than the majority of other traffic flowing through the VCF. The VCF is better able to manage the elephant flows by splicing them into smaller flowlets using ALB.

ALB is supported on a non-mixed VCF composed entirely of QFX5100 switches only. You should enable ALB using flowlets in non-mixed VCFs in environments where a small number of traffic flows are disproportionately larger than the majority of the other traffic flows.

**Related Documentation** • [Understanding Virtual Chassis Fabric Components on page 5243](#)

## Understanding Software Upgrades in a Virtual Chassis Fabric

This topic provides an overview of software upgrades on Virtual Chassis Fabric (VCF).

It contains the following sections:

- [Virtual Chassis Fabric Software Basics on page 5259](#)
- [Nonstop Software Upgrade \(NSSU\) on page 5259](#)
- [Automatic Software Update on page 5259](#)
- [Traditional Software Upgrade on page 5259](#)

---

### Virtual Chassis Fabric Software Basics

---

VCF is initially supported in Junos OS Release 13.2X51-D20. All devices in a VCF must be running the same version of Junos OS that supports VCF.

At initial VCF configuration, you should configure all devices to the same Junos OS release before interconnecting them into a VCF.

When you are adding a device to an existing VCF, you should update the Junos OS release on the new device to the Junos OS release running in the VCF before interconnecting it into the VCF. Updating the Junos OS on the device before interconnecting it helps ensure the device is gracefully added to the VCF, without the downtime that is required to reboot the device after an automatic software update or the troubleshooting that is required if the device isn't added to the VCF due to mismatched software releases.

Before you interconnect a device into a VCF, you should upgrade the software on the device being added to the VCF to the version of Junos OS running on the VCF.

---

### Nonstop Software Upgrade (NSSU)

---

Nonstop software upgrade (NSSU) enables you to upgrade the software running on all member devices in a VCF with minimal network traffic disruption during the upgrade.

NSSU upgrades the software on each device individually while all other devices continue normal operations.

For additional information on NSSU in a VCF, see *Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric*.

---

### Automatic Software Update

---

Automatic software update automatically upgrades the Junos OS running on a device joining a VCF to the version of Junos OS running on the VCF at the moment the new device is cabled into the VCF.

Automatic software update is enabled using the **set virtual-chassis auto-sw-update** statement.

---

### Traditional Software Upgrade

---

You can upgrade software on a VCF using the traditional method of upgrading software for Junos OS by logging onto the master Routing Engine and using the **request system software add** command to initiate the upgrade on a non-mixed VCF or the **request system software add set [package-name package-name ...]** to initiate the upgrade on a mixed VCF, where *package-name* is the path to an image for one device family.

When you upgrade Junos OS on a VCF using the traditional software upgrade, the entire system is down until the upgrade is complete.

**Related  
Documentation**

- *Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade*
- [Upgrading Software for a Virtual Chassis Fabric on page 5278](#)
- [Adding a Device to a Virtual Chassis Fabric on page 5270](#)

## CHAPTER 89

# Configuration

- [Configuration Tasks on page 5261](#)
- [Configuration Statements on page 5280](#)

### Configuration Tasks

---

- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- [Configuring a Nonprovisioned Virtual Chassis Fabric on page 5267](#)
- [Adding a Device to a Virtual Chassis Fabric on page 5270](#)
- [Removing a Device From a Virtual Chassis Fabric on page 5277](#)
- [Upgrading Software for a Virtual Chassis Fabric on page 5278](#)

### Autoprovisioning a Virtual Chassis Fabric

Autoprovisioning a Virtual Chassis Fabric (VCF) enables you to “plug and play” devices into your VCF after minimal initial configuration.

Update all devices to the same version of Junos OS that supports VCF. See *Upgrading Software* or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To configure a VCF using autoprovisioning:

1. Make a list of the serial numbers of all the spine devices in the VCF. The spine devices must be QFX5100 devices. You can configure up to four spine devices in a VCF. You can get the device's serial number in the **show virtual-chassis** output or by following the instructions in *Locating the Serial Number on a QFX5100 Device or Component*.
2. Configure each device into fabric mode. If needed, configure the devices into mixed mode.

Configure the device to reboot as part of the procedure to complete this configuration step.

Configure mixed mode if your VCF includes QFX3600, QFX3500, or EX4300 devices as leaf devices.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```

If you are configuring a VCF composed of two or more types of devices:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** A spine device whose fabric or mixed mode setting is improperly set cannot join a VCF. You can check the mode settings by using the `show virtual-chassis mode` command.

We recommend that you set the fabric and mixed mode settings before you interconnect your spine devices into the VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

You can, however, use the `request virtual-chassis mode fabric local` or `request virtual-chassis mode mixed local` commands to set a spine device into fabric or mixed mode after interconnecting your VCF.

The fabric and mixed mode settings are automatically updated for a leaf device when it is interconnected into an autoprovisioned VCF. If the fabric or mixed mode settings are changed when a leaf device is interconnected into a VCF, the leaf device reboots before joining the VCF.

---

3. When the reboot is complete, log in to one of the spine devices in your VCF.

4. Set the configuration mode to autoprovisioned:

```
[edit]
```

```
user@device# set virtual-chassis auto-provisioned
```

5. Configure your spine devices into the Routing Engine role:

```
[edit virtual-chassis]
```

```
user@device# set member member-id serial-number serial-number role routing-engine
```

For instance, to configure the four spine devices with the serial numbers “SERIALNUMB00”, “SERIALNUMB01”, “SERIALNUMB02”, and “SERIALNUMB03” into the Routing Engine role as members 0 through 3:

```
[edit virtual-chassis]
```

```
user@device# set member 0 serial-number SERIALNUMB00 role routing-engine
```

```
user@device# set member 1 serial-number SERIALNUMB01 role routing-engine
```

```
user@device# set member 2 serial-number SERIALNUMB02 role routing-engine
```

```
user@device# set member 3 serial-number SERIALNUMB03 role routing-engine
```

6. (Recommended) Configure a virtual management Ethernet (VME) interface for management of the VCF configuration:

```
[edit]
```

```
user@device# set interfaces vme unit 0 family inet address /ip-address/mask/
```



**NOTE:** A VME accesses the device in the master Routing Engine role using a management port, so cable management port em0 or em1 on each spine device in your VCF so the VME is available regardless of which spine device assumes the master Routing Engine role. See *Connecting a QFX Series Device to a Management Console*

7. Commit the configuration:

```
user@device# commit
```

8. Cable your VCF.

After your autoprovisioned VCF configuration is committed, you can cable any EX4300, QFX3500, QFX3600, or QFX5100 device that is zeroized or that has never been configured to a spine device using a supported SFP+ or QSFP+ interface. The device that is zeroized or in factory-default mode is added to the VCF as a leaf device. All VCPs are configured as part of this process.



**NOTE:** Mixed mode and fabric mode are checked and, if needed, set automatically on the device as part of this process. If the mixed or fabric mode has to be changed to become part of the VCF, the device reboots. The device participates in the VCF with no further user intervention after this reboot is complete.



**NOTE:** Automatic VCP conversion only works when the interfaces on both ends of the link are not configured into VCPs.

The 40-Gbps QSFP+ interfaces on EX4300 switches are configured as VCPs, by default. You must, therefore, delete the VCP on the 40-Gbps QSFP+ interface using the `request virtual-chassis vc-port delete` command before interconnecting it into the VCF in order for the link to be converted into a VCP. You can also manually configure the link into a VCP using the `request virtual-chassis vc-port set` command.

The device joins the VCF immediately without a reboot if the mixed or fabric mode setting does not need to be changed.

9. Install the VCF feature licenses.

For a VCF deployment, two license keys are recommended for redundancy—one for the device in the master role and the other for the device in the backup role.

To purchase a feature license for VCF, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with the feature license files and license keys. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the `show virtual-chassis` command.

After obtaining the licenses, follow the instructions in *Generating License Keys*.

**Related Documentation**

- [Understanding Virtual Chassis Fabric Configuration on page 5251](#)
- [Adding a Device to a Virtual Chassis Fabric on page 5270](#)
- [Removing a Device From a Virtual Chassis Fabric on page 5277](#)
- [Understanding Virtual Chassis Fabric Components on page 5243](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)

## Preprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric (VCF) configuration allows you to assign the member ID and role for each device in the VCF.

Update all devices to the same version of Junos OS that supports VCF. See *Upgrading Software* or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To preprovision a VCF:

1. Make a list of the serial numbers of all the devices to be connected in the VCF. You can get a device's serial number in the **show virtual-chassis** output or by following the instructions in *Locating the Serial Number on a QFX5100 Device or Component*, *Locating the Serial Number on a QFX3600 or QFX3600-I Device or Component*, *Locating the Serial Number on a QFX3500 Device or Component*, or *Locating the Serial Number on an EX4300 Switch or Component*.
2. Decide the desired role (**routing-engine** or **line-card**) for each device.

In a VCF, you configure up to four QFX5100 devices into the Routing Engine role as spine devices. All other devices are configured into the linecard role as leaf devices.



**NOTE:** Only two devices can simultaneously operate in the Routing Engine role within a VCF. A VCF, however, supports up to four spine devices. In scenarios where a preprovisioned VCF has three or more spine devices, the devices that are not operating in the Routing Engine role operate in the linecard role.

A spine device operating in the linecard role assumes the Routing Engine role when an active Routing Engine fails.

3. Configure each individual device into fabric mode. If needed, configure the devices into mixed mode.

Reboot each device to complete this configuration step.

Mixed mode must be configured if your VCF includes QFX3500, QFX3600, or EX4300 devices as leaf nodes.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```



If you are configuring a VCF that includes EX4300, QFX3500, or QFX3600 devices as leaf nodes:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** A device whose fabric or mixed mode setting is improperly set cannot join a VCF. You can check the mode settings using the `show virtual-chassis mode` command.

We recommend that you set the fabric and mixed mode before you interconnect your devices into a VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

You can, however, use the `request virtual-chassis mode fabric local` or `request virtual-chassis mode mixed local` commands to set a device into fabric or mixed mode after interconnecting your VCF.

4. Log in to one of your spine devices after the reboot has completed.

5. Specify the preprovisioned configuration mode:

```
[edit virtual-chassis]
user@device# set preprovisioned
```

6. Associate a member ID with a serial number for each device in your VCF, and configure the role for each device:

```
[edit virtual-chassis]
user@device# set member member-id serial-number serial-number role (line-card |
routing-engine)
```

Configure your spine devices into the Routing Engine role. You must use QFX5100 devices as your spine devices.

Configure your leaf devices into the linecard role.

For instance, if you wanted to preprovision a VCF with twenty member devices:

```
[edit virtual-chassis]
user@device# set member 0 serial-number SERIALNUMB00 role routing-engine
user@device# set member 1 serial-number SERIALNUMB01 role routing-engine
user@device# set member 2 serial-number SERIALNUMB02 role routing-engine
user@device# set member 3 serial-number SERIALNUMB03 role routing-engine
user@device# set member 4 serial-number SERIALNUMB04 role line-card
user@device# set member 5 serial-number SERIALNUMB05 role line-card
user@device# set member 6 serial-number SERIALNUMB06 role line-card
user@device# set member 7 serial-number SERIALNUMB07 role line-card
user@device# set member 8 serial-number SERIALNUMB08 role line-card
user@device# set member 9 serial-number SERIALNUMB09 role line-card
user@device# set member 10 serial-number SERIALNUMB10 role line-card
user@device# set member 11 serial-number SERIALNUMB11 role line-card
user@device# set member 12 serial-number SERIALNUMB12 role line-card
user@device# set member 13 serial-number SERIALNUMB13 role line-card
user@device# set member 14 serial-number SERIALNUMB14 role line-card
```

```

user@device# set member 15 serial-number SERIALNUMB15 role line-card
user@device# set member 16 serial-number SERIALNUMB16 role line-card
user@device# set member 17 serial-number SERIALNUMB17 role line-card
user@device# set member 18 serial-number SERIALNUMB18 role line-card
user@device# set member 19 serial-number SERIALNUMB19 role line-card

```

7. (Recommended) Configure a virtual management Ethernet (VME) interface for management of the VCF configuration:

```
[edit]
```

```
user@device# set interfaces vme unit 0 family inet address /ip-address/mask/
```



**NOTE:** A VME accesses the device in the master Routing Engine role using a management port, so cable management port em0 or em1 on each spine device in your VCF so the VME is available regardless of which spine device assumes the master Routing Engine role. See *Connecting a QFX Series Device to a Management Console*

8. Commit the configuration:

```
user@device# commit
```

9. Interconnect the spine device that you configured in the previous steps to all leaf devices by using the supported SFP+ and QSFP+ interfaces.



**NOTE:** The automatic Virtual Chassis port (VCP) conversion feature is enabled and automatically configures SFP+ and QSFP+ interfaces into VCPs when the VCF configuration mode is set to preprovisioned. You do not need to manually configure VCPs.

If you want to configure an SFP+ or QSFP+ interface into a network interface, disable LLDP on that interface. See *Configuring LLDP*.



**NOTE:** Automatic VCP conversion only works when the interfaces on both ends of the link are not configured into VCPs.

The 40-Gbps QSFP+ interfaces on EX4300 switches are configured as VCPs, by default. You must, therefore, delete the VCP on the 40-Gbps QSFP+ interface using the `request virtual-chassis vc-port delete` command before interconnecting it into the VCF in order for the link to be converted into a VCP. You can also manually configure the link into a VCP using the `request virtual-chassis vc-port set` command.

10. Interconnect all other spine devices to all other leaf devices using the supported SFP+ and QSFP+ interfaces.
11. Install the VCF feature licenses.

For a VCF deployment, two license keys are recommended for redundancy—one for the device in the master role and the other for the device in the backup role.

To purchase a feature license for VCF, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with the feature license files and license keys. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show virtual-chassis** command.

After obtaining the licenses, follow the instructions in *Generating License Keys*.

#### Related Documentation

- [Understanding Virtual Chassis Fabric Configuration on page 5251](#)
- [Adding a Device to a Virtual Chassis Fabric on page 5270](#)
- [Removing a Device From a Virtual Chassis Fabric on page 5277](#)
- [Understanding Virtual Chassis Fabric Components on page 5243](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)

## Configuring a Nonprovisioned Virtual Chassis Fabric



**CAUTION:** Configure your Virtual Chassis Fabric (VCF) using autoprovisioning or preprovisioning unless you have a compelling reason to use nonprovisioned configuration. You can configure all aspects of a VCF using autoprovisioned or preprovisioned configuration. The process for autoprovisioning your VCF is described in “[Autoprovisioning a Virtual Chassis Fabric](#)” on page 5261 and the process for preprovisioning your VCF is described in “[Preprovisioning a Virtual Chassis Fabric](#)” on page 5264.

Nonprovisioned VCF configuration is highly discouraged. Nonprovisioned VCF configuration should only be used by VCF experts in specialized scenarios.

A nonprovisioned VCF is the configuration mode used when a VCF has not been configured into autoprovisioned or preprovisioned mode.

In a nonprovisioned VCF, you configure the device roles by setting the mastership priority value of each device. If no mastership priority values are set, a master election algorithm process runs and selects the role for each device.

You must manually configure all Virtual Chassis ports (VCPs) in a nonprovisioned VCF. The automatic VCP conversion feature, which automatically configures supported 10-Gbps SFP+ links and 40-Gbps QSFP+ links into VCPs on autoprovisioned and preprovisioned VCFs, is not supported on nonprovisioned VCFs.

Update all devices to the same version of Junos OS that supports VCF. See *Upgrading Software or Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To configure a nonprovisioned VCF:

1. Power on the devices.
2. Configure each individual device into fabric mode. If needed, configure the devices into mixed mode.

Reboot each device to complete this configuration step.

A VCF must have QFX5100 devices in the spine role, and operates most efficiently when the leaf nodes are also QFX5100 devices. Mixed mode must be configured if your VCF also includes at least one QFX3600, QFX3500, or EX4300 device in the leaf role.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```

If you are configuring a VCF using at least one QFX3600, QFX3500, or EX4300 device as a leaf device:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** A device whose fabric or mixed mode setting is improperly set cannot join a VCF. You can check the mode settings using the `show virtual-chassis mode` command.

We recommend setting the fabric and mixed mode settings before interconnecting your devices into a VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

We strongly recommend configuring the mixed and fabric settings before you interconnect a device into a VCF. You can, however, use the `request virtual-chassis mode fabric local` or `request virtual-chassis mode mixed local` commands to set a device into fabric or mixed mode after you have interconnected your VCF.

3. After the device reboots are complete, cable your spine devices to your leaf devices using supported SFP+ and QSFP+ interfaces.
4. (Recommended) Configure a virtual management Ethernet (VME) interface for management of the VCF configuration:

[edit]

```
user@device# set interfaces vme unit 0 family inet address /ip-address/mask/
```



**NOTE:** A VME accesses the device in the master Routing Engine role using a management port, so cable management port em0 or em1 on each spine device in your VCF so the VME is available regardless of which spine device assumes the master Routing Engine role. See *Connecting a QFX Series Device to a Management Console*

5. Configure the desired SFP+ and QSFP+ interfaces into Virtual Chassis ports (VCPs):

```
user@device> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
user@device> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
```

The `show virtual-chassis vc-port` must be issued on the ports at both ends of the link in order for that link to be configured into a VCP.

6. Enter the `show virtual-chassis` command to confirm that the VCPs are operational and to learn the member ID of each member device in your VCF.

If you want to change the member ID that has been assigned to a member device, use the `request virtual-chassis renumber` command.

7. (Optional) Configure the mastership priority for each member device:

```
[edit virtual-chassis]
user@device# set member member-id mastership-priority number
```

In a nonprovisioned VCF, member roles are determined by a mastership election algorithm. The first value checked by the mastership election algorithm is the mastership priority value. The two QFX5100 devices with the highest mastership priority values assume the master and backup Routing Engine role, which must be used by the spine devices in a VCF. All other devices assume the linecard role.

QFX5100 devices assume the Routing Engine role, regardless of mastership priority settings. QFX5100 devices can also assume the linecard role.

QFX3600, QFX3500, and EX4300 devices always assume the linecard role in a VCF, regardless of the mastership priority settings.



**NOTE:** A spine device that isn't selected as master or backup Routing Engine assumes the linecard role. The spine devices should still be configured with a higher mastership priority value than the leaf devices to assure a spine device assumes the Routing Engine role when the master or backup Routing Engine fails.

If two or more devices have the same mastership priority value and are candidates for the Routing Engine role, the mastership election algorithm uses other parameters to determine which device is elected into the Routing Engine role. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 5078](#).

A device with a mastership priority of 0 never assumes the master or backup role.

For instance, to configure the mastership priority for member devices 0 through 19 in your VCF.

```
[edit virtual-chassis]
```

```
user@device# set member 0 mastership-priority 255
user@device# set member 1 mastership-priority 255
user@device# set member 2 mastership-priority 255
user@device# set member 3 mastership-priority 255
user@device# set member 4 mastership-priority 100
user@device# set member 5 mastership-priority 95
user@device# set member 6 mastership-priority 90
user@device# set member 7 mastership-priority 85
user@device# set member 8 mastership-priority 80
user@device# set member 9 mastership-priority 75
user@device# set member 10 mastership-priority 70
user@device# set member 11 mastership-priority 65
user@device# set member 12 mastership-priority 60
user@device# set member 13 mastership-priority 55
user@device# set member 14 mastership-priority 50
user@device# set member 15 mastership-priority 45
user@device# set member 16 mastership-priority 40
user@device# set member 17 mastership-priority 35
user@device# set member 18 mastership-priority 30
user@device# set member 19 mastership-priority 25
```

8. Install the VCF feature licenses.

For a VCF deployment, two license keys are recommended for redundancy—one for the device in the master role and the other for the device in the backup role.

To purchase a feature license for VCF, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with the feature license files and license keys. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show virtual-chassis** command.

After obtaining the licenses, follow the instructions in *Generating License Keys*.

**Related  
Documentation**

- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- [Adding a Device to a Virtual Chassis Fabric on page 5270](#)
- [Removing a Device From a Virtual Chassis Fabric on page 5277](#)
- [Understanding Virtual Chassis Fabric Configuration on page 5251](#)
- [Understanding Virtual Chassis Fabric Components on page 5243](#)

## Adding a Device to a Virtual Chassis Fabric

This topic describes how to add a device to a Virtual Chassis Fabric (VCF).

It contains the following sections:

- [Adding a Leaf Device to an Autoprovisioned Virtual Chassis Fabric on page 5271](#)
- [Adding a Spine Device to an Autoprovisioned Virtual Chassis Fabric on page 5272](#)
- [Adding a Spine or Leaf Device to a Preprovisioned Virtual Chassis Fabric on page 5273](#)
- [Adding a Spine or Leaf Device to a Nonprovisioned Virtual Chassis Fabric on page 5275](#)

### Adding a Leaf Device to an Autoprovisioned Virtual Chassis Fabric

Update your device to the same version of Junos OS running on the devices in the VCF. See *Upgrading Software* or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To add a leaf device to an autoprovisioned VCF:

1. Log onto the device that you are adding to the VCF.
2. (Optional) Perform this step if you want to avoid the downtime associated with an extra reboot when your device is interconnected into your VCF. If you do not perform this step, the VCF auto-detects the fabric and mixed mode settings and, if needed, reboots the device as part of the process of changing these settings.

Configure the leaf device into fabric mode. Configure your device into mixed mode if your VCF includes QFX3600, QFX3500, or EX4300 devices as leaf devices..

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local
```

If you are configuring a VCF composed of two or more types of devices:

```
user@device> request virtual-chassis mode fabric mixed local
```

3. If the leaf device that you are adding to the VCF has not previously been configured, proceed to the next step.

If your device has been configured, zeroize your device and reboot:

```
user@device> request system zeroize
```

```
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (yes) yes
```



**NOTE:** You must zeroize your device if you have previously entered one or more configuration commands, including basic configuration commands.

Your device will not join the VCF if it contains any configuration until it has been zeroized.



**NOTE:** The `request virtual-chassis mode fabric local` and `request virtual-chassis mode fabric mixed local` commands are entered in operational mode, so those settings are maintained when the device is zeroized.

You cannot use other methods to set a device into factory default mode before inserting it into a VCF if it was previously configured in another Virtual Chassis or VCF. You must use **`request system zeroize`**.

For additional information on this procedure, see [“Reverting to the Default Factory Configuration for the EX Series Switch” on page 620](#) or [Reverting to the Default Factory Configuration](#).

4. (Required only if you are adding a device that turns a non-mixed VCF into a mixed VCF) Log in to the VCF and set all devices in the VCF to mixed mode. Configure all devices to reboot to complete this procedure.

```
user@device> request virtual-chassis mode mixed all-members reboot
```

The VCF experiences downtime as part of the reboot procedure.

5. Interconnect your leaf device into the existing spine devices, using at least one 10-Gbps SFP+ interface or 40-Gbps QSFP+ interface to connect to each spine device in the VCF.

An autoprovisioned VCF automatically adds a supported device in factory-default mode to the VCF when it is connected to a spine devices using a supported SFP+ or QSFP+ link. The SFP+ or QSFP+ link is automatically converted into a Virtual Chassis port (VCP) as part of this process.

No further configuration is required.

---

### Adding a Spine Device to an Autoprovisioned Virtual Chassis Fabric

---

Update your device to the same version of Junos OS running on the devices in the VCF before interconnecting it into the VCF. See *Upgrading Software* or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To add a spine device to an autoprovisioned VCF:

1. Log in to your VCF.
2. If you are replacing a spine device that is already part of the VCF, power off the spine device in the VCF.

Follow the steps in [“Removing a Device From a Virtual Chassis Fabric” on page 5277](#) to remove the device from the VCF.

3. Modify the configuration.

If your new spine device is replacing an existing spine, modify the configuration to remove the old spine.

You can skip this step if you are not replacing an existing spine device.

```
[edit virtual-chassis]
```

```
user@device# delete member member-id
```

where *member-id* is the member ID of the spine that is removed from this procedure.

Add the spine device to the configuration:

```
[edit virtual-chassis]
```

```
user@device# set member member-id serial-number serial-number role routing-engine
```

For instance, to configure a spine device with the serial number OU81234567890 as member 3:

```
[edit virtual-chassis]
```

```
user@device# set member 3 serial-number OU81234567890 role routing-engine
```

4. Commit the configuration.

```
[edit]
```

```
user@device# commit
```



5. Log in to the device that is going to be added to the VCF.
6. Configure the device into fabric mode. If needed, also configure the device into mixed mode.

Reboot the device to complete this configuration step.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```

If you are configuring a VCF composed of QFX5100 devices and at least one other type of device:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** We recommend setting the fabric and mixed mode settings before interconnecting your devices into a VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

You can, however, use the `request virtual-chassis mode fabric local` or `request virtual-chassis mode mixed local` commands to set a device into fabric or mixed mode after interconnecting your VCF.

7. (Required only if you are adding a device that turns a non-mixed VCF into a mixed VCF) Log in to the VCF and set all devices in the VCF to mixed mode. Configure all devices to reboot to complete this procedure.

```
user@device> request virtual-chassis mode mixed all-members reboot
```

The VCF experiences downtime as part of the reboot procedure.

8. After the device reboots, interconnect the new device into the VCF by cabling the device to the leaf devices in the VCF using supported SFP+ or QSFP+ interfaces.

The SFP+ or QSFP+ links are converted into VCPs automatically.

The new spine device should be operational once the cabling is complete.

### Adding a Spine or Leaf Device to a Preprovisioned Virtual Chassis Fabric

Update your device to the same version of Junos OS running on the devices in the VCF before interconnecting it into the VCF. See *Upgrading Software* or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To add a spine or leaf device to a preprovisioned VCF:

1. Log in to your VCF.
2. If you are replacing a device that is already part of the VCF, power off the device in the VCF.

Follow the steps in [“Removing a Device From a Virtual Chassis Fabric”](#) on page 5277 to remove the device from the VCF.

3. Modify the configuration.

If your new device is replacing an existing device, modify the configuration to remove the old device.

You can skip this portion of the procedure if you are not replacing an existing device.

```
[edit virtual-chassis]
user@device# delete member member-id
```

where *member-id* is the member ID of the spine that is removed from this procedure.

Add the new device to the configuration:

```
[edit virtual-chassis]
user@device# set member member-id serial-number serial-number role routing-engine
```

For instance, to configure a device with the serial number OU81234567890 into the Routine Engine role as member 3:

```
[edit virtual-chassis]
user@device# set member 3 serial-number OU81234567890 role routing-engine
```

4. Commit the configuration.

```
[edit]
user@device# commit
```

5. Log in to the device that is going to be added to the VCF.

6. Configure the device into fabric mode. If needed, also configure the device into mixed mode.

Reboot the device to complete this configuration step.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```

If you are configuring a VCF composed of two or more types of devices:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** If you are adding a QFX3600, QFX3500, or EX4300 device to a VCF that is composed entirely of QFX5100 devices, you must also log in to the VCF and set all of the devices in the VCF into mixed mode.

Log in to the VCF and enter the **request virtual-chassis mode mixed all-members reboot** command to perform this task.

The VCF reboots and incurs downtime to complete this procedure.

---



**NOTE:** We recommend that you set the fabric and mixed mode settings before you interconnect your devices into a VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

You can, however, use the `request virtual-chassis mode fabric local` or `request virtual-chassis mode mixed local` commands to recover a device that was not set into fabric or mixed mode before you interconnect it into your VCF.

7. (Required only if you are adding a device that turns a non-mixed VCF into a mixed VCF) Log in to the VCF and set all devices in the VCF to mixed mode. Configure all devices to reboot to complete this procedure.

```
user@device> request virtual-chassis mode mixed all-members reboot
```

The VCF experiences downtime as part of the reboot procedure.

8. After the device reboots, interconnect the new device into the VCF using supported SFP+ or QSFP+ interfaces.

The SFP+ or QSFP+ links are converted into VCPs automatically.

The new device should be operational shortly after the cabling is complete.

### Adding a Spine or Leaf Device to a Nonprovisioned Virtual Chassis Fabric



**CAUTION:** Configure your VCF using autoprovisioning or preprovisioning unless you have a compelling reason to use nonprovisioned configuration. You can configure all aspects of a VCF using autoprovisioned or preprovisioned configuration.

Nonprovisioned VCF configuration is highly discouraged. Nonprovisioned VCF configuration should only be used by VCF experts in specialized scenarios.

Update your device to the same version of Junos OS running on the devices in the VCF before interconnecting it into the VCF. See *Upgrading Software* or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To add a spine or leaf device to a nonprovisioned VCF:

1. Log in to your VCF.
2. If you are replacing a device that is already part of the VCF, power off the device in the VCF. Uncable the device once the power off is complete.

You can skip this step if you are adding a new device without replacing an existing device. You must skip this step if there is no configuration for the device that you are removing from the VCF.

If the device is configured, delete the device from the VCF configuration:

```
[edit virtual-chassis]
user@device# delete member member-id
```

where *member-id* is the member ID of the device that you are removing.

3. Log in to the device that you are going to add to the VCF.
4. Configure the device into fabric mode. If needed, also configure the device into mixed mode.

Reboot the device to complete this configuration step.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```

If you are configuring a VCF that includes at least one QFX3600, QFX3500, or EX4300 devices as a leaf device:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** If you are adding a QFX3600, QFX3500, or EX4300 device to a VCF that is composed entirely of QFX5100 devices, you must also log in to the VCF and set all of the devices in the VCF into mixed mode.

Log in to the VCF and enter the **request virtual-chassis mode mixed all-members reboot** command to perform this task.

The VCF reboots and incurs downtime to complete this procedure.



**NOTE:** We recommend that you set the fabric and mixed mode settings before you interconnect your devices into a VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

You can, however, use the **request virtual-chassis mode fabric local** or **request virtual-chassis mode mixed local** commands to set a device into fabric or mixed mode after interconnecting your VCF.

5. (Required only if you are adding a device that turns a non-mixed VCF into a mixed VCF) Log in to the VCF and set all devices in the VCF to mixed mode, Configure all devices to reboot to complete this procedure.

```
user@device> request virtual-chassis mode mixed all-members reboot
```

The VCF experiences downtime as part of the reboot procedure.

6. After the device reboots, interconnect it into the VCF using supported SFP+ or QSFP+ interfaces.

7. Configure the SFP+ or QSFP+ interfaces into Virtual Chassis ports (VCPs):

```
user@device> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
user@device> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
```

The **request virtual-chassis vc-port** must be configured on the ports at both ends of the link in order for that link to be configured into a VCP.

8. (Optional) Log in to the VCF and set the mastership priority of the new device:

```
[edit virtual-chassis]
user@device# set member member-id mastership-priority number
```

If needed, enter the **show virtual-chassis** command to learn the member ID of the new member device in the VCF.

#### Related Documentation

- [Removing a Device From a Virtual Chassis Fabric on page 5277](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- [Understanding Virtual Chassis Fabric Configuration on page 5251](#)
- [Understanding Virtual Chassis Fabric Components on page 5243](#)

## Removing a Device From a Virtual Chassis Fabric

This topic describes how to remove a device from a Virtual Chassis Fabric (VCF):

To remove a device from a VCF:

1. Power off the device that you are removing from the VCF.
2. Uncable the device that you are removing from the VCF.
3. Log in to the Virtual Management ethernet (VME) interface. Remove the device from the VCF configuration.

You can skip this step if you are removing a device that was never configured.

```
[edit virtual-chassis]
user@device# delete member member-id
```

4. Delete the Virtual Chassis port (VCP) or ports on the devices that are still in the VCF but were connected to the removed device.

```
user@device> request virtual-chassis vc-port delete pic-slot pic-slot port port-number member member-id
```

When a device is removed from a VCF, the interface on the other end of the VCP link that was connected to the removed device remains configured as a VCP.

You can check the results of this command using the **show virtual-chassis vc-port** command.

5. (Required only if you are removing a device that turns a mixed VCF into a homogenous VCF) Log in to the VCF and disable mixed mode for all of the devices in the VCF, Configure all devices to reboot to complete this procedure.

```
user@device> request virtual-chassis mode mixed disable all-members reboot
```

This step should only be taken if you are removing a QFX3600, QFX3500, or EX4300 device from a mixed VCF and the only devices remaining in the VCF are QFX5100 devices.

The VCF experiences downtime as part of the reboot procedure.

6. Commit the configuration.

```
[edit]
user@device# commit
```

7. Power on the device that was removed from the VCF, and log in to it.

8. (Optional, but recommended) Delete the VCP or VCPs on the device that was removed:

```
user@device> request virtual-chassis vc-port delete pic-slot pic-slot port port-number member member-id
```

9. (Optional, but recommended) Reset the fabric and mixed mode settings.

If you are removing a device that was part of a VCF composed entirely of the same device:

```
user@device> request virtual-chassis mode fabric disable reboot
```

If you are removing a device that was part of a VCF composed of two or more device types:

```
user@device> request virtual-chassis mode fabric mixed disable reboot
```

Reboot the device to complete the process.

We recommend resetting the fabric and mixed mode settings immediately after removing it from the VCF to avoid any potential issues with your device if it is placed in your network in another role.

#### Related Documentation

- [Adding a Device to a Virtual Chassis Fabric on page 5270](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- [Understanding Virtual Chassis Fabric Configuration on page 5251](#)
- [Understanding Virtual Chassis Fabric Components on page 5243](#)

## Upgrading Software for a Virtual Chassis Fabric

This topic describes the processes that can be used to update software on an operational Virtual Chassis Fabric (VCF).

You should update the software on each device before initially interconnecting your VCF. This process describes the options that are available for upgrading software after a VCF is setup.

It contains the following sections:

- [NSSU on page 5279](#)
- [Automatic Software Update on page 5279](#)
- [Standard Upgrade on page 5279](#)

## NSSU

Nonstop software upgrade (NSSU) enables you to upgrade the software running on all member devices in a VCF with minimal network traffic disruption during the upgrade.

See *Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade*.

## Automatic Software Update

Automatic software update automatically upgrades the Junos OS running on a device joining a VCF to the version of Junos OS running on the VCF at the moment the new device is cabled into the VCF.

To configure the automatic software update feature for a VCF composed entirely of QFX5100 devices:

[edit]

```
user@device# set virtual-chassis auto-sw-update package-name package-name
```

To configure the automatic software update feature on a VCF composed of QFX5100 devices and at least one other type of device:

[edit]

```
user@device# set virtual-chassis auto-sw-update qfx-5 package-name package-name
```

```
user@device# set virtual-chassis auto-sw-update qfx-3 package-name package-name
```

```
user@device# set virtual-chassis auto-sw-update ex-4300 package-name package-name
```

where **qfx-5** specifies the path to the Junos OS used to run a QFX5100 devices, **qfx-3** specifies the path to the Junos OS used to run QFX3600 and QFX3500 devices, and **ex4300** specifies the path to the Junos OS used to run EX4300 switches.

If the software package is located on a local directory on the switch, use the following format for **package-name**:

**/pathname/package-name**

If the software package is to be downloaded and installed from a remote location, use one of the following formats:

**ftp://hostname/pathname/package-name**

**ftp://username:prompt@ftp.hostname.net/package-name**

**http://hostname/pathname/package-name**

## Standard Upgrade

You can upgrade software on a VCF using the traditional method of upgrading software for Junos OS by logging onto the master Routing Engine and using the **request system software add** command to initiate the upgrade on a non-mixed VCF or the **request system**

**software add set** [*package-name package-name ...*] to initiate the upgrade on a mixed VCF, where *package-name* is the path to an image for one device family.

When you upgrade Junos OS on a VCF using the traditional software upgrade, each device in the VCF must reboot. The entire system is down until the upgrade process is complete.

For information on performing this procedure, see *Upgrading Software*.

**Related  
Documentation**

- [Adding a Device to a Virtual Chassis Fabric on page 5270](#)
- [Understanding Software Upgrades in a Virtual Chassis Fabric on page 5258](#)

---

## Configuration Statements

- [\[edit virtual-chassis\] Configuration Statement Hierarchy on page 5280](#)
- [aliases \(Virtual Chassis\) on page 5283](#)
- [alias-name \(Virtual Chassis aliases\) on page 5284](#)
- [auto-provisioned on page 5285](#)
- [auto-sw-update on page 5286](#)
- [enhanced-hash-key on page 5288](#)
- [fabric-load-balance on page 5290](#)
- [id on page 5291](#)
- [inactivity-interval \(Fabric Load Balance\) on page 5292](#)
- [location \(Virtual Chassis\) on page 5293](#)
- [mac-persistence-timer on page 5294](#)
- [mastership-priority on page 5295](#)
- [member on page 5297](#)
- [no-management-vlan on page 5298](#)
- [no-split-detection on page 5299](#)
- [package-name on page 5300](#)
- [preprovisioned on page 5301](#)
- [role on page 5302](#)
- [serial-number on page 5305](#)
- [serial-number \(Virtual Chassis aliases\) on page 5306](#)
- [traceoptions \(Virtual Chassis\) on page 5307](#)
- [virtual-chassis on page 5310](#)

### [edit virtual-chassis] Configuration Statement Hierarchy

This topic lists supported and unsupported configuration statements in the **[edit virtual-chassis]** hierarchy level on EX Series and QFX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.



- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms.

For detailed information about feature support on specific EX Series or QFX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit virtual-chassis\] Hierarchy Level on page 5281](#)
- [Unsupported Statements in the \[edit virtual-chassis\] Hierarchy Level on page 5282](#)

### Supported Statements in the [edit virtual-chassis] Hierarchy Level

The following hierarchy shows the **[edit virtual-chassis]** configuration statements supported on EX Series or QFX Series switches:

```
virtual-chassis {
  aliases {
    serial-number serial-number {
      alias-name alias-name;
    }
  }
  auto-provisioned;
  auto-sw-update {
    (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
    package-name package-name;
  }
  fast-failover (ge | vcp disable | xe);
  graceful-restart {
    disable;
  }
  id id;
  mac-persistence-timer [minutes | disable];;
  member member-id {
    location location;
    mastership-priority number;
    no-management-vlan;
    role (line-card | routing-engine);
    serial-number;
  }
  no-split-detection;
  preprovisioned;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
      regex>;
    flag flag ;
  }
  vc-port {
    lag-hash (packet-based | source-port-based);
  }
  vcp-no-hold-time;
}
```

### Unsupported Statements in the [edit virtual-chassis] Hierarchy Level

All statements in the **[edit virtual-chassis]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### **Related Documentation**

- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4300 Virtual Chassis \(CLI Procedure\) on page 5097](#)
- [Configuring an EX2200 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX3300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches \(CLI Procedure\)](#)
- [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#)

## aliases (Virtual Chassis)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>aliases {   serial-number serial-number {     alias-name alias-name;   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). An alias allows you to more clearly identify the member switches in your Virtual Chassis or VCF by assigning a text label to a member switch's serial number.</p> <p>An alias is not specified for a device until the alias name is specified using the <b>alias-name</b> keyword.</p> <p>The alias appears in the <b>Alias-Name</b> field in the <b>show virtual-chassis</b> command.</p> <p>Alias usage is optional and aliases are used for administrative purposes only. Setting an alias has no effect on the operation of the member switch.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Understanding Virtual Chassis Fabric Components on page 5243</a></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis Components</a></li> </ul>                                                                                                                                                                                                                                        |

## alias-name (Virtual Chassis aliases)

**Syntax** `alias-name alias-name;`

**Hierarchy Level** `[edit virtual-chassis aliases serial-number serial-number]`

**Release Information** Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series switches.

**Description** Create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). An alias allows you to more clearly identify the member switches in your Virtual Chassis or VCF by assigning a text label to a member switch's serial number.

The alias appears in the **Alias-Name** field in the **show virtual-chassis** command.

Alias usage is optional and aliases are used for administrative purposes only. Setting an alias has no effect on the operation of the member switch.

In the following example, the **dc-floor-1** alias name is assigned to the member switch with the serial number AB0123456789.

### set serial-number

```
[edit virtual-chassis aliases]
user@switch# set serial-number AB0123456789 alias-name dc-floor-1
```

### show virtual-chassis

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 9d5d.5556.919a
Fabric Mode: Enabled

Member ID  Status   Serial No   Alias-Name   Model          Mstr  prio  Role
0 (FPC 0)  Prsnt    AB0123456789 dc-floor-1   qfx5100-48s-6q 129   Master
<additional output removed for brevity>
```

**Options** *alias-name*—The text label, or alias, assigned to the member switch by the user.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

- Related Documentation**
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
  - [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
  - [Configuring a QFX Series Virtual Chassis \(CLI Procedure\)](#)
  - [Understanding Virtual Chassis Fabric Components on page 5243](#)
  - [Understanding QFX Series Virtual Chassis Components](#)

## auto-provisioned

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | auto-provisioned;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices in a Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Enable the auto-provisioned configuration mode for a Virtual Chassis Fabric (VCF).</p> <p>When a VCF is autoprovisioned, you can plug and play leaf devices that have not been configured or are zeroized into your VCF without user configuration. The leaf devices are automatically configured into the linecard role and all other VCF configuration—configuring Virtual Chassis ports (VCPs), the member ID, fabric mode, mixed mode (if applicable), and other parameters—is completed without further user action when a supported spine device interconnects to the leaf device by using a 10-Gbps SFP+ or 40-Gbps QSFP+ link that can be converted into a VCP.</p> <p>A leaf device whose fabric or mixed mode setting is changed as part of the autoprovisioning process automatically reboots. You can avoid this reboot by configuring the fabric or mixed mode setting on the leaf device before interconnecting into the VCF.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li> <li>• <a href="#">Understanding Virtual Chassis Fabric Configuration on page 5251</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## auto-sw-update

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>auto-sw-update {<br/>    (ex-4200   ex-4300   ex-4500   ex-4600   qfx-3   qfx-5)<br/>    package-name package-name;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>     | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>The <b>ex-4200</b> and <b>ex-4500</b> options introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>The <b>ex-4300</b>, <b>qfx-3</b>, and <b>qfx-5</b> options introduced in Junos OS Release 13.2X51-D20.</p> <p>The <b>ex-4600</b> option introduced in Junos OS Release 13.2X51-D25.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>         | <p>Enable the automatic software update feature for Virtual Chassis or Virtual Chassis Fabric (VCF) configurations.</p> <p>You should only use the keywords that specify a device—<b>ex-4300</b>, <b>ex-4600</b>, <b>qfx-3</b>, and <b>qfx-5</b>—when configuring automatic software update on a mixed Virtual Chassis or Virtual Chassis Fabric (VCF). You can simply specify the <i>package-name</i> without specifying the device keywords in non-mixed Virtual Chassis or VCF topologies.</p> <p>You must enter the <b>auto-sw-update</b> statement multiple times—once for each device family in your mixed Virtual Chassis or VCF—in most scenarios when enabling the automatic software update for a mixed Virtual Chassis or VCF.</p> <p>The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. You do not, therefore, need to specify the <b>ex-4500</b> keyword when configuring automatic software update for a mixed Virtual Chassis that include EX4500 and EX4550 switches only. You also only have to enter the <b>ex-4500</b> keyword once to configure automatic software update for all EX4500 and EX4550 member switches in the same mixed Virtual Chassis.</p> <p>The Junos OS package for a QFX3500 device updates the software for QFX3500 and QFX3600 devices. You do not, therefore, need to specify the <b>qfx-3</b> keyword when configuring automatic software update for a Virtual Chassis composed entirely of QFX3500 and QFX3600 devices. You also have to enter the <b>qfx-3</b> keyword only once to configure automatic software update for all QFX3500 and QFX3600 member devices in the same mixed Virtual Chassis.</p> <p>The remaining statement is explained separately.</p> |
| <b>Default</b>             | The automatic software update feature is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b>package-name package-name</b>—Specify a path to a Junos OS software image.</p> <p><b>ex-4200</b>—Specify a path to a Junos OS image for an EX4200 switch when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

**ex-4300**—Specify a path to a Junos OS image for an EX4300 switch when enabling automatic software update for a mixed Virtual Chassis or VCF.

**ex-4500**—Specify a path to a Junos OS image for an EX4500 switch, an EX4550 switch, or both types of switches when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. Therefore, you only enter this command once to upgrade the EX4500 and EX4550 member switches in the same mixed Virtual Chassis.

The **ex-4500** keyword also does not need to be specified when configuring automatic software update for a mixed EX4500 and EX4550 Virtual Chassis.

**ex-4600**—Specify a path to a Junos OS image for an EX4600 switch when enabling automatic software update for a mixed Virtual Chassis.

**qfx-3**—Specify a path to a Junos OS image for a QFX3500, QFX3600, or both types of devices when enabling automatic software update for a mixed VCF.

**qfx-5**—Specify a path to a Junos OS image for a QFX5100 device when enabling automatic software update for a mixed VCF.

|                                 |                                                            |
|---------------------------------|------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.        |
|                                 | system-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches</i></li> <li>• <a href="#">Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure) on page 5114</a></li> <li>• <a href="#">Understanding Software Upgrades in a Virtual Chassis Fabric on page 5258</a></li> </ul> |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## enhanced-hash-key

---

**Syntax**    enhanced-hash-key {  
              ecmp-resilient-hash;  
              fabric-load-balance {  
                  flowlet {  
                      inactivity-interval *interval*;  
                  }  
              per-packet;  
          }  
              hash-mode {  
                  layer2-header;  
                  layer2-payload;  
              }  
              inet {  
                  no-ipv4-destination-address;  
                  no-ipv4-source-address;  
                  no-l4-destination-port;  
                  no-l4-source-port;  
                  no-protocol;  
                  vlan-id;  
              }  
              inet6 {  
                  no-ipv6-destination-address;  
                  no-ipv6-source-address;  
                  no-l4-destination-port;  
                  no-l4-source-port;  
                  no-next-header;  
                  vlan-id;  
              }  
              layer2 {  
                  no-destination-mac-address;  
                  no-ether-type;  
                  no-source-mac-address;  
                  vlan-id;  
              }  
          }  
      }

**Hierarchy Level**    [edit forwarding-options]

**Release Information**    Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.  
                              Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.  
                              The **fabric-load-balance** statement introduced in Junos OS Release 14.1X53-D10.

**Description**    Configure the hashing key used to hash link aggregation group (LAG) and equal-cost multipath (ECMP) traffic, or enable adaptive load balancing (ALB) in a Virtual Chassis Fabric (VCF).

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.



For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The remaining statements are explained separately.

|                           |                                                               |
|---------------------------|---------------------------------------------------------------|
| <b>Required Privilege</b> | interface—To view this statement in the configuration.        |
| <b>Level</b>              | interface-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                             |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2686</a></li><li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2585</a></li></ul> |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## fabric-load-balance

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>fabric-load-balance {<br/>    flowlet {<br/>        inactivity-interval interval;<br/>    }<br/>    per-packet;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>     | [edit forwarding-options <b>enhanced-hash-key</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b> | Statement introduced in Junos OS Release 14.1X53-D10.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>         | <p>Enable adaptive load balancing (ALB) for a VCF, and specify how ALB is implemented.</p> <p>When ALB is enabled, the Virtual Chassis ports (VCPs) are reset. Packets are dropped and might potentially arrive out of order for a brief period of time as a result of this VCP reset. Normal operation of the VCF resumes after the VCP reset with no further user action.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>             | <p>ALB is disabled, by default.</p> <p>If you do not specify a mode when enabled ALB, ALB is enabled using flowlet mode with an inactivity timer of 16 microseconds.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>             | <p><b>flowlet</b>—Implement ALB by using flowlets.</p> <p>When ALB is implemented using flowlets, traffic flows that enter the VCF are spliced into smaller flows—flowlets—and individually forwarded across the VCF to the same destination device over different paths when the inactivity time between packet bursts on the sending interface exceeds the user-configurable inactivity interval.</p> <p>The inactivity interval is 16 microseconds by default, and can be configured using the <b>inactivity-interval</b> statement. You should configure the inactivity interval to ensure in-order packet delivery, so that overall performance is not negatively impacted by the packet reordering process at the receiving device. To ensure in-order packet delivery, the inactivity interval should be larger than the largest latency skew among all the paths in the VCF from any node to any other node.</p> <p>Implementing ALB using flowlets is especially effective in environments that periodically experience extremely large traffic flows—<i>elephant flows</i>—that are substantially larger than the majority of other traffic flowing through the VCF. The VCF is better able to manage elephant flows by splicing them into smaller flowlets using ALB.</p> <p><b>per-packet</b>—Implement ALB using per-packet mode.</p> <p>When per-packet mode is enabled, the VCF forwarding algorithm dynamically monitors all paths in the VCF and forwards packets to destination devices using the best available path at that moment. Flows are reordered at the destination node when per-packet mode is used to enable ALB, so some performance impact due to packet reordering is experienced.</p> |

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Traffic Flow Through a Virtual Chassis Fabric on page 5258](#)

## id

**Syntax** `id id;`

**Hierarchy Level** [edit [virtual-chassis](#)]

**Release Information** Statement introduced in Junos OS Release 9.3 for EX Series switches.  
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

**Description** Configure the alphanumeric string that identifies a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.

**Options** *id*—Virtual Chassis ID (VCID), which uses the ISO family address format—for example, **9622.6ac8.5345**.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge*
- [Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge \(CLI Procedure\) on page 5116](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\)](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#)
- [Understanding Virtual Chassis Member ID Numbering in an EX8200 Virtual Chassis](#)

## inactivity-interval (Fabric Load Balance)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>inactivity-interval <i>interval</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit forwarding-options <a href="#">enhanced-hash-key fabric-load-balance</a> flowlet]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Set the inactivity interval for adaptive load balancing (ALB) using flowlets within a VCF.</p> <p>The inactivity interval is the amount of time that occurs between packet bursts on a sending interface before a traffic flow is spliced into smaller traffic flows—flowlets—when ALB is implemented using flowlets. The flowlets are then individually forwarded across the VCF to the same destination device over different paths.</p> <p>You should configure the inactivity interval to ensure in-order packet delivery, so that overall performance is not negatively impacted by the packet re-ordering process at the receiving device. To ensure in-order packet delivery, the inactivity interval should be larger than the largest latency skew among all the paths in the VCF from any node to any other node.</p> |
| <b>Default</b>                  | <p>ALB is disabled, by default.</p> <p>If ALB is enabled without specifying a mode, ALB is enabled using flowlet mode with an inactivity interval of 16 microseconds.</p> <p>If ALB is enabled using flowlet mode without specifying an inactivity interval, the inactivity interval is set to 16 microseconds.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b><i>interval</i></b>—The amount of time that occurs between packet bursts on a sending interface before a traffic flow is spliced into flowlets.</p> <p><b>Range:</b> 16 microseconds (<b>16us</b>) to 32 milliseconds(<b>32ms</b>).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Traffic Flow Through a Virtual Chassis Fabric on page 5258</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## location (Virtual Chassis)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>location location;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis member member-id</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Set a description of the location of the Virtual Chassis or VCF member switch or external Routing Engine.</p> <p>The <b>Location</b> field is visible to users who enter the <b>show virtual-chassis status detail</b> command.</p> <p>Setting this description has no effect on the operation of the member device.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>location</b> —Location of the current member switch or external Routing Engine. The <b>location</b> can be any single word.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</a></li> <li>• <a href="#">Example: Configuring a Preprovisioned Mixed EX4200 and EX4500 Virtual Chassis</a></li> <li>• <a href="#">Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</a></li> <li>• <a href="#">Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX8200 Virtual Chassis (CLI Procedure)</a></li> </ul> |

## mac-persistence-timer

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mac-persistence-timer [<i>minutes</i>   <b>disable</b>];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>disable</b> introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>The maximum timer limit changed from no maximum timer limit to 60 minutes in Junos OS Release 12.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                           |
| <b>Description</b>              | <p>Specify how long the Virtual Chassis or VCF continues to use the MAC address of the switch that was originally configured in the master role as the system MAC base address after the original master switch is removed from the Virtual Chassis or VCF. The system MAC base address does not change in the event of a switchover provided the switch originally configured in the master role remains a member of the Virtual Chassis or VCF.</p> <p>The maximum timer limit is 60 minutes starting in Junos OS Release 12.2. There are no minimum or maximum timer limits in prior Junos OS releases.</p> |
| <b>Default</b>                  | The MAC persistence timer is set to 10 minutes by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>minutes</b>—Time in minutes that the member switch in the backup role continues to use the system MAC base address of the old master before using its own system MAC base address after the switch in the master role is physically disconnected or removed from the Virtual Chassis or VCF.</p> <p><b>disable</b>—Disable the MAC persistence timer. The system MAC base address never changes when the MAC persistence timer is disabled, even when the switch in the master role is physically disconnected or removed from the Virtual Chassis or VCF.</p>                                           |
| <b>Required Privilege Level</b> | <p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis (CLI Procedure) on page 5113</a></li><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li></ul>                                                                                                                                                                                                                                 |

## mastership-priority

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>mastership-priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis member</a> <i>member-id</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Mastership priority option <b>0</b> introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>The mastership priority value is the most important factor in determining the role of the member switch within a nonprovisioned Virtual Chassis or VCF configuration. Other factors (see <a href="#">“Understanding How the Master in a Virtual Chassis Is Elected” on page 5078</a>) also affect the election of the master.</p> <p>The mastership priority value takes the highest precedence in the master election algorithm. The member switch with highest mastership priority assumes the master Routing Engine role of the Virtual Chassis or VCF. Toggling back and forth between master and backup status in failover conditions is undesirable, so we recommend that you assign the same mastership priority value to both the master and the backup. Secondary factors in the master election algorithm determine which of these two members (that is, the two members that are assigned the highest mastership priority value) functions as the master of the Virtual Chassis or VCF.</p> <p>This statement is not used for the EX8200 Virtual Chassis, which determines mastership by external Routing Engine uptime. See <i>Understanding Virtual Chassis Roles in an EX8200 Virtual Chassis</i>.</p> <p>A switch with a mastership priority of <b>0</b> never takes the master or backup role.</p> |
| <b>Default</b>                  | 128                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><i>number</i>—Mastership priority value.</p> <p><b>Range:</b> 0 through 255</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX4300 Virtual Chassis (CLI Procedure) on page 5097</a></li> <li>• <a href="#">Example: Configuring an EX3300 Virtual Chassis with a Master and Backup</a></li> <li>• <a href="#">Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

- *Example: Configuring an EX4200 Virtual Chassis Interconnected Across Multiple Wiring Closets*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*



## member

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>member <i>member-id</i> {   location <i>location</i>;   mastership-priority <i>number</i>;   no-management-vlan;   serial-number <i>serial-number</i>;   role <i>role</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure a switch or an XRE200 External Routing Engine as a member of a Virtual Chassis or a Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | <p>When an EX Series switch or a QFX Series devices configured in standalone mode is powered on but not interconnected through its Virtual Chassis ports (VCPs) with other member switches, its default member ID is 0.</p> <p>There is no default member ID in an EX8200 or EX9200 Virtual Chassis. An EX8200 or EX9200 Virtual Chassis must be preprovisioned, and that process configures the member IDs.</p>                                                                                                                                                             |
| <b>Options</b>                  | <p><b><i>member-id</i></b>—Identifies a specific member switch of a Virtual Chassis or VCF configuration.</p> <p>The exact range for a specific Virtual Chassis or VCF depends on the number of switches allowed in the Virtual Chassis or VCF.</p> <p>In an EX8200 Virtual Chassis, member IDs 0 through 7 are reserved for EX8200 member switches and member IDs 8 and 9 are reserved for the master and backup external Routing Engines.</p> <p>The remaining statements are explained separately.</p>                                                                    |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</a></li> <li>• <a href="#">Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</a></li> </ul> |

- *Configuring an EX3300 Virtual Chassis (CLI Procedure)*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- *Configuring an EX8200 Virtual Chassis (CLI Procedure)*
- *Configuring an EX9200 Virtual Chassis*
- *Configuring a QFX Series Virtual Chassis (CLI Procedure)*

---

## no-management-vlan

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-management-vlan;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <b>virtual-chassis member</b> <i>member-id</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Remove the specified member's out-of-band management port from the virtual management Ethernet (VME) global management VLAN of the Virtual Chassis or VCF configuration.</p> <p>For a member that is functioning in a linecard role, you can use this configuration to reserve the member's management Ethernet port for local troubleshooting:</p> <pre>virtual-chassis {<br/>  member 2 {<br/>    no-management-vlan;<br/>  }<br/>}</pre> <p>You cannot configure the IP address for a local management Ethernet port using the CLI or the J-Web interface. To do this, you need to use the shell <b>ifconfig</b> command.</p> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Setting Up a Multimember EX4200 Virtual Chassis Access Switch with a Default Configuration</i></li><li>• <i>Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)</i></li><li>• <a href="#">Understanding Global Management of a Virtual Chassis on page 5080</a></li><li>• <a href="#">Understanding Virtual Chassis Fabric Configuration on page 5251</a></li></ul>                                                                                                                                     |

## no-split-detection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-split-detection;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Disable the split and merge feature in a Virtual Chassis or VCF configuration.</p> <p>We recommend using this statement to disable the split and merge feature when configuring a two-member Virtual Chassis. Enabling this statement on a two-member Virtual Chassis ensures that both switches remain in the correct Virtual Chassis roles in the event of a Virtual Chassis split.</p>                                                                                                                          |
| <b>Default</b>                  | The split and merge feature is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge</i></li> <li>• <a href="#">Disabling Split and Merge in a Virtual Chassis (CLI Procedure) on page 5114</a></li> <li>• <a href="#">Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure) on page 5116</a></li> <li>• <a href="#">Understanding Split and Merge in a Virtual Chassis on page 5087</a></li> </ul> |

## package-name

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>package-name <i>package-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit virtual-chassis <a href="#">auto-sw-update</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Specify the software package name or location of the software package to be used by the automatic software update feature for Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>                  | No package name is specified.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b><i>package-name</i></b>—Name of the software package or the URL to the software package to be used.</p> <ul style="list-style-type: none"><li>If the software package is located on a local directory on the switch, use the following format for <b><i>package-name</i></b>:<br/><br/><b><i>/pathname/package-name</i></b></li><li>If the software package is to be downloaded and installed from a remote location, use one of the following formats:<br/><br/><b><i>ftp://hostname/pathname/package-name</i></b><br/><b><i>ftp://username:prompt@ftp.hostname.net/package-name</i></b><br/><b><i>http://hostname/pathname/package-name</i></b></li></ul> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><i>Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches</i></li><li><a href="#">Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure)</a> on page 5114</li><li><a href="#">Understanding Software Upgrades in a Virtual Chassis Fabric</a> on page 5258</li></ul>                                                                                                                                                                                                                                                                                    |

## preprovisioned

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | preprovisioned;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Enable the preprovisioned configuration mode for a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.</p> <p>When the preprovisioned configuration mode is enabled, you cannot use the CLI or the J-Web interface to change the mastership priority or member ID of member switches.</p> <p>You must use this statement to configure an EX8200 Virtual Chassis. Nonprovisioned configuration of an EX8200 Virtual Chassis is not supported.</p>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <i>Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</i></li> <li>• <i>Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</i></li> <li>• <i>Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</i></li> <li>• <i>Configuring an EX8200 Virtual Chassis (CLI Procedure)</i></li> <li>• <i>Configuring an EX9200 Virtual Chassis</i></li> <li>• <i>Configuring a QFX Series Virtual Chassis (CLI Procedure)</i></li> <li>• <a href="#">Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 5104</a></li> </ul> |

## role

---

|                            |                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>role (line-card   routing-engine);</code>                                                                                                                                                                                            |
| <b>Hierarchy Level</b>     | [edit <a href="#">virtual-chassis preprovisioned member</a> <i>member-id</i> ]                                                                                                                                                             |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF). |
| <b>Description</b>         | Specify the roles of the members of the Virtual Chassis or a Virtual Chassis Fabric (VCF) in a preprovisioned Virtual Chassis.                                                                                                             |

### Virtual Chassis Fabric

Specify the role to be performed by each switch. In a VCF, the spine devices are configured into the Routing Engine role and the leaf devices are configured into the line card role. You can configure several devices into the Routine Engine role, but only two will operate in the Routing Engine role at a time. The role must be associated with the member's serial number.

### EX Series (except EX8200 Virtual Chassis) and QFX Series Virtual Chassis

Specify the role to be performed by each member switch. Associate the role with the member's serial number.

When you use a preprovisioned configuration, you cannot modify the mastership priority or member ID of member switches through the user interfaces. The mastership priority value is generated by the software, based on the assigned role:

- A member configured as **routing-engine** is assigned the mastership priority **129**.
- A member configured as **line-card** is assigned the mastership priority **0**.
- A member listed in the preprovisioned configuration without an explicitly specified role is assigned the mastership priority **128**.

The configured role specifications are permanent. If both **routing-engine** members fail, a **line-card** member cannot take over as master of the Virtual Chassis configuration. You must delete the preprovisioned configuration to change the specified roles in a Virtual Chassis.

Explicitly configure two members as **routing-engine** and configure additional switches as members of the preprovisioned Virtual Chassis by specifying only their serial numbers. If you do not explicitly configure the role of the additional members, they function in a linecard role by default. In that case, a member that is functioning in a linecard role can take over mastership if the members functioning as master and backup (**routing-engine** role) both fail.

### EX8200 Virtual Chassis

Specify the role to be performed by each XRE200 External Routing Engine and each EX8200 member switch. Associate the role with the member's serial number. An EX8200 Virtual Chassis cannot function when both external Routing Engines, which must be configured in the **routing-engine** role, have failed.

- Options**
- **line-card**—Enables the member to be eligible to function only in the linecard role. Any member of the Virtual Chassis or VCF configuration other than the master or backup functions in the linecard role and runs only a subset of Junos OS for EX Series switches. A member functioning in the linecard role does not run the control protocols or the chassis management processes.

A Virtual Chassis must have at least three members for one member to function in the linecard role.

In an EX8200 Virtual Chassis configuration, all member switches must be in the linecard role.

- **routing-engine**—Enables the member to function as a master or backup of the Virtual Chassis or VCF configuration. The master manages all members and runs the chassis management processes and control protocols. The backup synchronizes with the master in terms of protocol states, forwarding tables, and so forth, so that it is prepared to preserve routing information and maintain network connectivity without disruption in case the master is unavailable.

(All Virtual Chassis composed of EX Series switches, except EX8200 switches, or QFX Series devices) Specify two and only two members as **routing-engine**. The software determines which of the two members assigned the **routing-engine** role functions as master, based on the master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 5078](#). In these Virtual Chassis, the **routing-engine** role is associated with a switch.

(EX8200 Virtual Chassis) All XRE200 External Routing Engines must be in the **routing-engine** role.

|                                 |                                                            |
|---------------------------------|------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.        |
|                                 | system-control—To add this statement to the configuration. |

**Related  
Documentation**

- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
- [Preprovisioning a Virtual Chassis Fabric on page 5264](#)
- *Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File*
- *Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines*
- *Configuring an EX3300 Virtual Chassis (CLI Procedure)*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- *Configuring an EX8200 Virtual Chassis (CLI Procedure)*
- *Configuring an EX9200 Virtual Chassis*
- *Configuring a QFX Series Virtual Chassis (CLI Procedure)*
- *Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)*
- *Adding a New EX4200 Switch to an Existing EX4200 Virtual Chassis (CLI Procedure)*
- [Replacing a Member Switch of a Virtual Chassis Configuration \(CLI Procedure\) on page 5104](#)



## serial-number


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>serial-number <i>serial-number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis preprovisioned member</a> <i>member-id</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>In a preprovisioned Virtual Chassis or Virtual Chassis Fabric (VCF), specify the serial number of each member switch to be included in the configuration. If you do not include the serial number within the configuration, the switch cannot be recognized as a member of a preprovisioned configuration.</p> <p>In an EX8200 Virtual Chassis configuration, specify the serial number of each XRE200 External Routing Engine and each EX8200 member switch to be included in the Virtual Chassis configuration. If you do not include the serial number within the Virtual Chassis configuration, the external Routing Engine or switch cannot be recognized as a member of the configuration.</p>                                                                                                                                                                                                                 |
| <b>Options</b>                  | <i>serial-number</i> —Permanent serial number for the external Routing Engine or for the member switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <a href="#">Configuring an EX2200 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX3300 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX4300 Virtual Chassis (CLI Procedure) on page 5097</a></li> <li>• <a href="#">Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX8200 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX9200 Virtual Chassis</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)</a></li> </ul> |

## serial-number (Virtual Chassis aliases)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>serial-number <i>serial-number</i> {<br/>    <i>alias-name</i> <i>alias-name</i>;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis aliases</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series Virtual Chassis and Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Specify the serial number that will be labeled with an alias in a Virtual Chassis or Virtual Chassis Fabric (VCF).</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b><i>serial-number</i></b>—Permanent serial number for the member switch in the Virtual Chassis or VCF.</p> <p>You can retrieve the serial number for any device in your Virtual Chassis or VCF by entering the <b>show virtual-chassis</b> command and reviewing the output in the <b>Serial No</b> field.</p>                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Understanding Virtual Chassis Fabric Components on page 5243</a></li><li>• <a href="#">Understanding QFX Series Virtual Chassis Components</a></li></ul> |

## tracoptions (Virtual Chassis)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>tracoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i> &lt;detail&gt; &lt;disable&gt; &lt;receive&gt; &lt;send&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>     | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>detail</b> added in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>         | Define tracing operations for the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>             | Tracing operations are disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>             | <p><b>detail</b>—(Optional) Generate detailed trace information for a flag.</p> <p><b>disable</b>—(Optional) Disable a flag.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> <li><b>all</b>—All tracing operations.</li> </ul> |
|                            | <p> <b>TIP:</b> The all flag displays a subset of logs that are useful in debugging most issues. For more detailed information, use all detail.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|                            | <ul style="list-style-type: none"> <li><b>auto-configuration</b>—Trace Virtual Chassis ports (VCPs) that have been automatically configured.</li> <li><b>csn</b>—Trace Virtual Chassis complete sequence number (CSN) packets.</li> <li><b>error</b>—Trace Virtual Chassis errored packets.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- **hello**—Trace Virtual Chassis hello packets.
- **krt**—Trace Virtual Chassis KRT events.
- **lsp**—Trace Virtual Chassis link-state packets.
- **lsp-generation**—Trace Virtual Chassis link-state packet generation.
- **me**—Trace Virtual Chassis ME events.
- **normal**—Trace normal events.
- **packets**—Trace Virtual Chassis packets.
- **parse**—Trace reading of the configuration.
- **psn**—Trace partial sequence number (PSN) packets.
- **route**—Trace Virtual Chassis routing information.
- **spf**—Trace Virtual Chassis SPF events.
- **state**—Trace Virtual Chassis state transitions.
- **task**—Trace Virtual Chassis task operations.

**no-stamp**—(Optional) Do not place a timestamp on any trace file.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**receive**—(Optional) Trace received packets.

**replace**—(Optional) Replace a trace file rather than appending information to it.

**send**—(Optional) Trace transmitted packets.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

|                                 |                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | <b>system</b> —To view this statement in the configuration.<br><b>system-control</b> —To add this statement to the configuration. |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|

**Related  
Documentation**

- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157](#)
- [Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 5154](#)
- [Verifying That Virtual Chassis Ports Are Operational on page 5155](#)
- [\*Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis\*](#)
- [Troubleshooting an EX Series Virtual Chassis on page 5235](#)

## virtual-chassis

```
Syntax  virtual-chassis {
        aliases {
            serial-number serial-number {
                alias-name alias-name;
            }
        }
        auto-provisioned
        auto-sw-update {
            (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
            package-name package-name;
        }
        fast-failover (ge | vcp disable | xe);
        graceful-restart {
            disable;
        }
        id id;
        mac-persistence-timer [minutes | disable];;
        member member-id {
            location location;
            mastership-priority number;
            no-management-vlan;
            serial-number;
            role;
        }
        no-split-detection;
        preprovisioned;
        traceoptions (Virtual Chassis) {
            file filename <files number> <size size> <world-readable | no-world-readable> <match
                regex>;
            flag flag ;
        }
        vc-port {
            lag-hash (packet-based | source-port-based);
        }
        vcp-no-hold-time;
    }
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

**Description** Configure a Virtual Chassis or a Virtual Chassis Fabric (VCF).

The remaining statements are explained separately.

**Default** A standalone EX Series switch is a Virtual Chassis by default. It has a default member ID of 0, a default mastership priority of 128, and a default role as master.

A QFX Series device configured in standalone mode is a Virtual Chassis by default. It has a default member ID of 0, a default mastership priority of 128, and a default role as master.

A standalone XRE200 External Routing Engine or EX8200 switch is not part of an EX8200 Virtual Chassis until a Virtual Chassis configuration is set up.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 5261</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 5264</a></li> <li>• <i>Configuring a QFX Series Virtual Chassis (CLI Procedure)</i></li> <li>• <i>Example: Configuring an EX3300 Virtual Chassis with a Master and Backup</i></li> <li>• <i>Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet</i></li> <li>• <i>Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</i></li> <li>• <i>Configuring an EX3300 Virtual Chassis (CLI Procedure)</i></li> <li>• <i>Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</i></li> <li>• <i>Configuring an EX8200 Virtual Chassis (CLI Procedure)</i></li> <li>• <i>Configuring an EX9200 Virtual Chassis</i></li> </ul> |





## CHAPTER 90

# Administration

- [Routine Monitoring on page 5313](#)
- [Operational Commands on page 5315](#)

## Routine Monitoring

---

- [Verifying the Member ID, Role, Status, and Neighbor Member Connections of a Virtual Chassis Fabric Member Device on page 5313](#)
- [Verifying Virtual Chassis Port Connections in a Virtual Chassis Fabric on page 5314](#)
- [Verifying the Virtual Chassis Fabric Mode Settings on page 5315](#)

### Verifying the Member ID, Role, Status, and Neighbor Member Connections of a Virtual Chassis Fabric Member Device

**Purpose** Use this procedure to learn the current member ID, role, status, Virtual Chassis port (VCP) connections, and other information for the devices in your VCF.

Understanding the current member IDs, roles, device statuses, and VCP connections is required for routine monitoring of your VCF. You'll often need to identify this basic operational information to confirm a device or a VCP is working properly in the VCF, or how the VCF topology changed as a result of a configuration change or network error.

**Action** To display VCF status using the CLI:

#### [show virtual-chassis \(Virtual Chassis Fabric\)](#)

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 0282.5fa0.3f08
Fabric Mode: Enabled
```

|             |        |              |             | Mstr |         | Mixed Route Neighbor |      |    |   |
|-------------|--------|--------------|-------------|------|---------|----------------------|------|----|---|
| List        |        |              |             | prio | Role    | Mode                 | Mode | ID |   |
| Member ID   | Status | Serial No    | Model       |      |         |                      |      |    |   |
| Interface   |        |              |             |      |         |                      |      |    |   |
| 0 (FPC 0)   | Prsnt  | AB3112430001 | qfx5100-48s | 129  | Master* | N                    | F    | 3  |   |
| vcp-255/1/0 |        |              |             |      |         |                      |      |    | 2 |
| vcp-255/1/1 |        |              |             |      |         |                      |      |    | 4 |
| vcp-255/1/2 |        |              |             |      |         |                      |      |    | 4 |
| vcp-255/1/3 |        |              |             |      |         |                      |      |    | 4 |

```

1 (FPC 1) Prsnt    AB3112230001 qfx5100-48s    129    Backup      N  F    3
vcp-255/1/0
   2
vcp-255/1/1
   4
vcp-255/1/2
   4
vcp-255/1/3
2 (FPC 2) Prsnt    AB3112460011 qfx5100-48s      0    Linecard    N  F    1
vcp-255/1/0
   0
vcp-255/1/1
3 (FPC 3) Prsnt    AB3112460011 qfx5100-48s      0    Linecard    N  F    1
vcp-255/1/0
   0
vcp-255/1/1
4 (FPC 4) Prsnt    AB3112430011 qfx5100-48s      0    Linecard    N  F    1
vcp-255/1/0
   0
vcp-255/1/1

```

**Meaning** This output verifies that fabric mode is enabled and that all devices in the VCF are participating in the fabric, as shown by the **Prsnt** status output for each device.

The Neighbor ID and Interface outputs show that all VCPs are operating correctly.

- Related Documentation**
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
  - [Preprovisioning a Virtual Chassis Fabric on page 5264](#)

## Verifying Virtual Chassis Port Connections in a Virtual Chassis Fabric

**Purpose** Verify the Virtual Chassis Ports (VCPs) in your Virtual Chassis Fabric (VCF).

You should use this command if you suspect a VCP link in your VCF is broken.

**Action** To display the VCPs of a device:

```
user@switch> show virtual-chassis vc-port member 4
fpc4:
```

```

-----
Interface  Type           Trunk  Status  Speed  Neighbor
or         / Port         ID      Up      (mbps)  ID  Interface
0/48      Auto-Configured -1     Up      40000   0   vcp-255/0/2
0/49      Auto-Configured -1     Up      40000   1   vcp-255/0/2
0/50      Auto-Configured -1     Up      40000   2   vcp-255/0/2
0/51      Auto-Configured -1     Up      40000   3   vcp-255/0/2

```

**Meaning** All of the VCPs on this device are up and active.

If the **Status** of an interface is **Absent** or the interface that you thought was a VCP does not appear in the command output, you likely have a problem with a link that has not been converted into a VCP. In this scenario, configure the interface on the link into a VCP using the **request virtual-chassis vc-port** command.

- Related Documentation**
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
  - [Preprovisioning a Virtual Chassis Fabric on page 5264](#)

## Verifying the Virtual Chassis Fabric Mode Settings

**Purpose** Verify the Virtual Chassis Fabric (VCF) mode settings on a device.

You must configure all devices in a VCF into fabric mode using the **request virtual-chassis mode fabric** command for the devices to operate in a VCF.

All VCFs use QFX5100 devices in the spine role. If a VCF uses a QFX3500, QFX3600, or EX4300 devices as a leaf node, you must also configure each device into mixed mode using the **request virtual-chassis mode mixed** command.

You must also configure a device out of mixed and fabric mode if it is removed from a VCF and placed into your network in a different role.

**Action** To display the current mode of a device:

```
user@switch> show virtual-chassis mode
fpc0:
```

```
-----
Current mode : Fabric with mixed devices
Future mode after reboot : Fabric with mixed devices
```

**Meaning** The output indicates that the switch is currently in mixed and fabric mode.

The output also indicates that the mode will not change when the device is rebooted without further configuration. You must reboot the device to change the fabric or mixed mode, so the **Future mode after reboot** output differs from the **Current mode** output when the mode has been changed but the device has not been rebooted.

- Related Documentation**
- [Autoprovisioning a Virtual Chassis Fabric on page 5261](#)
  - [Preprovisioning a Virtual Chassis Fabric on page 5264](#)

## Operational Commands

- [clear virtual-chassis vc-port statistics](#)
- [request session member](#)
- [request virtual-chassis mode](#)
- [request virtual-chassis reactivate](#)
- [request virtual-chassis vc-port](#)
- [request virtual-chassis vc-port diagnostics optics](#)
- [show forwarding-options enhanced-hash-key](#)
- [show virtual-chassis active-topology](#)
- [show virtual-chassis device-topology](#)

- [show virtual-chassis login](#)
- [show virtual-chassis mode](#)
- [show virtual-chassis protocol adjacency](#)
- [show virtual-chassis protocol database](#)
- [show virtual-chassis protocol interface](#)
- [show virtual-chassis protocol route](#)
- [show virtual-chassis protocol statistics](#)
- [show virtual-chassis](#)
- [show virtual-chassis vc-port](#)
- [show virtual-chassis vc-port diagnostics optics](#)
- [show virtual-chassis vc-port statistics](#)

## clear virtual-chassis vc-port statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear virtual-chassis vc-port statistics<br><all-members><br><interface-name><br><local><br><member member-id>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The options <b>all-members</b> and <b>local</b> were added in Junos OS Release 9.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric.</p>                                                                                                                                                                                                                              |
| <b>Description</b>              | Clear—reset to zero (0)—the traffic statistics counters on Virtual Chassis ports (VCPs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Clear traffic statistics for the specified VCP.</p> <p><b>local</b>—(Optional) Clear traffic statistics for VCPs from the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Clear traffic statistics for VCPs from the specified member of a Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis vc-port statistics on page 5227</a></li> <li>• <a href="#">show virtual-chassis vc-port on page 5209</a></li> <li>• <a href="#">Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157</a></li> </ul>                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">clear virtual-chassis vc-port statistics (EX4200 Virtual Chassis) on page 5317</a><br><a href="#">clear virtual-chassis vc-port statistics (EX8200 Virtual Chassis) on page 5318</a><br><a href="#">clear virtual-chassis vc-port statistics member 3 on page 5318</a>                                                                                                                                                                                                                                                                                                           |

### Sample Output

#### clear virtual-chassis vc-port statistics (EX4200 Virtual Chassis)

```
user@switch> clear virtual-chassis vc-port statistics
fpc0:
-----
Statistics cleared
```

### clear virtual-chassis vc-port statistics (EX8200 Virtual Chassis)

```
user@external-routing-engine> clear virtual-chassis vc-port statistics
```

```
member0:
```

```
-----  
Statistics cleared
```

```
member1:
```

```
-----  
Statistics cleared
```

```
member8:
```

```
-----  
Statistics cleared
```

```
member9:
```

```
-----  
Statistics cleared
```

### clear virtual-chassis vc-port statistics member 3

```
user@switch> clear virtual-chassis vc-port statistics member 3
```

```
Cleared statistics on member 3
```

---

## request session member

---

|                                 |                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>request session member <i>member-id</i></code>                                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                          |
| <b>Description</b>              | Start a session with the specified member of a Virtual Chassis or a VCF.                                                                                                                                                                                      |
| <b>Options</b>                  | <i>member-id</i> —Member ID for the specific member of the Virtual Chassis or VCF.                                                                                                                                                                            |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">member on page 5131</a></li><li>• <a href="#">Understanding EX Series Virtual Chassis Components on page 5067</a></li><li>• <a href="#">Understanding QFX Series Virtual Chassis Components</a></li></ul> |

## request virtual-chassis mode

---

**Syntax**    request virtual-chassis mode  
              fabric  
              mixed  
              <disable>  
              <reboot>  
              <all-members>  
              <local>  
              <member *member-id*>

**Release Information**    Command introduced in Junos OS Release 11.1 for EX Series switches.  
                              Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.  
                              The **fabric** keyword introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices in a Virtual Chassis Fabric (VCF).  
                              Command introduced in Junos OS Release 13.2X51-D20 for VCF.

**Description**    Configure the mode for a device or multiple devices in a Virtual Chassis or a VCF.

A device must be configured in fabric mode to participate as a member device in a VCF.

A device must be configured in mixed mode when it is participating in a Virtual Chassis or a VCF with different types of devices.

Do not enable the **request virtual-chassis mode mixed** command for a standalone device or for a member switch that is intended to remain in a non-mixed Virtual Chassis or VCF. Enabling this command reduces the maximum scaling numbers for some features on the switch, Virtual Chassis, or VCF.

You do not need to configure mixed mode if the only devices in your Virtual Chassis are EX4500 and EX4550 switches.

To avoid potential traffic disruptions and configuration issues for a mixed Virtual Chassis, we recommend configuring mixed mode on your device before cabling it into your Virtual Chassis. We recommend rebooting your device to complete this configuration procedure before interconnecting your device into the Virtual Chassis.

To avoid potential traffic disruptions and configuration issues, we recommend configuring the fabric and, if applicable, the mixed mode settings on your device before cabling it into a VCF. We recommend rebooting your device to complete this configuration procedure before interconnecting your device into the VCF. You can change the fabric and mixed mode settings after the device has been added to a Virtual Chassis or VCF, however.

If you set some of the devices in a mixed Virtual Chassis or VCF to mixed mode using this command but not others, the mixed Virtual Chassis or VCF might not form. If you experience this issue, enter the **request virtual-chassis mode mixed all-members** command to set the Virtual Chassis mode to mixed for all devices in the Virtual Chassis or VCF. You then need to reboot the devices that have been set into mixed mode to complete the procedure. The Virtual Chassis or VCF forms after the devices have rebooted.

When you do not use this command to set any of the switches in a mixed EX4200 and EX4500 Virtual Chassis to mixed mode, a mixed EX4200 and EX4500 Virtual Chassis



forms with one of the switches assuming the master role if the switches are running Junos OS Release 11.4 or later. All other switches in the mixed EX4200 and EX4500 Virtual Chassis are placed into the linecard role. If you experience this behavior, enter the **request virtual-chassis mode mixed all-members** command to set the Virtual Chassis mode to mixed for all switches in the Virtual Chassis. You will then need to reboot the switches to complete the procedure. The Virtual Chassis will form after all of the switches have rebooted.

The Virtual Chassis mode setting is maintained through reboots even though it is set in operational mode.

- Options**
- none**—Set the Virtual Chassis mode for all members of the Virtual Chassis or VCF.
  - all-members**—(Optional) Set the Virtual Chassis mode for all members of the Virtual Chassis or VCF.
  - disable**—Disable the Virtual Chassis fabric or mixed mode setting if it was previously enabled.
  - fabric**—Set the device into fabric mode so that the device can participate in a VCF.
  - local**—(Optional) Set the Virtual Chassis mode on the member device where the command is issued.
  - member *member-id***—(Optional) Set the Virtual Chassis mode to mixed on the specified member of the Virtual Chassis or VCF.
  - mixed**—Set the device into mixed mode so that the device can participate in a mixed Virtual Chassis or mixed VCF.



**NOTE:** You do not need to set mixed mode if the only devices in your Virtual Chassis are QFX3500 and QFX3600 devices.

You do not need to configure mixed mode if the only devices in your Virtual Chassis are EX4500 and EX4550 switches.

**Required Privilege Level** system-control

- Related Documentation**
- [Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches \(CLI Procedure\)](#)
  - [Verifying the Virtual Chassis Fabric Mode Settings on page 5315](#)
  - [Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 5154](#)

**List of Sample Output** [request virtual-chassis mode mixed on page 5322](#)  
[request virtual-chassis mode fabric mixed reboot on page 5322](#)

## Sample Output

request virtual-chassis mode mixed

```
user@switch> request virtual-chassis mode mixed
```

## Sample Output

request virtual-chassis mode fabric mixed reboot

```
user@switch> request virtual-chassis mode fabric mixed reboot
```

---

## request virtual-chassis reactivate

---

|                                 |                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>request virtual-chassis reactivate</code>                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.3 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                    |
| <b>Description</b>              | <p>Reactivate a device that has been assigned a member ID but is not currently connected to the Virtual Chassis or VCF.</p> <p>You can use this command to reactivate a device that was previously part of the Virtual Chassis or VCF but whose status is no longer <b>Prsnt</b>.</p>         |
| <b>Required Privilege Level</b> | system-control                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 5154</a></li><li>• <i>Verifying the Member ID, Role, and Neighbor Member Connections of an EX8200 Virtual Chassis Member</i></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">request virtual-chassis reactivate on page 5323</a>                                                                                                                                                                                                                               |

### Sample Output

#### request virtual-chassis reactivate

```
user@switch> request virtual-chassis reactivate
```

## request virtual-chassis vc-port

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>request virtual-chassis vc-port set   delete</b> <fpc-slot <i>fpc-slot</i> > pic-slot <i>pic-slot</i> port <i>port-number</i> <member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>fpc-slot</b> introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Enable or disable an optical port as a Virtual Chassis port (VCP).</p> <p>If you omit <b>member <i>member-id</i></b>, this command defaults to enabling or disabling the uplink VCP or SFP network port configured as a VCP on the switch where the command is issued.</p> <p>On an EX3300 switch, uplink ports 2 and 3 are configured as VCPs by default. No other uplink ports on any other EX Series switches are configured as VCPs by default.</p> <p>You might experience a temporary traffic disruption immediately after creating or deleting a user-configured VCP in an EX8200 Virtual Chassis.</p>                |
| <b>Options</b>                  | <p><b>pic-slot <i>pic-slot</i></b>—Number of the PIC slot for the port on the switch.</p> <p><b>port <i>port-number</i></b>—Number of the port that is to be enabled or disabled as a VCP.</p> <p><b>member <i>member-id</i></b>—(Optional) Enable or disable the specified VCP on the specified member of the Virtual Chassis or VCF.</p>                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | system-control                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">request virtual-chassis vc-port</a> (dedicated port)</li><li>• <a href="#">show virtual-chassis vc-port on page 5209</a></li><li>• <a href="#">show virtual-chassis vc-port statistics on page 5227</a></li><li>• <a href="#">clear virtual-chassis vc-port statistics on page 5161</a></li><li>• <a href="#">Virtual Chassis Port (VCP) Interface Names in an EX8200 Virtual Chassis</a></li><li>• <a href="#">Understanding EX Series Virtual Chassis Components on page 5067</a></li><li>• <a href="#">Understanding QFX Series Virtual Chassis Components</a></li></ul> |
| <b>List of Sample Output</b>    | <p><a href="#">request virtual-chassis vc-port set pic-slot 1 port 0 on page 5325</a></p> <p><a href="#">request virtual-chassis vc-port set pic-slot 1 port 1 member 3 on page 5325</a></p> <p><a href="#">request virtual-chassis vc-port delete pic-slot 1 port 1 member 3 on page 5325</a></p>                                                                                                                                                                                                                                                                                                                              |

## Sample Output

**request virtual-chassis vc-port set pic-slot 1 port 0**

user@switch> request virtual-chassis vc-port set pic-slot 1 port 0

To check the results of this command, use the [show virtual-chassis vc-port](#) command.

**request virtual-chassis vc-port set pic-slot 1 port 1 member 3**

user@switch> request virtual-chassis vc-port set pic-slot 1 port 1 member 3

To check the results of this command, use the [show virtual-chassis vc-port](#) command.

**request virtual-chassis vc-port delete pic-slot 1 port 1 member 3**

user@switch> request virtual-chassis vc-port delete pic-slot 1 port 1 member 3

To check the results of this command, use the [show virtual-chassis vc-port](#) command.

## request virtual-chassis vc-port diagnostics optics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>request virtual-chassis vc-port diagnostics optics</b>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Run a digital optical monitoring (DOM) scan on the optical ports configured as Virtual Chassis ports (VCPs).</p> <p>Enter the <b>show virtual-chassis vc-port diagnostics optics</b> command to view the results of the diagnostic scan.</p> <p>On certain EX Series switches, the <b>request virtual-chassis vc-port diagnostics optics</b> command must be entered to run a diagnostic scan before you can gather the <b>show virtual-chassis vc-port diagnostics optics</b> output.</p> |
| <b>Required Privilege Level</b> | system-control                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show virtual-chassis vc-port diagnostics optics on page 5213</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                |

## Sample Output

### request virtual-chassis vc-port diagnostics optics

```
user@switch> request virtual-chassis vc-port diagnostics optics
fpc0:
-----
vc-port Diagnostics Optics Done
```

## show forwarding-options enhanced-hash-key

|                                 |                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show forwarding-options enhanced-hash-key</b>                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p><b>Fabric Load Balancing Options</b> output fields introduced in Junos OS Release 14.1X53-D10.</p>                                                                                                |
| <b>Description</b>              | <p>Display information about which packet fields are used by the hashing algorithm to make hashing decisions.</p> <p>You can configure the fields that are inspected by the hashing algorithm to make hashing decisions for traffic entering a LAG bundle using the <b>forwarding-options enhanced-hash-key</b> statement.</p>                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2686</a></li> <li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2585</a></li> <li>• <a href="#">enhanced-hash-key on page 2764</a></li> </ul> |
| <b>List of Sample Output</b>    | <p><a href="#">show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode) on page 5328</a></p> <p><a href="#">show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode) on page 5329</a></p> <p><a href="#">show forwarding-options enhanced-hash-key (Fabric Load Balancing Options) on page 5329</a></p>                                         |
| <b>Output Fields</b>            | <p><a href="#">Table 305 on page 2861</a> lists the output fields for the <b>show forwarding-options enhanced-hash-key</b> command. Output fields are listed in the approximate order in which they first appear.</p>                                                                                                                                                    |

**Table 586: show forwarding-options enhanced-hash-key Output Fields**

| Field Name                   | Field Description                                                                                        |
|------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Hash-Mode</b>             | Current hash mode: Layer 2 header or Layer 2 payload.                                                    |
| <b>Protocol</b>              | Indicates whether the Protocol field is or is not used by the hashing algorithm: Yes or No.              |
| <b>Destination L4 Port</b>   | Indicates whether the Destination L4 Port field is or is not used by the hashing algorithm: Yes or No.   |
| <b>Source L4 Port</b>        | Indicates whether the Source L4 Port field is or is not used by the hashing algorithm: Yes or No.        |
| <b>Destination IPv4 Addr</b> | Indicates whether the Destination IPv4 Addr field is or is not used by the hashing algorithm: Yes or No. |

Table 586: show forwarding-options enhanced-hash-key Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                    |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Source IPv4 Addr</b>        | Indicates whether the Source IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                  |
| <b>Vlan id</b>                 | Indicates whether the Vlan id field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                           |
| <b>Next Hdr</b>                | Indicates whether the Next Hdr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                          |
| <b>Destination IPv6 Addr</b>   | Indicates whether the Destination IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                             |
| <b>Source IPv6 Addr</b>        | Indicates whether the Source IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                  |
| <b>Ether Type</b>              | Indicates whether the Ether Type field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                        |
| <b>Destination MAC Address</b> | Indicates whether the Destination MAC Address field is or is not used by the hashing algorithm: Yes or No.                                                                                                                           |
| <b>Source MAC Address</b>      | Indicates whether the Source MAC Address field is or is not used by the hashing algorithm: Yes or No.                                                                                                                                |
| <b>Load Balancing Method</b>   | Indicates the load balancing method for adaptive load balancing (ALB): flowlet or per-packet.<br><br>The load balancing method is flowlet by default, and can be configured using the <a href="#">fabric-load-balance</a> statement. |
| <b>Fabric Link Scale</b>       | Indicates the fabric link scale, in mbps.                                                                                                                                                                                            |
| <b>Inactivity Interval</b>     | Indicates the fabric load balance inactivity interval, in microseconds (us).<br><br>The inactivity interval is 16 microseconds by default, and can be configured using the <a href="#">inactivity-interval</a> statement.            |
| <b>Hash Region Size/Trunk</b>  | Indicates the hash region size, in buckets per fabric trunk.                                                                                                                                                                         |

## Sample Output

### show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0
```

```
Current Hash Settings
-----
```



```
Hash-Mode                               :layer2-payload
```

```
inet Hash settings-
```

```
-----
```

```
inet packet fields
```

```
Protocol                               : Yes
Destination L4 Port                     : Yes
Source L4 Port                           : Yes
Destination IPv4 Addr                   : Yes
Source IPv4 Addr                         : Yes
Vlan id                                 : No
```

```
inet6 Hash settings-
```

```
-----
```

```
inet6 packet fields
```

```
Next Hdr                               : Yes
Destination L4 Port                     : Yes
Source L4 Port                           : Yes
Destination IPv6 Addr                   : Yes
Source IPv6 Addr                         : Yes
Vlan id                                 : No
```

#### show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0
```

```
Current Hash Settings
```

```
-----
```

```
Hash-Mode                               : layer2-header
```

```
layer2 Hash settings-
```

```
-----
```

```
layer2 packet fields
```

```
Ether Type                             : Yes
Destination MAC Address                  : Yes
Source MAC Address                       : Yes
VLAN ID                                 : No
```

#### show forwarding-options enhanced-hash-key (Fabric Load Balancing Options)

```
user@switch> show forwarding-options enhanced-hash-key
<some output removed for brevity>
```

```
Fabric Load Balancing Options
```

```
-----
```

```
Load Balancing Method : Flowlet
Fabric Link Scale      : 40960 (mbps)
Inactivity Interval    : 16 (us)
Hash Region Size/Trunk : 1024 (buckets)
```

## show virtual-chassis active-topology

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis active-topology<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display the active topology of the Virtual Chassis or VCF with next-hop reachability information.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>none</b>—Display the active topology of the member switch where the command is issued.</p> <p><b>all-members</b>—(Optional) Display the active topology of all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the active topology of the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the active topology of the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157</a></li> <li>• <a href="#">Understanding EX Series Virtual Chassis Configuration on page 5086</a></li> </ul>                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis active-topology (EX4200 Virtual Chassis) on page 5331</a><br><a href="#">show virtual-chassis active-topology (EX8200 Virtual Chassis) on page 5331</a><br><a href="#">show virtual-chassis active-topology (Virtual Chassis Fabric) on page 5332</a>                                                                                                                                                                                                    |
| <b>Output Fields</b>            | Table 572 on page 5174 lists the output fields for the <b>show virtual-chassis active-topology</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                       |

**Table 587: show virtual-chassis active-topology Output Fields**

| Field Name            | Field Description                                                                                                                                                                              |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Destination ID</b> | Specifies the member ID of the destination.                                                                                                                                                    |
| <b>Next-hop</b>       | <p>Specifies the member ID and Virtual Chassis port (VCP) of the next hop to which packets for the destination ID are forwarded.</p> <p>The next hop can be more than one device in a VCF.</p> |

## Sample Output

### show virtual-chassis active-topology (EX4200 Virtual Chassis)

```

user@switch> show virtual-chassis active-topology
 1                      1(vcp-1)

 2                      1(vcp-1)

 3                      1(vcp-1)

 4                      1(vcp-1)

 5                      8(vcp-0) 1(vcp-1)

 6                      8(vcp-0)

 7                      8(vcp-0)

 8                      8(vcp-0)

```

### show virtual-chassis active-topology (EX8200 Virtual Chassis)

```

user@external-routing-engine> show virtual-chassis active-topology
member0:

```

| Destination ID | Next-hop           |
|----------------|--------------------|
| 1              | 1(vcp-4/0/4.32768) |
| 8              | 8(vcp-0/0.32768)   |
| 9              | 8(vcp-0/0.32768)   |

```
member1:
```

| Destination ID | Next-hop           |
|----------------|--------------------|
| 0              | 0(vcp-3/0/4.32768) |
| 8              | 8(vcp-0/0.32768)   |
| 9              | 8(vcp-0/0.32768)   |

```
member8:
```

---

| Destination ID | Next-hop         |
|----------------|------------------|
| 0              | 0(vcp-1/1.32768) |
| 1              | 1(vcp-1/2.32768) |
| 9              | 9(vcp-2/1.32768) |

member9:

---

| Destination ID | Next-hop         |
|----------------|------------------|
| 0              | 8(vcp-1/2.32768) |
| 1              | 8(vcp-1/2.32768) |
| 8              | 8(vcp-1/2.32768) |

### show virtual-chassis active-topology (Virtual Chassis Fabric)

user@device> show virtual-chassis active-topology  
fpc0:

---

| Destination ID            | Next-hop                                  |
|---------------------------|-------------------------------------------|
| 1<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 2<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 3<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 4                         | 4(vcp-255/0/2.32768)                      |
| 5                         | 5(vcp-255/0/3.32768)                      |
| 6                         | 6(vcp-255/0/1.32768)                      |

fpc1:

---

| Destination ID            | Next-hop                                  |
|---------------------------|-------------------------------------------|
| 0<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 2<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 3<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 4                         | 4(vcp-255/0/2.32768)                      |
| 5                         | 5(vcp-255/0/3.32768)                      |
| 6                         | 6(vcp-255/0/1.32768)                      |

fpc2:

| Destination ID            | Next-hop                                  |
|---------------------------|-------------------------------------------|
| 0<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 1<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 3<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 4                         | 4(vcp-255/0/2.32768)                      |
| 5                         | 5(vcp-255/0/3.32768)                      |
| 6                         | 6(vcp-255/0/1.32768)                      |

fpc3:

| Destination ID            | Next-hop                                  |
|---------------------------|-------------------------------------------|
| 0<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 1<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 2<br>6(vcp-255/0/1.32768) | 4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768) |
| 4                         | 4(vcp-255/0/2.32768)                      |
| 5                         | 5(vcp-255/0/3.32768)                      |
| 6                         | 6(vcp-255/0/1.32768)                      |

fpc4:

| Destination ID             | Next-hop                                                             |
|----------------------------|----------------------------------------------------------------------|
| 0                          | 0(vcp-255/0/48.32768)                                                |
| 1                          | 1(vcp-255/0/49.32768)                                                |
| 2                          | 2(vcp-255/0/50.32768)                                                |
| 3                          | 3(vcp-255/0/51.32768)                                                |
| 5<br>0(vcp-255/0/48.32768) | 3(vcp-255/0/51.32768) 2(vcp-255/0/50.32768)<br>1(vcp-255/0/49.32768) |
| 6<br>0(vcp-255/0/48.32768) | 3(vcp-255/0/51.32768) 2(vcp-255/0/50.32768)<br>1(vcp-255/0/49.32768) |

fpc5:

| Destination ID | Next-hop              |
|----------------|-----------------------|
| 0              | 0(vcp-255/0/48.32768) |

|                       |                       |                       |
|-----------------------|-----------------------|-----------------------|
| 1                     | 1(vcp-255/0/49.32768) |                       |
| 2                     | 2(vcp-255/0/50.32768) |                       |
| 3                     | 3(vcp-255/0/51.32768) |                       |
| 4                     | 3(vcp-255/0/51.32768) | 2(vcp-255/0/50.32768) |
| 0(vcp-255/0/48.32768) | 1(vcp-255/0/49.32768) |                       |
| 6                     | 3(vcp-255/0/51.32768) | 2(vcp-255/0/50.32768) |
| 0(vcp-255/0/48.32768) | 1(vcp-255/0/49.32768) |                       |

fpc6:

| Destination ID       | Next-hop                                  |
|----------------------|-------------------------------------------|
| 0                    | 0(vcp-255/0/0.32768)                      |
| 1                    | 1(vcp-255/0/1.32768)                      |
| 2                    | 2(vcp-255/0/2.32768)                      |
| 3                    | 3(vcp-255/0/3.32768)                      |
| 4                    | 3(vcp-255/0/3.32768) 2(vcp-255/0/2.32768) |
| 0(vcp-255/0/0.32768) | 1(vcp-255/0/1.32768)                      |
| 5                    | 3(vcp-255/0/3.32768) 2(vcp-255/0/2.32768) |
| 0(vcp-255/0/0.32768) | 1(vcp-255/0/1.32768)                      |

## show virtual-chassis device-topology

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis device-topology<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Display the device topology—the member and system IDs, the VCP numbers, and device status—for all hardware devices in the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>none</b>—Display the device topology for all members of the Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Display the device topology for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the device topology for the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the device topology for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding EX Series Virtual Chassis Port Link Aggregation on page 5083</a></li> <li>• <a href="#">Understanding EX8200 Virtual Chassis Topologies</a></li> </ul>                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>            | <a href="#">Table 573 on page 5179</a> lists the output fields for the <b>show virtual-chassis device-topology</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                       |

**Table 588: show virtual-chassis device-topology Output Fields**

| Field Name    | Field Description                                                                                                                                                                                                                                                                                             |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Member</b> | Assigned member ID.                                                                                                                                                                                                                                                                                           |
| <b>Device</b> | Assigned device ID.<br><br>For an EX8200 Virtual Chassis, the member ID and the device ID are always identical.                                                                                                                                                                                               |
| <b>Status</b> | The status of the device within the Virtual Chassis or VCF. Outputs include: <ul style="list-style-type: none"> <li>• <b>Prsnt</b>—Device is currently connected to and participating in the Virtual Chassis or VCF.</li> <li>• <b>NotPrsnt</b>—Device is assigned but is not currently connected.</li> </ul> |

Table 588: show virtual-chassis device-topology Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                      |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <b>System ID</b>                 | System ID of the device.<br><br>The system ID of the device is the device's MAC address.                                               |
| <b>Member (Neighbor List)</b>    | Assigned member ID of the neighbor device.                                                                                             |
| <b>Device (Neighbor List)</b>    | Assigned device ID of the neighbor device.<br><br>For an EX8200 Virtual Chassis, the member ID and the device ID are always identical. |
| <b>Interface (Neighbor List)</b> | The interface connecting the device to the neighbor.                                                                                   |

## Sample Output

### show virtual-chassis device-topology

```
user@switch> show virtual-chassis device-topology
```

```
member0:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
0         0      Prsnt   0021.59f7.d000  8         8      vcp-0/0
                                1         1      vcp-4/0/1
1         1      Prsnt   0026.888d.6800  8         8      vcp-0/0
                                9         9      vcp-0/1
                                0         0      vcp-3/0/4
8         8      Prsnt   0000.4a75.9b7c  9         9      vcp-1/0
                                0         0      vcp-1/1
                                1         1      vcp-1/2
9         9      Prsnt   0000.73e9.9a57  8         8      vcp-1/0
                                1         1      vcp-1/1
```

```
member1:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
0         0      Prsnt   0021.59f7.d000  8         8      vcp-0/0
                                1         1      vcp-4/0/1
1         1      Prsnt   0026.888d.6800  8         8      vcp-0/0
                                9         9      vcp-0/1
                                0         0      vcp-3/0/4
8         8      Prsnt   0000.4a75.9b7c  9         9      vcp-1/0
                                0         0      vcp-1/1
                                1         1      vcp-1/2
9         9      Prsnt   0000.73e9.9a57  8         8      vcp-1/0
                                1         1      vcp-1/1
```

```
member8:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
```



|   |   |       |                |   |   |           |
|---|---|-------|----------------|---|---|-----------|
| 0 | 0 | Prsnt | 0021.59f7.d000 | 8 | 8 | vcp-0/0   |
|   |   |       |                | 1 | 1 | vcp-4/0/1 |
| 1 | 1 | Prsnt | 0026.888d.6800 | 8 | 8 | vcp-0/0   |
|   |   |       |                | 9 | 9 | vcp-0/1   |
|   |   |       |                | 0 | 0 | vcp-3/0/4 |
| 8 | 8 | Prsnt | 0000.4a75.9b7c | 9 | 9 | vcp-1/0   |
|   |   |       |                | 0 | 0 | vcp-1/1   |
|   |   |       |                | 1 | 1 | vcp-1/2   |
| 9 | 9 | Prsnt | 0000.73e9.9a57 | 8 | 8 | vcp-1/0   |
|   |   |       |                | 1 | 1 | vcp-1/1   |

member9:

|        |        |        |                | Neighbor List |        |           |
|--------|--------|--------|----------------|---------------|--------|-----------|
| Member | Device | Status | System ID      | Member        | Device | Interface |
| 0      | 0      | Prsnt  | 0021.59f7.d000 | 8             | 8      | vcp-0/0   |
|        |        |        |                | 1             | 1      | vcp-4/0/1 |
| 1      | 1      | Prsnt  | 0026.888d.6800 | 8             | 8      | vcp-0/0   |
|        |        |        |                | 9             | 9      | vcp-0/1   |
|        |        |        |                | 0             | 0      | vcp-3/0/4 |
| 8      | 8      | Prsnt  | 0000.4a75.9b7c | 9             | 9      | vcp-1/0   |
|        |        |        |                | 0             | 0      | vcp-1/1   |
|        |        |        |                | 1             | 1      | vcp-1/2   |
| 9      | 9      | Prsnt  | 0000.73e9.9a57 | 8             | 8      | vcp-1/0   |
|        |        |        |                | 1             | 1      | vcp-1/1   |

#### show virtual-chassis device-topology (Virtual Chassis Fabric)

user@device> show virtual-chassis device-topology  
fpc0:

|        |        |        |                | Neighbor List |        |              |
|--------|--------|--------|----------------|---------------|--------|--------------|
| Member | Device | Status | System ID      | Member        | Device | Interface    |
| 0      | 0      | Prsnt  | 100e.7eb6.a900 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 1      | 1      | Prsnt  | 100e.7eb8.3a40 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 2      | 2      | Prsnt  | 100e.7eb5.d700 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 3      | 3      | Prsnt  | 100e.7eb5.c440 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 4      | 4      | Prsnt  | 100e.7eb5.7e40 | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
| 5      | 5      | Prsnt  | 100e.7eb5.80c0 | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
|        |        |        |                | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
| 6      | 6      | Prsnt  | 100e.7eb6.3b00 | 3             | 3      | vcp-255/0/3  |
|        |        |        |                | 2             | 2      | vcp-255/0/2  |
|        |        |        |                | 0             | 0      | vcp-255/0/0  |
|        |        |        |                | 1             | 1      | vcp-255/0/1  |

fpc1:

Neighbor List

| Member | Device | Status | System ID      | Member | Device | Interface    |
|--------|--------|--------|----------------|--------|--------|--------------|
| 0      | 0      | Prsnt  | 100e.7eb6.a900 | 4      | 4      | vcp-255/0/2  |
|        |        |        |                | 5      | 5      | vcp-255/0/3  |
|        |        |        |                | 6      | 6      | vcp-255/0/1  |
| 1      | 1      | Prsnt  | 100e.7eb8.3a40 | 4      | 4      | vcp-255/0/2  |
|        |        |        |                | 5      | 5      | vcp-255/0/3  |
|        |        |        |                | 6      | 6      | vcp-255/0/1  |
| 2      | 2      | Prsnt  | 100e.7eb5.d700 | 4      | 4      | vcp-255/0/2  |
|        |        |        |                | 5      | 5      | vcp-255/0/3  |
|        |        |        |                | 6      | 6      | vcp-255/0/1  |
| 3      | 3      | Prsnt  | 100e.7eb5.c440 | 4      | 4      | vcp-255/0/2  |
|        |        |        |                | 5      | 5      | vcp-255/0/3  |
|        |        |        |                | 6      | 6      | vcp-255/0/1  |
| 4      | 4      | Prsnt  | 100e.7eb5.7e40 | 3      | 3      | vcp-255/0/51 |
|        |        |        |                | 2      | 2      | vcp-255/0/50 |
|        |        |        |                | 0      | 0      | vcp-255/0/48 |
|        |        |        |                | 1      | 1      | vcp-255/0/49 |
| 5      | 5      | Prsnt  | 100e.7eb5.80c0 | 3      | 3      | vcp-255/0/51 |
|        |        |        |                | 2      | 2      | vcp-255/0/50 |
|        |        |        |                | 1      | 1      | vcp-255/0/49 |
|        |        |        |                | 0      | 0      | vcp-255/0/48 |
| 6      | 6      | Prsnt  | 100e.7eb6.3b00 | 3      | 3      | vcp-255/0/3  |
|        |        |        |                | 2      | 2      | vcp-255/0/2  |
|        |        |        |                | 0      | 0      | vcp-255/0/0  |
|        |        |        |                | 1      | 1      | vcp-255/0/1  |

fpc2:

| Neighbor List |        |        |                |        |        |              |
|---------------|--------|--------|----------------|--------|--------|--------------|
| Member        | Device | Status | System ID      | Member | Device | Interface    |
| 0             | 0      | Prsnt  | 100e.7eb6.a900 | 4      | 4      | vcp-255/0/2  |
|               |        |        |                | 5      | 5      | vcp-255/0/3  |
|               |        |        |                | 6      | 6      | vcp-255/0/1  |
| 1             | 1      | Prsnt  | 100e.7eb8.3a40 | 4      | 4      | vcp-255/0/2  |
|               |        |        |                | 5      | 5      | vcp-255/0/3  |
|               |        |        |                | 6      | 6      | vcp-255/0/1  |
| 2             | 2      | Prsnt  | 100e.7eb5.d700 | 4      | 4      | vcp-255/0/2  |
|               |        |        |                | 5      | 5      | vcp-255/0/3  |
|               |        |        |                | 6      | 6      | vcp-255/0/1  |
| 3             | 3      | Prsnt  | 100e.7eb5.c440 | 4      | 4      | vcp-255/0/2  |
|               |        |        |                | 5      | 5      | vcp-255/0/3  |
|               |        |        |                | 6      | 6      | vcp-255/0/1  |
| 4             | 4      | Prsnt  | 100e.7eb5.7e40 | 3      | 3      | vcp-255/0/51 |
|               |        |        |                | 2      | 2      | vcp-255/0/50 |
|               |        |        |                | 0      | 0      | vcp-255/0/48 |
|               |        |        |                | 1      | 1      | vcp-255/0/49 |
| 5             | 5      | Prsnt  | 100e.7eb5.80c0 | 3      | 3      | vcp-255/0/51 |
|               |        |        |                | 2      | 2      | vcp-255/0/50 |
|               |        |        |                | 1      | 1      | vcp-255/0/49 |
|               |        |        |                | 0      | 0      | vcp-255/0/48 |
| 6             | 6      | Prsnt  | 100e.7eb6.3b00 | 3      | 3      | vcp-255/0/3  |
|               |        |        |                | 2      | 2      | vcp-255/0/2  |
|               |        |        |                | 0      | 0      | vcp-255/0/0  |
|               |        |        |                | 1      | 1      | vcp-255/0/1  |

fpc3:

| Neighbor List |        |        |                |        |        |             |
|---------------|--------|--------|----------------|--------|--------|-------------|
| Member        | Device | Status | System ID      | Member | Device | Interface   |
| 0             | 0      | Prsnt  | 100e.7eb6.a900 | 4      | 4      | vcp-255/0/2 |
|               |        |        |                | 5      | 5      | vcp-255/0/3 |

|   |   |       |                |   |   |              |
|---|---|-------|----------------|---|---|--------------|
| 1 | 1 | Prsnt | 100e.7eb8.3a40 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 4 | 4 | vcp-255/0/2  |
|   |   |       |                | 5 | 5 | vcp-255/0/3  |
| 2 | 2 | Prsnt | 100e.7eb5.d700 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 4 | 4 | vcp-255/0/2  |
|   |   |       |                | 5 | 5 | vcp-255/0/3  |
| 3 | 3 | Prsnt | 100e.7eb5.c440 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 4 | 4 | vcp-255/0/2  |
|   |   |       |                | 5 | 5 | vcp-255/0/3  |
| 4 | 4 | Prsnt | 100e.7eb5.7e40 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 3 | 3 | vcp-255/0/51 |
|   |   |       |                | 2 | 2 | vcp-255/0/50 |
| 5 | 5 | Prsnt | 100e.7eb5.80c0 | 0 | 0 | vcp-255/0/48 |
|   |   |       |                | 1 | 1 | vcp-255/0/49 |
|   |   |       |                | 3 | 3 | vcp-255/0/51 |
|   |   |       |                | 2 | 2 | vcp-255/0/50 |
|   |   |       |                | 1 | 1 | vcp-255/0/49 |
| 6 | 6 | Prsnt | 100e.7eb6.3b00 | 0 | 0 | vcp-255/0/48 |
|   |   |       |                | 3 | 3 | vcp-255/0/3  |
|   |   |       |                | 2 | 2 | vcp-255/0/2  |
|   |   |       |                | 0 | 0 | vcp-255/0/0  |
|   |   |       |                | 1 | 1 | vcp-255/0/1  |

fpc4:

|        |        |        |                | Neighbor List |        |              |
|--------|--------|--------|----------------|---------------|--------|--------------|
| Member | Device | Status | System ID      | Member        | Device | Interface    |
| 0      | 0      | Prsnt  | 100e.7eb6.a900 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 1      | 1      | Prsnt  | 100e.7eb8.3a40 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 2      | 2      | Prsnt  | 100e.7eb5.d700 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 3      | 3      | Prsnt  | 100e.7eb5.c440 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 4      | 4      | Prsnt  | 100e.7eb5.7e40 | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
| 5      | 5      | Prsnt  | 100e.7eb5.80c0 | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
|        |        |        |                | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
| 6      | 6      | Prsnt  | 100e.7eb6.3b00 | 3             | 3      | vcp-255/0/3  |
|        |        |        |                | 2             | 2      | vcp-255/0/2  |
|        |        |        |                | 0             | 0      | vcp-255/0/0  |
|        |        |        |                | 1             | 1      | vcp-255/0/1  |

fpc5:

|        |        |        |                | Neighbor List |        |             |
|--------|--------|--------|----------------|---------------|--------|-------------|
| Member | Device | Status | System ID      | Member        | Device | Interface   |
| 0      | 0      | Prsnt  | 100e.7eb6.a900 | 4             | 4      | vcp-255/0/2 |
|        |        |        |                | 5             | 5      | vcp-255/0/3 |
|        |        |        |                | 6             | 6      | vcp-255/0/1 |
| 1      | 1      | Prsnt  | 100e.7eb8.3a40 | 4             | 4      | vcp-255/0/2 |
|        |        |        |                | 5             | 5      | vcp-255/0/3 |

|   |   |       |                |   |   |              |
|---|---|-------|----------------|---|---|--------------|
| 2 | 2 | Prsnt | 100e.7eb5.d700 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 4 | 4 | vcp-255/0/2  |
|   |   |       |                | 5 | 5 | vcp-255/0/3  |
| 3 | 3 | Prsnt | 100e.7eb5.c440 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 4 | 4 | vcp-255/0/2  |
|   |   |       |                | 5 | 5 | vcp-255/0/3  |
| 4 | 4 | Prsnt | 100e.7eb5.7e40 | 6 | 6 | vcp-255/0/1  |
|   |   |       |                | 3 | 3 | vcp-255/0/51 |
|   |   |       |                | 2 | 2 | vcp-255/0/50 |
| 5 | 5 | Prsnt | 100e.7eb5.80c0 | 0 | 0 | vcp-255/0/48 |
|   |   |       |                | 1 | 1 | vcp-255/0/49 |
|   |   |       |                | 3 | 3 | vcp-255/0/51 |
| 6 | 6 | Prsnt | 100e.7eb6.3b00 | 2 | 2 | vcp-255/0/50 |
|   |   |       |                | 1 | 1 | vcp-255/0/49 |
|   |   |       |                | 0 | 0 | vcp-255/0/48 |
|   |   |       |                | 3 | 3 | vcp-255/0/3  |
|   |   |       |                | 2 | 2 | vcp-255/0/2  |
|   |   |       |                | 0 | 0 | vcp-255/0/0  |
|   |   |       |                | 1 | 1 | vcp-255/0/1  |

fpc6:

|        |        |        |                | Neighbor List |        |              |
|--------|--------|--------|----------------|---------------|--------|--------------|
| Member | Device | Status | System ID      | Member        | Device | Interface    |
| 0      | 0      | Prsnt  | 100e.7eb6.a900 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 1      | 1      | Prsnt  | 100e.7eb8.3a40 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 2      | 2      | Prsnt  | 100e.7eb5.d700 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 3      | 3      | Prsnt  | 100e.7eb5.c440 | 4             | 4      | vcp-255/0/2  |
|        |        |        |                | 5             | 5      | vcp-255/0/3  |
|        |        |        |                | 6             | 6      | vcp-255/0/1  |
| 4      | 4      | Prsnt  | 100e.7eb5.7e40 | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
| 5      | 5      | Prsnt  | 100e.7eb5.80c0 | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 3             | 3      | vcp-255/0/51 |
|        |        |        |                | 2             | 2      | vcp-255/0/50 |
| 6      | 6      | Prsnt  | 100e.7eb6.3b00 | 1             | 1      | vcp-255/0/49 |
|        |        |        |                | 0             | 0      | vcp-255/0/48 |
|        |        |        |                | 3             | 3      | vcp-255/0/3  |
|        |        |        |                | 2             | 2      | vcp-255/0/2  |
|        |        |        |                | 0             | 0      | vcp-255/0/0  |
|        |        |        |                | 1             | 1      | vcp-255/0/1  |

## show virtual-chassis login

|                                 |                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show virtual-chassis login</b>                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                |
| <b>Description</b>              | <p>Supply the address of the host that logged into the Virtual Chassis or VCF, or identify the location of the member switch that redirected the current session to a different member switch.</p> <p>You might need this information for tracing or troubleshooting purposes.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">request session member on page 5163</a></li> <li>• <a href="#">Understanding Global Management of a Virtual Chassis on page 5080</a></li> </ul>                                                                               |
| <b>List of Sample Output</b>    | <p><a href="#">show virtual-chassis login (Direct Login to the Master Console Port) on page 5341</a></p> <p><a href="#">show virtual-chassis login (Backup Console Session Redirected to the Master Console Port) on page 5341</a></p>                                             |

### Sample Output

#### show virtual-chassis login (Direct Login to the Master Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host 248.1.2.3
```

#### show virtual-chassis login (Backup Console Session Redirected to the Master Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host backup
```

## show virtual-chassis mode

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show virtual-chassis mode</b><br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).<br><b>Current mode</b> and <b>Future mode after reboot</b> fields introduced in Junos OS Release 13.2X51-D20.                                                                                                                                                                                             |
| <b>Description</b>              | Display the Virtual Chassis or Virtual Chassis Fabric (VCF) mixed mode status.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>none</b>—Display the Virtual Chassis or VCF mixed mode status for the device on which the command is entered.</p> <p><b>all-members</b>—(Optional) Display the Virtual Chassis or VCF mixed mode status for all member devices in the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the Virtual Chassis or VCF mixed mode status for the device on which the command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the Virtual Chassis or VCF mixed mode status for the specified member device..</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">request virtual-chassis mode on page 5320</a></li> <li>• <a href="#">Verifying the Virtual Chassis Fabric Mode Settings on page 5315</a></li> <li>• <a href="#">Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)</a></li> </ul>                                                                                                                                                                                                             |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis mode (EX4200) on page 5343</a><br><a href="#">show virtual-chassis mode (QFX5100) on page 5343</a>                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>            | <a href="#">Table 589 on page 5342</a> lists the output fields for the <b>show virtual-chassis mode</b> command.                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Table 589: show virtual-chassis mode Output Fields

| Field Name        | Field Description                                                                                              |
|-------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Mixed Mode</b> | Specifies the mixed mode status of the member switch. Mixed mode is either <b>Enabled</b> or <b>Disabled</b> . |

Table 589: show virtual-chassis mode Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Current mode</b>             | <p>Specifies the current mixed and fabric mode settings running on the member device or devices.</p> <p>A device reboot is required to change the fabric or mixed mode. The <b>Current mode</b> and <b>Future mode after reboot</b> are different when the mode has been changed but the device has not been rebooted.</p> <p>Outputs include:</p> <ul style="list-style-type: none"> <li>• <b>Fabric with mixed devices</b>—Fabric mode and mixed mode are enabled.</li> <li>• <b>Fabric with similar devices</b>—Fabric mode is enabled and mixed mode is disabled.</li> <li>• <b>Virtual Chassis with mixed devices</b>—Fabric mode is disabled and mixed mode is enabled.</li> <li>• <b>Virtual Chassis with similar devices</b>—Fabric mode is disabled and mixed mode is disabled.</li> </ul> |
| <b>Future mode after reboot</b> | <p>Specifies the mixed and fabric mode settings running on the member device or devices.</p> <p>A device reboot is required to change the fabric or mixed mode. The <b>Current mode</b> and <b>Future mode after reboot</b> are different when the mode has been changed but the device has not been rebooted.</p> <p>Outputs include:</p> <ul style="list-style-type: none"> <li>• <b>Fabric with mixed devices</b>—Fabric mode and mixed mode are enabled.</li> <li>• <b>Fabric with similar devices</b>—Fabric mode is enabled and mixed mode is disabled.</li> <li>• <b>Virtual Chassis with mixed devices</b>—Fabric mode is disabled and mixed mode is enabled.</li> <li>• <b>Virtual Chassis with similar devices</b>—Fabric mode is disabled and mixed mode is disabled.</li> </ul>         |

## Sample Output

### show virtual-chassis mode (EX4200)

```
user@switch>show virtual-chassis mode
fpc0:
-----
Mixed Mode: Disabled
```

## Sample Output

### show virtual-chassis mode (QFX5100)

```
user@switch>show virtual-chassis mode
fpc0:
-----
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc1:
-----
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc2:
-----
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc3:
```

-----  
Current mode : Fabric with similar devices  
Future mode after reboot : Fabric with similar devices

fpc4:

-----  
Current mode : Fabric with similar devices  
Future mode after reboot : Fabric with similar devices



## show virtual-chassis protocol adjacency

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show virtual-chassis protocol adjacency &lt;brief   detail   extensive&gt; &lt;all-members&gt; &lt;local&gt; &lt;member member-id&gt; &lt;system-id&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Display the Virtual Chassis Control Protocol (VCCP) adjacency statistics in the Virtual Chassis or VCF for all hardware devices.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>none</b>—Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> and <b>extensive</b> options provide identical displays.</p> <p><b>all-members</b>—(Optional) Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display VCCP adjacency statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.</p> <p><b>system-id</b>—(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding EX Series Virtual Chassis Port Link Aggregation on page 5083</a></li> <li>• <a href="#">Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>    | <p><a href="#">show virtual-chassis protocol adjacency on page 5346</a></p> <p><a href="#">show virtual-chassis protocol adjacency detail on page 5347</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | Table 574 on page 5187 lists the output fields for the <b>show virtual-chassis protocol adjacency</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 590: show virtual-chassis protocol adjacency Output Fields

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                        | Level of Output |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>           | Name of the Virtual Chassis port (VCP) interface.                                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>System</b>              | The MAC address of the device on the receiving side of the VCP link.                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>State</b>               | State of the link. Outputs include: <ul style="list-style-type: none"> <li>• <b>Up</b>—The link is up.</li> <li>• <b>Down</b>—The link is down.</li> <li>• <b>New</b>—The link is new.</li> <li>• <b>One-way</b>—The link is transmitting traffic in one direction.</li> <li>• <b>Initializing</b>—The link is initializing.</li> <li>• <b>Rejected</b>—The link is rejected.</li> </ul> | All levels      |
| <b>Hold, Expires in</b>    | Remaining holdtime of the adjacency.                                                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>Priority</b>            | Priority to become the designated intermediary system.                                                                                                                                                                                                                                                                                                                                   | detail          |
| <b>Up/Down Transitions</b> | Count of adjacency status transition changes from up to down or down to up.                                                                                                                                                                                                                                                                                                              | detail          |
| <b>Last transition</b>     | Time of the last up/down transition.                                                                                                                                                                                                                                                                                                                                                     | detail          |

## Sample Output

### show virtual-chassis protocol adjacency

```
user@switch> show virtual-chassis protocol adjacency
```

```
member0:
```

```
-----
Interface      System      State      Hold (secs)
vcp-0/0.32768  0000.4a75.9b7c Up          57
vcp-0/1.32768  0000.4a75.9b7c Up          59
vcp-4/0/1.32768 0026.888d.6800 Up          57
```

```
member1:
```

```
-----
Interface      System      State      Hold (secs)
vcp-0/0.32768  0000.4a75.9b7c Up          58
vcp-0/1.32768  0000.73e9.9a57 Up          59
vcp-3/0/4.32768 0021.59f7.d000 Up          58
```

```
member8:
```

```
-----
Interface      System      State      Hold (secs)
vcp-1/0.32768  0000.73e9.9a57 Up          58
vcp-1/1.32768  0021.59f7.d000 Up          58
vcp-1/2.32768  0026.888d.6800 Up          59
vcp-2/0.32768  0021.59f7.d000 Up          59
```

```
member9:
```

```
-----
Interface      System      State      Hold (secs)
```

|               |                   |    |
|---------------|-------------------|----|
| vcp-1/0.32768 | 0000.4a75.9b7c Up | 58 |
| vcp-1/1.32768 | 0026.888d.6800 Up | 59 |

### show virtual-chassis protocol adjacency detail

```
user@switch> show virtual-chassis protocol adjacency detail
```

```
member0:
```

```
-----
0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago
```

```
0000.4a75.9b7c
  interface-name: vcp-0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago
```

```
0026.888d.6800
  interface-name: vcp-4/0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago
```

```
member1:
```

```
-----
0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0000.73e9.9a57
  interface-name: vcp-0/1.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago
```

```
0021.59f7.d000
  interface-name: vcp-3/0/4.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago
```

```
member8:
```

```
-----
0000.73e9.9a57
  interface-name: vcp-1/0.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0021.59f7.d000
  interface-name: vcp-1/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0026.888d.6800
  interface-name: vcp-1/2.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0021.59f7.d000
  interface-name: vcp-2/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
member9:
```

```
-----
0000.4a75.9b7c
  interface-name: vcp-1/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0026.888d.6800
  interface-name: vcp-1/1.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago
```

## show virtual-chassis protocol database

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis protocol database<br><brief   detail   extensive><br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display the Virtual Chassis Control Protocol (VCCP) database statistics for all hardware devices within the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>none</b>—Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> option provides more output than the <b>brief</b> option. The <b>extensive</b> option provides all output and is most useful for customer support personnel.</p> <p><b>all-members</b>—(Optional) Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display VCCP database statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display VCCP database statistics for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</a></li> <li>• <a href="#">Understanding EX Series Virtual Chassis Components on page 5067</a></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis Components</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis protocol database on page 5350</a><br><a href="#">show virtual-chassis protocol database detail on page 5351</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | <a href="#">Table 575 on page 5190</a> lists the output fields for the <b>show virtual-chassis protocol database</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 591: show virtual-chassis protocol database Output Fields

| Field Name | Field Description                               | Level of Output |
|------------|-------------------------------------------------|-----------------|
| LSP ID     | Link-state protocol (LSP) data unit identifier. | All levels      |

Table 591: show virtual-chassis protocol database Output Fields (*continued*)

| Field Name       | Field Description                                      | Level of Output |
|------------------|--------------------------------------------------------|-----------------|
| <b>Sequence</b>  | Sequence number of the LSP.                            | All levels      |
| <b>Checksum</b>  | Checksum value of the LSP.                             | All levels      |
| <b>Lifetime</b>  | Remaining lifetime of the LSP, in seconds.             | All levels      |
| <b>Neighbor</b>  | MAC address of the neighbor on the advertising system. | detail          |
| <b>Interface</b> | Virtual Chassis port (VCP) interface name.             | detail          |
| <b>Metric</b>    | Metric of the prefix or neighbor.                      | detail          |

The **extensive** output was omitted from this list. The **extensive** output is useful for customer support personnel only.

## Sample Output

### show virtual-chassis protocol database

```
user@switch> show virtual-chassis protocol database
```

```
member0:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361  0x27e8   113
0021.59f7.d000.00-00  0x16882  0x3993   118
0026.888d.6800.00-00  0x1691f  0x82b7   116
  4 LSPs
```

```
member1:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361  0x27e8   114
0021.59f7.d000.00-00  0x16883  0x289    116
0026.888d.6800.00-00  0x1691f  0x82b7   118
  4 LSPs
```

```
member8:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   118
0000.73e9.9a57.00-00  0xf361  0x27e8   114
0021.59f7.d000.00-00  0x16883  0x289    116
0026.888d.6800.00-00  0x16920  0xa335   116
  4 LSPs
```

```
member9:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361  0x27e8   116
0021.59f7.d000.00-00  0x16883  0x289    114
```

```
0026.888d.6800.00-00      0x16920   0xa335      116
4 LSPs
```

### show virtual-chassis protocol database detail

```
user@switch> show virtual-chassis protocol database detail
member0:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 114 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 118 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15
```

```
0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15
```

```
member1:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15
```

```
0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 117 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15
```

```
member8:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 118 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
```

```
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15

0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15
```

member9:

```
-----

0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 116 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150

0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 117 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 113 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15

0026.888d.6800.00-00 Sequence: 0x1694f, Checksum: 0xa61a, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15
```



## show virtual-chassis protocol interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show virtual-chassis protocol interface &lt;brief   detail&gt; &lt;all-members&gt; &lt;interface-name&gt; &lt;local&gt; &lt;member member-id&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Display information about Virtual Chassis Control Protocol (VCCP) statistics for VCCP-enabled interfaces within the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>none</b>—Display the VCCP interface statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail</b> —(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> option provides more output than the <b>brief</b> option.</p> <p><b>all-members</b>—(Optional) Display VCCP interface statistics for all members of the Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display VCCP interface statistics for the specified interface.</p> <p><b>local</b>—(Optional) Display VCCP interface statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display VCCP interface statistics for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>EX Series Virtual Chassis Overview</i></li> <li>• <i>Understanding QFX Series Virtual Chassis</i></li> <li>• <i>Understanding Virtual Chassis Ports in an EX8200 Virtual Chassis</i></li> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis protocol interface on page 5354</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | <a href="#">Table 576 on page 5195</a> lists the output fields for the <b>show virtual-chassis protocol interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 592: show virtual-chassis protocol interface Output Fields

| Field Name       | Field Description                                                                                                                                          | Level of Output |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b> | Name of the VCP.                                                                                                                                           | All levels      |
| <b>State</b>     | State of the link. Outputs include: <ul style="list-style-type: none"> <li>• <b>Up</b>—The link is up.</li> <li>• <b>Down</b>—The link is down.</li> </ul> | All levels      |
| <b>Metric</b>    | Metric of the prefix or neighbor.                                                                                                                          | All levels      |

## Sample Output

### show virtual-chassis protocol interface

```
user@switch> show virtual-chassis protocol interface
```

```
member0:
```

```
-----
IS-IS interface database:
```

| Interface       | State | Metric |
|-----------------|-------|--------|
| vcp-0/0.32768   | Up    | 150    |
| vcp-0/1.32768   | Up    | 150    |
| vcp-4/0/1.32768 | Up    | 15     |
| vcp-4/0/7.32768 | Down  | 15     |

```
member1:
```

```
-----
IS-IS interface database:
```

| Interface       | State | Metric |
|-----------------|-------|--------|
| vcp-0/0.32768   | Up    | 150    |
| vcp-0/1.32768   | Up    | 150    |
| vcp-3/0/4.32768 | Up    | 15     |

```
member8:
```

```
-----
IS-IS interface database:
```

| Interface     | State | Metric |
|---------------|-------|--------|
| vcp-0/0.32768 | Down  | 150    |
| vcp-1/0.32768 | Up    | 150    |
| vcp-1/1.32768 | Up    | 150    |
| vcp-1/2.32768 | Up    | 150    |
| vcp-1/3.32768 | Down  | 150    |
| vcp-2/0.32768 | Up    | 150    |
| vcp-2/1.32768 | Down  | 150    |
| vcp-2/2.32768 | Down  | 150    |
| vcp-2/3.32768 | Down  | 150    |

```
member9:
```

```
-----
IS-IS interface database:
```

| Interface     | State | Metric |
|---------------|-------|--------|
| vcp-0/0.32768 | Down  | 150    |
| vcp-1/0.32768 | Up    | 150    |
| vcp-1/1.32768 | Up    | 150    |
| vcp-1/2.32768 | Down  | 150    |
| vcp-1/3.32768 | Down  | 150    |



## show virtual-chassis protocol route

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis protocol route<br><all-members><br><destination-id><br><local><br><member member-id>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display the unicast and multicast Virtual Chassis Control Protocol (VCCP) routing tables within the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Display the unicast and multicast routing tables for all members of the Virtual Chassis.</p> <p><b>all-members</b>—(Optional) Display the unicast and multicast routing tables for all members of the Virtual Chassis or VCF.</p> <p><b>destination-id</b>—(Optional) Display the unicast and multicast routing tables to the specified destination member ID for each member of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the unicast and multicast routing tables on the device where this command is entered.</p> <p><b>member member-id</b>—(Optional) Display the unicast and multicast routing tables for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>EX Series Virtual Chassis Overview</i></li> <li>• <i>Understanding QFX Series Virtual Chassis</i></li> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis protocol route on page 5357</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | <a href="#">Table 577 on page 5197</a> lists the output fields for the <b>show virtual-chassis protocol route</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

**Table 593: show virtual-chassis protocol route Output Fields**

| Field Name     | Field Description                                                              |
|----------------|--------------------------------------------------------------------------------|
| <b>Dev</b>     | MAC address of the member storing the VCCP routing table.                      |
| <b>Version</b> | Version of the shortest-path-first algorithm that generated the routing table. |

Table 593: show virtual-chassis protocol route Output Fields (*continued*)

| Field Name       | Field Description                                                            |
|------------------|------------------------------------------------------------------------------|
| <b>System ID</b> | MAC address of the device.                                                   |
| <b>Version</b>   | Version of the shortest-path-first (SPF) algorithm that generated the route. |
| <b>Metric</b>    | The metric number to get to that device.                                     |
| <b>Interface</b> | Name of the Virtual Chassis port (VCP) interface connecting the devices.     |
| <b>Via</b>       | MAC address of the next-hop device, if applicable.                           |

## Sample Output

### show virtual-chassis protocol route

```

user@switch> show virtual-chassis protocol route
member0:
-----
Dev 0021.59f7.d000 ucast routing table           Current version: 21
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    21      150 vcp-0/1.32768 0000.4a75.9b7c
0000.73e9.9a57    21      165 vcp-4/0/1.32768 0026.888d.6800
0021.59f7.d000    21        0
0026.888d.6800    21      15 vcp-4/0/1.32768 0026.888d.6800

Dev 0021.59f7.d000 mcast routing table           Current version: 21
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    21
0000.73e9.9a57    21
0021.59f7.d000    21          vcp-4/0/1.32768
                   vcp-0/1.32768
0026.888d.6800    21

member1:
-----
Dev 0026.888d.6800 ucast routing table           Current version: 25
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    25      150 vcp-0/0.32768 0000.4a75.9b7c
0000.73e9.9a57    25      150 vcp-0/1.32768 0000.73e9.9a57
0021.59f7.d000    25        15 vcp-3/0/4.32768 0021.59f7.d000
0026.888d.6800    25        0

Dev 0026.888d.6800 mcast routing table           Current version: 25
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    25
0000.73e9.9a57    25          vcp-3/0/4.32768
0021.59f7.d000    25          vcp-0/1.32768
0026.888d.6800    25          vcp-3/0/4.32768
                   vcp-0/0.32768

```

vcp-0/1.32768

member8:

-----

Dev 0000.4a75.9b7c ucast routing table Current version: 39

-----

| System ID      | Version | Metric | Interface     | Via            |
|----------------|---------|--------|---------------|----------------|
| 0000.4a75.9b7c | 39      | 0      |               |                |
| 0000.73e9.9a57 | 39      | 150    | vcp-1/0.32768 | 0000.73e9.9a57 |
| 0021.59f7.d000 | 39      | 150    | vcp-2/0.32768 | 0021.59f7.d000 |
| 0026.888d.6800 | 39      | 150    | vcp-1/2.32768 | 0026.888d.6800 |

Dev 0000.4a75.9b7c mcast routing table Current version: 39

-----

| System ID      | Version | Metric | Interface     | Via |
|----------------|---------|--------|---------------|-----|
| 0000.4a75.9b7c | 39      |        | vcp-1/0.32768 |     |
|                |         |        | vcp-2/0.32768 |     |
|                |         |        | vcp-1/2.32768 |     |
| 0000.73e9.9a57 | 39      |        |               |     |
| 0021.59f7.d000 | 39      |        |               |     |
| 0026.888d.6800 | 39      |        |               |     |

member9:

-----

Dev 0000.73e9.9a57 ucast routing table Current version: 31

-----

| System ID      | Version | Metric | Interface     | Via            |
|----------------|---------|--------|---------------|----------------|
| 0000.4a75.9b7c | 31      | 150    | vcp-1/0.32768 | 0000.4a75.9b7c |
| 0000.73e9.9a57 | 31      | 0      |               |                |
| 0021.59f7.d000 | 31      | 165    | vcp-1/1.32768 | 0026.888d.6800 |
| 0026.888d.6800 | 31      | 150    | vcp-1/1.32768 | 0026.888d.6800 |

Dev 0000.73e9.9a57 mcast routing table Current version: 31

-----

| System ID      | Version | Metric | Interface     | Via |
|----------------|---------|--------|---------------|-----|
| 0000.4a75.9b7c | 31      |        |               |     |
| 0000.73e9.9a57 | 31      |        | vcp-1/0.32768 |     |
|                |         |        | vcp-1/1.32768 |     |
| 0021.59f7.d000 | 31      |        |               |     |
| 0026.888d.6800 | 31      |        |               |     |

## show virtual-chassis protocol statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis protocol statistics<br><all-members><br><interface-name><br><local><br><member member-id>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display the Virtual Chassis Control Protocol (VCCP) statistics for all hardware devices within the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>none</b>—Display VCCP statistics for all members of the Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Display VCCP statistics for all members of the Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display VCCP statistics for the specified interface.</p> <p><b>local</b>—(Optional) Display VCCP statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display VCCP statistics for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>EX Series Virtual Chassis Overview</i></li> <li>• <i>Understanding QFX Series Virtual Chassis</i></li> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> </ul>                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis protocol statistics on page 5360</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>            | <a href="#">Table 578 on page 5200</a> lists the output fields for the <b>show virtual-chassis protocol interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                          |

**Table 594: show virtual-chassis protocol statistics Output Fields**

| Field Name       | Field Description                                                                    |
|------------------|--------------------------------------------------------------------------------------|
| <b>PDU type</b>  | Protocol data unit type.                                                             |
| <b>Received</b>  | Number of PDUs received since VCCP started or since the statistics were set to zero. |
| <b>Processed</b> | Number of PDUs received minus the number of PDUs dropped.                            |

Table 594: show virtual-chassis protocol statistics Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                        |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Drops</b>                  | Number of PDUs dropped.                                                                                                                                  |
| <b>Sent</b>                   | Number of PDUs transmitted since VCCP started or since the statistics were set to zero.                                                                  |
| <b>Rexmit</b>                 | Number of PDUs retransmitted since VCCP started or since the statistics were set to zero.                                                                |
| <b>Total Packets Received</b> | Number of PDUs received since VCCP started or since the statistics were set to zero.                                                                     |
| <b>Total Packets Sent</b>     | Number of PDUs sent since VCCP started or since the statistics were set to zero.                                                                         |
| <b>LSP queue length</b>       | Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.                                                            |
| <b>SPF runs</b>               | Number of shortest-path-first (SPF) calculations that have been performed.                                                                               |
| <b>Fragments Rebuilt</b>      | Number of link-state PDU fragments that the local system has computed.                                                                                   |
| <b>LSP Regenerations</b>      | Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed. |
| <b>Purges initiated</b>       | Number of purges that the system initiated. A purge is initiated if the software determines that a link-state PDU must be removed from the network.      |

## Sample Output

### show virtual-chassis protocol statistics

```

user@switch> show virtual-chassis protocol statistics
member0:
-----
IS-IS statistics for 0021.59f7.d000:
PDU type      Received    Processed      Drops      Sent      Rexmit
LSP            8166        8166           0         4551         0
HELLO          1659        1659           0         1693         0
CSNP             2            2             0            3         0
PSNP           1909        1909           0         2293         0
Unknown         0            0             0            0         0
Totals        11736       11736           0         8540         0

Total packets received: 11736 Sent: 8540

LSP queue length: 0 Drops: 0
SPF runs: 9
Fragments rebuilt: 1640
LSP regenerations: 1
Purges initiated: 0

member1:
-----
IS-IS statistics for 0026.888d.6800:

```



| PDU type | Received | Processed | Drops | Sent  | Rexmit |
|----------|----------|-----------|-------|-------|--------|
| LSP      | 10909    | 10909     | 0     | 12088 | 0      |
| HELLO    | 1877     | 1877      | 0     | 2251  | 0      |
| CSNP     | 3        | 3         | 0     | 3     | 0      |
| PSNP     | 3846     | 3846      | 0     | 3732  | 0      |
| Unknown  | 0        | 0         | 0     | 0     | 0      |
| Totals   | 16635    | 16635     | 0     | 18074 | 0      |

Total packets received: 16635 Sent: 18074

LSP queue length: 0 Drops: 0  
 SPF runs: 13  
 Fragments rebuilt: 1871  
 LSP regenerations: 2  
 Purges initiated: 0

member8:

IS-IS statistics for 0000.4a75.9b7c:

| PDU type | Received | Processed | Drops | Sent  | Rexmit |
|----------|----------|-----------|-------|-------|--------|
| LSP      | 7935     | 7935      | 0     | 14865 | 0      |
| HELLO    | 2695     | 2695      | 0     | 7124  | 0      |
| CSNP     | 4        | 4         | 0     | 4     | 0      |
| PSNP     | 4398     | 4398      | 0     | 3666  | 0      |
| Unknown  | 0        | 0         | 0     | 0     | 0      |
| Totals   | 15032    | 15032     | 0     | 25659 | 0      |

Total packets received: 15032 Sent: 25659

LSP queue length: 0 Drops: 0  
 SPF runs: 26  
 Fragments rebuilt: 2666  
 LSP regenerations: 4  
 Purges initiated: 0

member9:

IS-IS statistics for 0000.73e9.9a57:

| PDU type | Received | Processed | Drops | Sent  | Rexmit |
|----------|----------|-----------|-------|-------|--------|
| LSP      | 10800    | 10800     | 0     | 6327  | 0      |
| HELLO    | 1492     | 1492      | 0     | 2356  | 0      |
| CSNP     | 2        | 2         | 0     | 2     | 0      |
| PSNP     | 2683     | 2683      | 0     | 3149  | 0      |
| Unknown  | 0        | 0         | 0     | 0     | 0      |
| Totals   | 14977    | 14977     | 0     | 11834 | 0      |

Total packets received: 14977 Sent: 11834

LSP queue length: 0 Drops: 0  
 SPF runs: 19  
 Fragments rebuilt: 1510  
 LSP regenerations: 6  
 Purges initiated: 0

## show virtual-chassis

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show virtual-chassis</b><br><b>&lt;status&gt;</b>                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p> <p><b>Fabric ID</b>, <b>Fabric Mode</b>, and <b>Route Mode</b> output fields introduced in Junos OS Release 13.2X51-D20.</p> <p><b>Alias-Name</b> output field introduced in Junos OS Release 14.1X53-D10.</p> |
| <b>Description</b>              | Display information about all members of the Virtual Chassis or VCF.                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b>none</b>—Display information about all Virtual Chassis or VCF member devices.</p> <p><b>status</b>—Same output as for <b>show virtual-chassis</b>.</p>                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis active-topology on page 5174</a></li> <li>• <a href="#">show virtual-chassis protocol adjacency on page 5186</a></li> <li>• <a href="#">show virtual-chassis vc-path on page 5207</a></li> <li>• <a href="#">Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157</a></li> </ul>                                                               |
| <b>List of Sample Output</b>    | <p><a href="#">show virtual-chassis (EX4200 Virtual Chassis) on page 5364</a></p> <p><a href="#">show virtual-chassis (EX8200 Virtual Chassis) on page 5364</a></p> <p><a href="#">show virtual-chassis (Virtual Chassis Fabric) on page 5365</a></p>                                                                                                                                                                                                              |
| <b>Output Fields</b>            | <p><a href="#">Table 579 on page 5203</a> lists the output fields for the <b>show virtual-chassis</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                         |

**Table 595: show virtual-chassis Output Fields**

| Field Name                  | Field Description                                                |
|-----------------------------|------------------------------------------------------------------|
| <b>Fabric ID</b>            | Assigned ID used to identify the VCF.                            |
| <b>Fabric Mode</b>          | Mode of the VCF: Enabled, Disabled, or Mixed.                    |
| <b>Virtual Chassis ID</b>   | Assigned ID that applies to the entire Virtual Chassis or VCF.   |
| <b>Virtual Chassis Mode</b> | Mode of the Virtual Chassis or VCF: Enabled, Disabled, or Mixed. |

Table 595: show virtual-chassis Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Member ID</b>           | Assigned member ID and FPC: <ul style="list-style-type: none"> <li>On all EX Series Virtual Chassis except EX8200 Virtual Chassis, and on a VCF, the FPC number refers to the member ID assigned to the switch.</li> <li>On EX8200 Virtual Chassis, member IDs are numbered 0 through 9. The FPC number indicates the slot number of the line card within the Virtual Chassis. The FPC number on member 0 is always 0 through 15. The FPC number on member 1 is always 16 through 31. The FPC number on member 2 is always 32 through 47; and so on for the members.</li> </ul>                                                                                                                         |
| <b>Status</b>              | For a nonprovisioned configuration: <ul style="list-style-type: none"> <li><b>Prsnt</b> for a member that is currently connected to the Virtual Chassis or VCF configuration.</li> <li><b>NotPrsnt</b> for a member ID that has been assigned but is not currently connected.</li> </ul> For a preprovisioned configuration: <ul style="list-style-type: none"> <li><b>Prsnt</b> for a member that is specified in the preprovisioned configuration file and is currently connected to the Virtual Chassis or VCF.</li> <li><b>Unprvsnd</b> for a member that is interconnected with the Virtual Chassis or VCF configuration but is not specified in the preprovisioned configuration file.</li> </ul> |
| <b>Serial No</b>           | Serial number of the member device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Alias-Name</b>          | The user-configured alias of the member device.<br><br>The <b>Alias-Name</b> field appears only if an alias has been configured for at least one device in the Virtual Chassis or VCF. Aliases are configured using the <b>alias-name</b> statement in the <code>[edit virtual-chassis aliases serial-number serial-number]</code> hierarchy.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Model</b>               | Model number of the member device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Mastership Priority</b> | Mastership priority value of the member device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Role</b>                | Role of the member device: master, backup, or linecard.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Mixed Mode</b>          | Mixed mode configuration status: <ul style="list-style-type: none"> <li><b>Y</b> for a member device configured in mixed mode.</li> <li><b>N</b> for a member device not configured in mixed mode.</li> <li><b>NA</b> for a member device that cannot be configured in mixed mode.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Route Mode</b>          | The route mode of the member device: fabric (F) or Virtual Chassis (V).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Location</b>            | Location of the member device.<br><br>If this field is empty, the location field was not set for the device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Neighbor List</b>       | Member ID of the neighbor member to which this member's Virtual Chassis port (VCP) is connected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## Sample Output

### show virtual-chassis (EX4200 Virtual Chassis)

```

user@switch> show virtual-chassis
Virtual Chassis ID: 0019.e250.47a0
Virtual Chassis Mode: Enabled

```

| Member ID | Status | Serial No    | Model      | Mastership<br>priority | Role     | Mixed<br>Mode | Neighbor List<br>ID | Interface |
|-----------|--------|--------------|------------|------------------------|----------|---------------|---------------------|-----------|
| 0 (FPC 0) | Prsnt  | AK0207360276 | ex4200-24t | 249                    | Master*  | N             | 8                   | vcp-0     |
|           |        |              |            |                        |          |               | 1                   | vcp-1     |
| 1 (FPC 1) | Prsnt  | AK0207360281 | ex4200-24t | 248                    | Backup   | N             | 0                   | vcp-0     |
|           |        |              |            |                        |          |               | 2                   | vcp-1     |
| 2 (FPC 2) | Prsnt  | AJ0207391130 | ex4200-48p | 247                    | Linecard | N             | 1                   | vcp-0     |
|           |        |              |            |                        |          |               | 3                   | vcp-1     |
| 3 (FPC 3) | Prsnt  | AK0207360280 | ex4200-24t | 246                    | Linecard | N             | 2                   | vcp-0     |
|           |        |              |            |                        |          |               | 4                   | vcp-1     |
| 4 (FPC 4) | Prsnt  | AJ0207391113 | ex4200-48p | 245                    | Linecard | N             | 3                   | vcp-0     |
|           |        |              |            |                        |          |               | 5                   | vcp-1     |
| 5 (FPC 5) | Prsnt  | BP0207452204 | ex4200-48t | 244                    | Linecard | N             | 4                   | vcp-0     |
|           |        |              |            |                        |          |               | 6                   | vcp-1     |
| 6 (FPC 6) | Prsnt  | BP0207452222 | ex4200-48t | 243                    | Linecard | N             | 5                   | vcp-0     |
|           |        |              |            |                        |          |               | 7                   | vcp-1     |
| 7 (FPC 7) | Prsnt  | BR0207432028 | ex4200-24f | 242                    | Linecard | N             | 6                   | vcp-0     |
|           |        |              |            |                        |          |               | 8                   | vcp-1     |
| 8 (FPC 8) | Prsnt  | BR0207431996 | ex4200-24f | 241                    | Linecard | N             | 7                   | vcp-0     |
|           |        |              |            |                        |          |               | 0                   | vcp-1     |

Member ID for next new member: 9 (FPC 9)

### show virtual-chassis (EX8200 Virtual Chassis)

```

user@external-routing-engine> show virtual-chassis
Virtual Chassis ID: c806.0842.de51
Virtual Chassis Mode: Enabled

```

| Member ID       | Status | Serial No    | Model  | Mastership<br>priority | Role     | Neighbor List<br>ID | Interface |
|-----------------|--------|--------------|--------|------------------------|----------|---------------------|-----------|
| 0 (FPC 0-15)    | Prsnt  | BA0908380001 | ex8216 | 0                      | Linecard | 8                   | vcp-0/0   |
|                 |        |              |        |                        |          | 8                   | vcp-0/1   |
|                 |        |              |        |                        |          | 1                   | vcp-4/0/4 |
| 1 (FPC 16-31)   | Prsnt  | BT0909411634 | ex8208 | 0                      | Linecard | 8                   | vcp-0/0   |
|                 |        |              |        |                        |          | 0                   | vcp-3/0/4 |
| 8 (FPC 128-143) | Prsnt  | 062009000021 | ex-xre | 128                    | Master   | 9                   | vcp-1/0   |
|                 |        |              |        |                        |          | 1                   | vcp-1/2   |

```

9 (FPC 144-159) Prsnt 062009000022 ex-xre 128 Backup*
9 vcp-1/3
0 vcp-2/0
9 vcp-2/1
0 vcp-1/1
8 vcp-1/0
8 vcp-1/2
8 vcp-1/3
8 vcp-1/3

```

### show virtual-chassis (Virtual Chassis Fabric)

```

user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 0282.5fa0.3f08
Fabric Mode: Enabled

```

| List        | Member ID | Status       | Serial No   | Model | Mstr<br>prio | Role | Mixed | Route | Neighbor |
|-------------|-----------|--------------|-------------|-------|--------------|------|-------|-------|----------|
| Interface   |           |              |             |       |              |      |       |       |          |
| 0 (FPC 0)   | Prsnt     | AB3112430001 | qfx5100-48s | 129   | Master*      | N    | F     | 3     |          |
| vcp-255/1/0 |           |              |             |       |              |      |       |       | 2        |
| vcp-255/1/1 |           |              |             |       |              |      |       |       | 4        |
| vcp-255/1/2 |           |              |             |       |              |      |       |       | 4        |
| vcp-255/1/3 |           |              |             |       |              |      |       |       | 4        |
| 1 (FPC 1)   | Prsnt     | AB3112230001 | qfx5100-48s | 129   | Backup       | N    | F     | 3     |          |
| vcp-255/1/0 |           |              |             |       |              |      |       |       | 2        |
| vcp-255/1/1 |           |              |             |       |              |      |       |       | 4        |
| vcp-255/1/2 |           |              |             |       |              |      |       |       | 4        |
| vcp-255/1/3 |           |              |             |       |              |      |       |       | 4        |
| 2 (FPC 2)   | Prsnt     | AB3112460011 | qfx5100-48s | 0     | Linecard     | N    | F     | 1     |          |
| vcp-255/1/0 |           |              |             |       |              |      |       |       | 0        |
| vcp-255/1/1 |           |              |             |       |              |      |       |       | 0        |
| 3 (FPC 3)   | Prsnt     | AB3112460011 | qfx5100-48s | 0     | Linecard     | N    | F     | 1     |          |
| vcp-255/1/0 |           |              |             |       |              |      |       |       | 0        |
| vcp-255/1/1 |           |              |             |       |              |      |       |       | 0        |
| 4 (FPC 4)   | Prsnt     | AB3112430011 | qfx5100-48s | 0     | Linecard     | N    | F     | 1     |          |
| vcp-255/1/0 |           |              |             |       |              |      |       |       | 0        |
| vcp-255/1/1 |           |              |             |       |              |      |       |       | 0        |

## show virtual-chassis vc-port

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show virtual-chassis vc-port<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Display the status of the Virtual Chassis ports (VCPs), including both the dedicated VCPs and the uplink ports configured as VCPs.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>none</b>—Display the operational status of all VCPs of the member switch where the command is issued.</p> <p><b>all-members</b>—(Optional) Display the operational status of all VCPs on all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the operational status of the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the operational status of all VCPs for the specified member of the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis vc-port statistics on page 5227</a></li> <li>• <a href="#">Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157</a></li> <li>• <a href="#">Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis</a></li> </ul>                                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis vc-port (EX4200 Virtual Chassis) on page 5368</a><br><a href="#">show virtual-chassis vc-port (EX8200 Virtual Chassis) on page 5368</a><br><a href="#">show virtual-chassis vc-port all-members on page 5369</a>                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>            | Table 581 on page 5209 lists the output fields for the <b>show virtual-chassis vc-port</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                |

**Table 596: show virtual-chassis vc-port Output Fields**

| Field Name       | Field Description                            |
|------------------|----------------------------------------------|
| <b>fpcnumber</b> | The FPC number is the same as the member ID. |

Table 596: show virtual-chassis vc-port Output Fields (*continued*)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface or PIC/Port | <p>VCP name.</p> <ul style="list-style-type: none"> <li>The dedicated VCPs in an EX4200 or EX4500 Virtual Chassis are <b>vcp-0</b> and <b>vcp-1</b>. The dedicated VCPs in an EX4550 Virtual Chassis are <b>VCP-1/0</b>, <b>VCP-1/1</b>, <b>VCP-2/0</b>, and <b>VCP-2/1</b>.</li> <li>Optical ports set as VCPs are named <b>1/0</b> and <b>1/1</b>, representing the PIC number and the port number.</li> <li>The native VCP (port 0) on an XRE200 External Routing Engine in an EX8200 Virtual Chassis is named <b>vcp-0</b>.</li> <li>The VCPs on each Virtual Chassis Control Interface (VCCI) module in an XRE200 External Routing Engine are named using the <b>vcp-slot-number/port-number</b> convention; for instance, <b>vcp-1/0</b>.</li> <li>The VCPs on EX8200 member switches are named using the <b>vcp-slot-number/pic-number/interface-number</b> convention; for instance, <b>vcp-3/0/2</b>.</li> <li>A <b>255</b> as the first number in your port number indicates that your VCP is part of a Link Aggregation group (LAG) bundle. For instance, a display of <b>vcp-255/1/0</b> indicates that the dedicated VCP named <b>vcp-1/0</b> is part of a LAG bundle. A display of <b>vcp-255/1/0</b> indicates that an uplink port that was previously named <b>xe-0/1/0</b> is now part of a VCP LAG bundle.</li> </ul> |
| Type                  | <p>Type of VCP:</p> <ul style="list-style-type: none"> <li><b>Dedicated</b>—The rear panel VCP on an EX4200, EX4500, or EX4550 switch, or any VCP link connected to an XRE200 External Routing Engine in an EX8200 Virtual Chassis.</li> <li><b>Configured</b>—Optical port configured as a VCP.</li> <li><b>Auto-Configured</b>—Optical port autoconfigured as a VCP.</li> </ul> <p>See “<a href="#">Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)</a>” on <a href="#">page 5109</a> or <i>Setting a 10-Gigabit Ethernet Port as a Virtual Chassis Port in an EX8200 Virtual Chassis (CLI Procedure)</i> for information about configuring VCPs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Trunk ID              | <p>A positive-number ID assigned to a link aggregation group (LAG) formed by the Virtual Chassis. The trunk ID value is –1 if no trunk is formed. A LAG between uplink VCPs requires that the link speed be the same on connected interfaces and that at least two VCPs on one member be connected to at least two VCPs on the other member in an EX4200 or EX4500 Virtual Chassis.</p> <p>Dedicated VCP LAGs are assigned trunk IDs 1 and 2. Trunk IDs for LAGs formed with uplink VCPs therefore have values of 3 or greater.</p> <p>The trunk ID value changes if the link-adjacency state between LAG members changes; trunk membership is then allocated or deallocated.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Status                | <p>Interface status:</p> <ul style="list-style-type: none"> <li><b>absent</b>—Interface is not a VCP link.</li> <li><b>down</b>—VCP link is down.</li> <li><b>up</b>—VCP link is up.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Speed (mbps)          | Speed of the interface in megabits per second.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Neighbor ID/Interface | The Virtual Chassis member ID and interface of a VCP on a member that is connected to the interface or PIC/Port field in the same row as this interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## Sample Output

### show virtual-chassis vc-port (EX4200 Virtual Chassis)

```
user@switch> show virtual-chassis vc-port
```

```
fpc0:
```

| Interface<br>or<br>PIC / Port | Type            | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface   |
|-------------------------------|-----------------|-------------|--------|-----------------|----------------|-------------|
| vcp-0                         | Dedicated       | 1           | Up     | 32000           | 1              | vcp-1       |
| vcp-1                         | Dedicated       | 2           | Up     | 32000           | 0              | vcp-0       |
| 1/0                           | Auto-Configured | 3           | Up     | 1000            | 2              | vcp-255/1/0 |
| 1/0                           | Auto-Configured | 3           | Up     | 1000            | 2              | vcp-255/1/1 |

### show virtual-chassis vc-port (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis vc-port
```

```
member0:
```

| Interface<br>or<br>Slot/PIC/Port | Type       | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface |
|----------------------------------|------------|-------------|--------|-----------------|----------------|-----------|
| vcp-0/0                          | Dedicated  | -1          | Up     | 1000            | 8              | vcp-1/1   |
| vcp-0/1                          | Dedicated  | -1          | Up     | 1000            | 8              | vcp-2/0   |
| 4/0/4                            | Configured | -1          | Up     | 10000           | 1              | vcp-3/0/4 |
| 4/0/7                            | Configured | -1          | Down   | 10000           |                |           |
| 4/0/3                            | Configured |             | Absent |                 |                |           |
| 4/0/2                            | Configured |             | Absent |                 |                |           |
| 4/0/5                            | Configured |             | Absent |                 |                |           |
| 4/0/6                            | Configured |             | Absent |                 |                |           |
| 4/0/1                            | Configured |             | Absent |                 |                |           |
| 4/0/0                            | Configured |             | Absent |                 |                |           |

```
member1:
```

| Interface<br>or<br>Slot/PIC/Port | Type       | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface |
|----------------------------------|------------|-------------|--------|-----------------|----------------|-----------|
| vcp-0/0                          | Dedicated  | -1          | Up     | 1000            | 8              | vcp-1/2   |
| 3/0/0                            | Configured | -1          | Down   | 10000           |                |           |
| 3/0/1                            | Configured | -1          | Down   | 10000           |                |           |
| 3/0/4                            | Configured | -1          | Up     | 10000           | 0              | vcp-4/0/4 |
| 3/0/5                            | Configured |             | Absent |                 |                |           |
| 4/0/5                            | Configured |             | Absent |                 |                |           |
| 4/0/4                            | Configured |             | Absent |                 |                |           |

```
member8:
```

| Interface<br>or<br>Slot/PIC/Port | Type      | Trunk<br>ID | Status | Speed<br>(mbps) | Neighbor<br>ID | Interface |
|----------------------------------|-----------|-------------|--------|-----------------|----------------|-----------|
| vcp-0/0                          | Dedicated | -1          | Down   | 1000            |                |           |
| vcp-1/0                          | Dedicated | -1          | Up     | 1000            | 9              | vcp-1/0   |
| vcp-1/1                          | Dedicated | -1          | Up     | 1000            | 0              | vcp-0/0   |
| vcp-1/2                          | Dedicated | -1          | Up     | 1000            | 1              | vcp-0/0   |
| vcp-1/3                          | Dedicated | -1          | Up     | 1000            | 9              | vcp-1/3   |
| vcp-2/0                          | Dedicated | -1          | Up     | 1000            | 0              | vcp-0/1   |
| vcp-2/1                          | Dedicated | -1          | Up     | 1000            | 9              | vcp-1/2   |
| vcp-2/2                          | Dedicated | -1          | Down   | 1000            |                |           |



```
vcp-2/3      Dedicated      -1   Down      1000
```

```
member9:
```

```
-----
Interface    Type           Trunk  Status    Speed    Neighbor
or           or              ID     Status    (mbps)   ID  Interface
Slot/PIC/Port
vcp-0/0      Dedicated      -1     Disabled  1000
vcp-1/0      Dedicated      -1     Up        1000      8   vcp-1/0
vcp-1/1      Dedicated      -1     Down      1000
vcp-1/2      Dedicated      -1     Up        1000      8   vcp-2/1
vcp-1/3      Dedicated      -1     Up        1000      8   vcp-1/3
```

### show virtual-chassis vc-port all-members

```
user@switch> show virtual-chassis vc-port all-members
```

```
fpc0:
```

```
-----
Interface    Type           Trunk  Status    Speed    Neighbor
or           or              ID     Status    (mbps)   ID  Interface
PIC / Port
vcp-0        Dedicated      1      Up        32000    1   vcp-1
vcp-1        Dedicated      2      Up        32000    0   vcp-0
1/0          Auto-Configured 3      Up        1000     2   vcp-255/1/0
1/1          Auto-Configured 3      Up        1000     2   vcp-255/1/1
```

```
fpc1:
```

```
-----
Interface    Type           Trunk  Status    Speed    Neighbor
or           or              ID     Status    (mbps)   ID  Interface
PIC / Port
vcp-0        Dedicated      1      Up        32000    0   vcp-1
vcp-1        Dedicated      2      Up        32000    0   vcp-0
1/0          Auto-Configured -1     Up        1000     3   vcp-255/1/0
```

```
fpc2:
```

```
-----
Interface    Type           Trunk  Status    Speed    Neighbor
or           or              ID     Status    (mbps)   ID  Interface
PIC / Port
vcp-0        Dedicated      1      Up        32000    3   vcp-1
vcp-1        Dedicated      2      Up        32000    3   vcp-0
1/0          Auto-Configured 3      Up        1000     0   vcp-255/1/0
1/1          Auto-Configured 3      Up        1000     0   vcp-255/1/1
```

```
fpc3:
```

```
-----
Interface    Type           Trunk  Status    Speed    Neighbor
or           or              ID     Status    (mbps)   ID  Interface
PIC / Port
vcp-0        Dedicated      1      Up        32000    2   vcp-0
vcp-1        Dedicated      2      Up        32000    2   vcp-1
1/0          Auto-Configured -1     Up        1000     1   vcp-255/1/0
```

## show virtual-chassis vc-port diagnostics optics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show virtual-chassis vc-port diagnostics optics &lt;all-members&gt; &lt;interface-name&gt; &lt;local&gt; &lt;member member-id&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.2 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Display diagnostics data and alarms for Ethernet optical transceivers installed in ports configured as Virtual Chassis Ports (VCPs) in an EX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that a transceiver is not operating properly. DOM information can be used to diagnose why a transceiver is not working.</p> <p>On some EX Series switches, the <b>request virtual-chassis vc-port diagnostics optics</b> command must be entered to run a diagnostic scan before you can gather the <b>show virtual-chassis vc-port diagnostics optics</b> output.</p> |
| <b>Options</b>                  | <p><b>none</b>—Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display diagnostics information for the transceiver installed in a specified VCP.</p> <p><b>local</b>—(Optional) Display diagnostics information for transceivers installed in VCPs on the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display diagnostics information for transceivers installed in VCPs on a specified member of a Virtual Chassis or VCF.</p>                                                       |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show virtual-chassis vc-port on page 5209</a></li><li>• <i>Installing a Transceiver in an EX Series Switch</i></li><li>• <i>Removing a Transceiver from an EX Series Switch</i></li><li>• <a href="#">Junos OS Ethernet Interfaces Configuration Guide</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">show virtual-chassis vc-port diagnostics optics on page 5373</a><br><a href="#">show virtual-chassis vc-port diagnostics optics (interface-name) on page 5378</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

[show virtual-chassis vc-port diagnostics optics local](#) on page 5380

[show virtual-chassis vc-port diagnostics optics \(member member-id\)](#) on page 5382

**Output Fields** [Table 582 on page 5214](#) lists the output fields for the **show virtual-chassis vc-port diagnostics optics** command. Output fields are listed in the approximate order in which they appear.

**Table 597: show virtual-chassis vc-port diagnostics optics Output Fields**

| Field Name                            | Field Description                                                                                                                                                       |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FPC                                   | Displays the FPC slot number.                                                                                                                                           |
| Virtual chassis port                  | Displays the name of the VCP.                                                                                                                                           |
| Laser bias current                    | Displays the magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents. |
| Laser output power                    | Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).                                                                              |
| Module temperature                    | Displays the temperature, in Celsius and Fahrenheit.                                                                                                                    |
| Module voltage                        | Displays the voltage, in Volts.                                                                                                                                         |
| Receiver signal average optical power | Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).                                                           |
| Laser bias current high alarm         | Displays whether the laser bias power setting high alarm is <i>On</i> or <i>Off</i> .                                                                                   |
| Laser bias current low alarm          | Displays whether the laser bias power setting low alarm is <i>On</i> or <i>Off</i> .                                                                                    |
| Laser bias current high warning       | Displays whether the laser bias power setting high warning is <i>On</i> or <i>Off</i> .                                                                                 |
| Laser bias current low warning        | Displays whether the laser bias power setting low warning is <i>On</i> or <i>Off</i> .                                                                                  |
| Laser output power high alarm         | Displays whether the laser output power high alarm is <i>On</i> or <i>Off</i> .                                                                                         |
| Laser output power low alarm          | Displays whether the laser output power low alarm is <i>On</i> or <i>Off</i> .                                                                                          |
| Laser output power high warning       | Displays whether the laser output power high warning is <i>On</i> or <i>Off</i> .                                                                                       |
| Laser output power low warning        | Displays whether the laser output power low warning is <i>On</i> or <i>Off</i> .                                                                                        |
| Module temperature high alarm         | Displays whether the module temperature high alarm is <i>On</i> or <i>Off</i> .                                                                                         |
| Module temperature low alarm          | Displays whether the module temperature low alarm is <i>On</i> or <i>Off</i> .                                                                                          |
| Module temperature high warning       | Displays whether the module temperature high warning is <i>On</i> or <i>Off</i> .                                                                                       |
| Module temperature low warning        | Displays whether the module temperature low warning is <i>On</i> or <i>Off</i> .                                                                                        |

Table 597: show virtual-chassis vc-port diagnostics optics Output Fields (*continued*)

| Field Name                                | Field Description                                                                  |
|-------------------------------------------|------------------------------------------------------------------------------------|
| Module voltage high alarm                 | Displays whether the module voltage high alarm is <i>On</i> or <i>Off</i> .        |
| Module voltage low alarm                  | Displays whether the module voltage low alarm is <i>On</i> or <i>Off</i> .         |
| Module voltage high warning               | Displays whether the module voltage high warning is <i>On</i> or <i>Off</i> .      |
| Module voltage low warning                | Displays whether the module voltage low warning is <i>On</i> or <i>Off</i> .       |
| Laser rx power high alarm                 | Displays whether the receive laser power high alarm is <i>On</i> or <i>Off</i> .   |
| Laser rx power low alarm                  | Displays whether the receive laser power low alarm is <i>On</i> or <i>Off</i> .    |
| Laser rx power high warning               | Displays whether the receive laser power high warning is <i>On</i> or <i>Off</i> . |
| Laser rx power low warning                | Displays whether the receive laser power low warning is <i>On</i> or <i>Off</i> .  |
| Laser bias current high alarm threshold   | Displays the vendor-specified threshold for the laser bias current high alarm.     |
| Laser bias current low alarm threshold    | Displays the vendor-specified threshold for the laser bias current low alarm.      |
| Laser bias current high warning threshold | Displays the vendor-specified threshold for the laser bias current high warning.   |
| Laser bias current low warning threshold  | Displays the vendor-specified threshold for the laser bias current low warning.    |
| Laser output power high alarm threshold   | Displays the vendor-specified threshold for the laser output power high alarm.     |
| Laser output power low alarm threshold    | Displays the vendor-specified threshold for the laser output power low alarm.      |
| Laser output power high warning threshold | Displays the vendor-specified threshold for the laser output power high warning.   |
| Laser output power low warning threshold  | Displays the vendor-specified threshold for the laser output power low warning.    |
| Module temperature high alarm threshold   | Displays the vendor-specified threshold for the module temperature high alarm.     |
| Module temperature low alarm threshold    | Displays the vendor-specified threshold for the module temperature low alarm.      |
| Module temperature high warning threshold | Displays the vendor-specified threshold for the module temperature high warning.   |
| Module temperature low warning threshold  | Displays the vendor-specified threshold for the module temperature low warning.    |

Table 597: show virtual-chassis vc-port diagnostics optics Output Fields (*continued*)

| Field Name                            | Field Description                                                            |
|---------------------------------------|------------------------------------------------------------------------------|
| Module voltage high alarm threshold   | Displays the vendor-specified threshold for the module voltage high alarm.   |
| Module voltage low alarm threshold    | Displays the vendor-specified threshold for the module voltage low alarm.    |
| Module voltage high warning threshold | Displays the vendor-specified threshold for the module voltage high warning. |
| Module voltage low warning threshold  | Displays the vendor-specified threshold for the module voltage low warning.  |
| Laser rx power high alarm threshold   | Displays the vendor-specified threshold for the laser rx power high alarm.   |
| Laser rx power low alarm threshold    | Displays the vendor-specified threshold for the laser rx power low alarm.    |
| Laser rx power high warning threshold | Displays the vendor-specified threshold for the laser rx power high warning. |
| Laser rx power low warning threshold  | Displays the vendor-specified threshold for the laser rx power low warning.  |

## Sample Output

### show virtual-chassis vc-port diagnostics optics

```

user@switch> show virtual-chassis vc-port diagnostics optics
fpc0:
-----
Virtual chassis port: vcp-0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-1
  Optical diagnostics                : N/A

fpc1:
-----
Virtual chassis port: vcp-0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-1
  Optical diagnostics                : N/A

fpc2:
-----
Virtual chassis port: vcp-2/0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current                 : 4.130 mA
  Laser output power                 : 0.2450 mW / -6.11 dBm
  Module temperature                 : 32 degrees C / 90 degrees F
  Module voltage                     : 3.3530 V
  Receiver signal average optical power : 0.0971 mW / -10.13 dBm
  Laser bias current high alarm      : Off
  Laser bias current low alarm       : Off

```

```
Laser bias current high warning      : Off
Laser bias current low warning       : Off
Laser output power high alarm        : Off
Laser output power low alarm         : Off
Laser output power high warning      : Off
Laser output power low warning       : Off
Module temperature high alarm        : Off
Module temperature low alarm         : Off
Module temperature high warning      : Off
Module temperature low warning       : Off
Module voltage high alarm            : Off
Module voltage low alarm             : Off
Module voltage high warning          : Off
Module voltage low warning           : Off
Laser rx power high alarm            : Off
Laser rx power low alarm             : Off
Laser rx power high warning          : Off
Laser rx power low warning           : Off
Laser bias current high alarm threshold : 14.998 mA
Laser bias current low alarm threshold : 0.998 mA
Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold : 1.198 mA
Laser output power high alarm threshold : 0.7940 mW / -1.00 dBm
Laser output power low alarm threshold : 0.0790 mW / -11.02 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0990 mW / -10.04 dBm
Module temperature high alarm threshold : 85 degrees C / 185 degrees F
Module temperature low alarm threshold : -10 degrees C / 14 degrees F
Module temperature high warning threshold : 80 degrees C / 176 degrees F
Module temperature low warning threshold : -5 degrees C / 23 degrees F
Module voltage high alarm threshold : 3.600 V
Module voltage low alarm threshold : 3.000 V
Module voltage high warning threshold : 3.499 V
Module voltage low warning threshold : 3.099 V
Laser rx power high alarm threshold : 1.5848 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                  : 5.428 mA
Laser output power                  : 0.4760 mW / -3.22 dBm
Module temperature                  : 28 degrees C / 83 degrees F
Module voltage                      : 3.3440 V
Receiver signal average optical power : 0.4002 mW / -3.98 dBm
Laser bias current high alarm       : Off
Laser bias current low alarm        : Off
Laser bias current high warning     : Off
Laser bias current low warning      : Off
Laser output power high alarm       : Off
Laser output power low alarm        : Off
Laser output power high warning     : Off
Laser output power low warning      : Off
Module temperature high alarm       : Off
Module temperature low alarm        : Off
Module temperature high warning     : Off
Module temperature low warning      : Off
Module voltage high alarm           : Off
Module voltage low alarm            : Off
Module voltage high warning         : Off
Module voltage low warning          : Off
Laser rx power high alarm           : Off
```

```

Laser rx power low alarm           : Off
Laser rx power high warning        : Off
Laser rx power low warning         : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

fpc3:

-----  
Virtual chassis port: vcp-255/0/2

```

Laser bias current           : 7.876 mA
Laser output power           : 0.5330 mW / -2.73 dBm
Module temperature            : 26 degrees C / 78 degrees F
Module voltage                : 3.3060 V
Receiver signal average optical power : 0.4885 mW / -3.11 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm  : Off
Laser output power low alarm   : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm  : Off
Module temperature low alarm   : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm      : Off
Module voltage low alarm       : Off
Module voltage high warning    : Off
Module voltage low warning     : Off
Laser rx power high alarm      : Off
Laser rx power low alarm       : Off
Laser rx power high warning    : Off
Laser rx power low warning     : Off
Laser bias current high alarm threshold : 14.500 mA
Laser bias current low alarm threshold : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F

```

```
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current : 5.052 mA
Laser output power : 0.5030 mW / -2.98 dBm
Module temperature : 24 degrees C / 75 degrees F
Module voltage : 3.2890 V
Receiver signal average optical power : 0.5028 mW / -2.99 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/4
Laser bias current : 7.978 mA
Laser output power : 0.5460 mW / -2.63 dBm
Module temperature : 24 degrees C / 76 degrees F
```



```

Module voltage : 3.3060 V
Receiver signal average optical power : 0.6305 mW / -2.00 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
Laser bias current high alarm threshold : 14.500 mA
Laser bias current low alarm threshold : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

fpc4:

```

-----
Virtual chassis port: vcp-0
  Optical diagnostics : N/A
Virtual chassis port: vcp-1
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/4
  Laser bias current : 7.860 mA
  Laser output power : 0.5370 mW / -2.70 dBm
  Module temperature : 24 degrees C / 75 degrees F
  Module voltage : 3.2920 V
  Receiver signal average optical power : 0.6271 mW / -2.03 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm : Off
  Laser output power high warning : Off

```

```

Laser output power low warning           : Off
Module temperature high alarm            : Off
Module temperature low alarm             : Off
Module temperature high warning          : Off
Module temperature low warning           : Off
Module voltage high alarm                : Off
Module voltage low alarm                 : Off
Module voltage high warning              : Off
Module voltage low warning               : Off
Laser rx power high alarm                : Off
Laser rx power low alarm                 : Off
Laser rx power high warning              : Off
Laser rx power low warning               : Off
Laser bias current high alarm threshold  : 14.500 mA
Laser bias current low alarm threshold   : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold  : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold   : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold  : 75 degrees C / 167 degrees F
Module temperature low alarm threshold   : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold       : 3.630 V
Module voltage low alarm threshold        : 2.970 V
Module voltage high warning threshold     : 3.465 V
Module voltage low warning threshold       : 3.135 V
Laser rx power high alarm threshold       : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold        : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold     : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold       : 0.1023 mW / -9.90 dBm

```

#### show virtual-chassis vc-port diagnostics optics (interface-name)

```

user@external-routing-engine> show virtual-chassis vc-port diagnostics optics vcp-255/0/3
fpc0:

```

```

fpc1:

```

```

fpc2:

```

```

Virtual chassis port: vcp-255/0/3
Laser bias current           : 5.448 mA
Laser output power           : 0.4770 mW / -3.21 dBm
Module temperature           : 28 degrees C / 82 degrees F
Module voltage                : 3.3450 V
Receiver signal average optical power : 0.3973 mW / -4.01 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm  : Off
Laser output power low alarm   : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm  : Off
Module temperature low alarm   : Off

```

```

Module temperature high warning      : Off
Module temperature low warning       : Off
Module voltage high alarm            : Off
Module voltage low alarm             : Off
Module voltage high warning          : Off
Module voltage low warning           : Off
Laser rx power high alarm            : Off
Laser rx power low alarm             : Off
Laser rx power high warning          : Off
Laser rx power low warning           : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold   : 3.630 V
Module voltage low alarm threshold    : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold  : 3.135 V
Laser rx power high alarm threshold   : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold    : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold  : 0.1023 mW / -9.90 dBm

```

fpc3:

-----  
Virtual chassis port: vcp-255/0/3

```

Laser bias current      : 5.040 mA
Laser output power      : 0.5020 mW / -2.99 dBm
Module temperature      : 24 degrees C / 74 degrees F
Module voltage          : 3.2870 V
Receiver signal average optical power : 0.5073 mW / -2.95 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm  : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm  : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm     : Off
Module voltage low alarm      : Off
Module voltage high warning   : Off
Module voltage low warning    : Off
Laser rx power high alarm     : Off
Laser rx power low alarm      : Off
Laser rx power high warning   : Off
Laser rx power low warning    : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold  : 2.000 mA

```

```

Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold  : 2.500 mA
Laser output power high alarm threshold   : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold    : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold  : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold    : 75 degrees C / 167 degrees F
Module temperature low alarm threshold     : -5 degrees C / 23 degrees F
Module temperature high warning threshold  : 70 degrees C / 158 degrees F
Module temperature low warning threshold   : 0 degrees C / 32 degrees F
Module voltage high alarm threshold        : 3.630 V
Module voltage low alarm threshold         : 2.970 V
Module voltage high warning threshold      : 3.465 V
Module voltage low warning threshold       : 3.135 V
Laser rx power high alarm threshold        : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold         : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold      : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold       : 0.1023 mW / -9.90 dBm

```

fpc4:

-----

#### show virtual-chassis vc-port diagnostics optics local

```

user@switch> show virtual-chassis vc-port diagnostics optics local
Virtual chassis port: vcp-2/0
  Optical diagnostics : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current : 4.130 mA
  Laser output power : 0.2450 mW / -6.11 dBm
  Module temperature : 32 degrees C / 90 degrees F
  Module voltage     : 3.3530 V
  Receiver signal average optical power : 0.0961 mW / -10.17 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm  : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm : Off
  Module temperature low alarm  : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Module voltage high alarm : Off
  Module voltage low alarm  : Off
  Module voltage high warning : Off
  Module voltage low warning : Off
  Laser rx power high alarm : Off
  Laser rx power low alarm  : Off
  Laser rx power high warning : Off
  Laser rx power low warning : Off
  Laser bias current high alarm threshold : 14.998 mA
  Laser bias current low alarm threshold  : 0.998 mA

```

```

Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold  : 1.198 mA
Laser output power high alarm threshold   : 0.7940 mW / -1.00 dBm
Laser output power low alarm threshold    : 0.0790 mW / -11.02 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold  : 0.0990 mW / -10.04 dBm
Module temperature high alarm threshold   : 85 degrees C / 185 degrees F
Module temperature low alarm threshold    : -10 degrees C / 14 degrees F
Module temperature high warning threshold : 80 degrees C / 176 degrees F
Module temperature low warning threshold  : -5 degrees C / 23 degrees F
Module voltage high alarm threshold       : 3.600 V
Module voltage low alarm threshold        : 3.000 V
Module voltage high warning threshold     : 3.499 V
Module voltage low warning threshold      : 3.099 V
Laser rx power high alarm threshold       : 1.5848 mW / 2.00 dBm
Laser rx power low alarm threshold        : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold     : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold      : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                       : 5.426 mA
Laser output power                       : 0.4760 mW / -3.22 dBm
Module temperature                       : 28 degrees C / 83 degrees F
Module voltage                           : 3.3450 V
Receiver signal average optical power    : 0.3955 mW / -4.03 dBm
Laser bias current high alarm            : Off
Laser bias current low alarm             : Off
Laser bias current high warning          : Off
Laser bias current low warning           : Off
Laser output power high alarm            : Off
Laser output power low alarm             : Off
Laser output power high warning          : Off
Laser output power low warning           : Off
Module temperature high alarm            : Off
Module temperature low alarm             : Off
Module temperature high warning          : Off
Module temperature low warning           : Off
Module voltage high alarm                : Off
Module voltage low alarm                 : Off
Module voltage high warning              : Off
Module voltage low warning               : Off
Laser rx power high alarm                : Off
Laser rx power low alarm                 : Off
Laser rx power high warning              : Off
Laser rx power low warning               : Off
Laser bias current high alarm threshold  : 10.500 mA
Laser bias current low alarm threshold   : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold  : 2.500 mA
Laser output power high alarm threshold  : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold    : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold  : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold   : 75 degrees C / 167 degrees F
Module temperature low alarm threshold    : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold  : 0 degrees C / 32 degrees F
Module voltage high alarm threshold       : 3.630 V
Module voltage low alarm threshold        : 2.970 V
Module voltage high warning threshold     : 3.465 V
Module voltage low warning threshold      : 3.135 V
Laser rx power high alarm threshold       : 1.5849 mW / 2.00 dBm

```

```

Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold    : 0.1023 mW / -9.90 dBm

```

### show virtual-chassis vc-port diagnostics optics (member member-id)

```

user@switch> show virtual-chassis vc-port diagnostics optics member 2
fpc2:

```

```

-----
Virtual chassis port: vcp-2/0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current                  : 4.130 mA
  Laser output power                  : 0.2450 mW / -6.11 dBm
  Module temperature                  : 31 degrees C / 88 degrees F
  Module voltage                      : 3.3530 V
  Receiver signal average optical power : 0.0961 mW / -10.17 dBm
  Laser bias current high alarm       : Off
  Laser bias current low alarm        : Off
  Laser bias current high warning     : Off
  Laser bias current low warning      : Off
  Laser output power high alarm       : Off
  Laser output power low alarm        : Off
  Laser output power high warning     : Off
  Laser output power low warning      : Off
  Module temperature high alarm       : Off
  Module temperature low alarm        : Off
  Module temperature high warning     : Off
  Module temperature low warning      : Off
  Module voltage high alarm           : Off
  Module voltage low alarm            : Off
  Module voltage high warning         : Off
  Module voltage low warning          : Off
  Laser rx power high alarm           : Off
  Laser rx power low alarm            : Off
  Laser rx power high warning         : Off
  Laser rx power low warning          : Off
  Laser bias current high alarm threshold : 14.998 mA
  Laser bias current low alarm threshold : 0.998 mA
  Laser bias current high warning threshold : 14.000 mA
  Laser bias current low warning threshold : 1.198 mA
  Laser output power high alarm threshold : 0.7940 mW / -1.00 dBm
  Laser output power low alarm threshold : 0.0790 mW / -11.02 dBm
  Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
  Laser output power low warning threshold : 0.0990 mW / -10.04 dBm
  Module temperature high alarm threshold : 85 degrees C / 185 degrees F
  Module temperature low alarm threshold : -10 degrees C / 14 degrees F
  Module temperature high warning threshold : 80 degrees C / 176 degrees F
  Module temperature low warning threshold : -5 degrees C / 23 degrees F
  Module voltage high alarm threshold : 3.600 V
  Module voltage low alarm threshold : 3.000 V
  Module voltage high warning threshold : 3.499 V
  Module voltage low warning threshold : 3.099 V
  Laser rx power high alarm threshold : 1.5848 mW / 2.00 dBm

```

```

Laser rx power low alarm threshold      : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold   : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold    : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                      : 5.418 mA
Laser output power                      : 0.4770 mW / -3.21 dBm
Module temperature                      : 28 degrees C / 83 degrees F
Module voltage                          : 3.3450 V
Receiver signal average optical power   : 0.3964 mW / -4.02 dBm
Laser bias current high alarm           : Off
Laser bias current low alarm            : Off
Laser bias current high warning         : Off
Laser bias current low warning          : Off
Laser output power high alarm           : Off
Laser output power low alarm            : Off
Laser output power high warning         : Off
Laser output power low warning          : Off
Module temperature high alarm           : Off
Module temperature low alarm            : Off
Module temperature high warning         : Off
Module temperature low warning          : Off
Module voltage high alarm               : Off
Module voltage low alarm                : Off
Module voltage high warning             : Off
Module voltage low warning              : Off
Laser rx power high alarm               : Off
Laser rx power low alarm                : Off
Laser rx power high warning             : Off
Laser rx power low warning              : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold  : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold  : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold      : 3.630 V
Module voltage low alarm threshold       : 2.970 V
Module voltage high warning threshold    : 3.465 V
Module voltage low warning threshold     : 3.135 V
Laser rx power high alarm threshold      : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold       : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold    : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold     : 0.1023 mW / -9.90 dBm

```

## show virtual-chassis vc-port statistics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show virtual-chassis vc-port statistics &lt;all-members&gt; &lt;brief   detail   extensive &gt; &lt;interface-name&gt; &lt;local&gt; &lt;member member-id&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The options <b>all-members</b>, <b>brief</b>, <b>detail</b>, <b>extensive</b>, and <b>local</b> were added in Junos OS Release 9.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Display the traffic statistics collected on Virtual Chassis ports (VCPs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>none</b>—Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> and <b>extensive</b> options provide identical displays.</p> <p><b>all-members</b>—(Optional) Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display traffic statistics for the specified VCP.</p> <p><b>local</b>—(Optional) Display traffic statistics for VCPs on the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display traffic statistics for VCPs on the specified member of a Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">clear virtual-chassis vc-port statistics on page 5161</a></li><li>• <a href="#">show virtual-chassis vc-port on page 5209</a></li><li>• <a href="#">Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis on page 5157</a></li><li>• <a href="#">Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <p><a href="#">show virtual-chassis vc-port statistics on page 5387</a></p> <p><a href="#">show virtual-chassis vc-port statistics (EX8200 Virtual Chassis) on page 5388</a></p> <p><a href="#">show virtual-chassis vc-port statistics brief on page 5388</a></p> <p><a href="#">show virtual-chassis vc-port statistics extensive on page 5388</a></p> <p><a href="#">show virtual-chassis vc-port statistics member 0 on page 5390</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                               |



**Output Fields** Table 582 on page 5214 lists the output fields for the **show virtual-chassis vc-port statistics** command. Output fields are listed in the approximate order in which they appear.

**Table 598: show virtual-chassis vc-port statistics Output Fields**

| Field Name                   | Field Description                                                                                                                                                                    | Level of Output            |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|
| <b>fpcnumber</b>             | (All Virtual Chassis except EX8200 Virtual Chassis. VCF) ID of the Virtual Chassis member. The FPC number is the same as the member ID.                                              | All levels                 |
| <b>member number</b>         | (EX8200 Virtual Chassis only) Member ID of the Virtual Chassis member.                                                                                                               | All levels                 |
| <b>Interface</b>             | VCP name.                                                                                                                                                                            | <b>brief</b>               |
| <b>Input Octets/Packets</b>  | Number of octets and packets received on the VCP.                                                                                                                                    | <b>brief, member, none</b> |
| <b>Output Octets/Packets</b> | Number of octets and packets transmitted on the VCP.                                                                                                                                 | <b>brief, member, none</b> |
| <b>master: number</b>        | Member ID of the master Routing Engine.                                                                                                                                              | All levels                 |
| <b>Port</b>                  | VCP for which RX (Receive) statistics, TX (Transmit) statistics, or both are reported by the VCP subsystem during a sampling interval—since the statistics counter was last cleared. | <b>detail, extensive</b>   |
| <b>Total octets</b>          | Total number of octets received and transmitted on the VCP.                                                                                                                          | <b>detail, extensive</b>   |
| <b>Total packets</b>         | Total number of packets received and transmitted on the VCP.                                                                                                                         | <b>detail, extensive</b>   |
| <b>Unicast packets</b>       | Number of unicast packets received and transmitted on the VCP.                                                                                                                       | <b>detail, extensive</b>   |
| <b>Broadcast packets</b>     | Number of broadcast packets received and transmitted on the VCP.                                                                                                                     | <b>detail, extensive</b>   |
| <b>Multicast packets</b>     | Number of multicast packets received and transmitted on the VCP.                                                                                                                     | <b>detail, extensive</b>   |
| <b>MAC control frames</b>    | Number of media access control (MAC) control frames received and transmitted on the VCP.                                                                                             | <b>detail, extensive</b>   |

Table 598: show virtual-chassis vc-port statistics Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output          |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>CRC alignment errors</b> | <p>Number of packets received on the VCP that had a length—excluding framing bits, but including frame check sequence (FCS) octets—of between 64 and 1518 octets, inclusive, and had one of the following errors:</p> <ul style="list-style-type: none"> <li>Invalid FCS with an integral number of octets (FCS error)</li> <li>Invalid FCS with a nonintegral number of octets (alignment error)</li> </ul>                                                                                                   | <b>detail, extensive</b> |
| <b>Oversize packets</b>     | Number of packets received on the VCP that were longer than 1518 octets (excluding framing bits, but including FCS octets) but were otherwise well formed.                                                                                                                                                                                                                                                                                                                                                     | <b>detail, extensive</b> |
| <b>Undersize packets</b>    | Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits but including FCS octets) and were otherwise well formed..                                                                                                                                                                                                                                                                                                                                                      | <b>detail, extensive</b> |
| <b>Jabber packets</b>       | <p>Number of packets received on the VCP that were longer than 1518 octets—excluding framing bits, but including FCS octets—and that had either an FCS error or an alignment error.</p> <p><b>NOTE:</b> This definition of <i>jabber</i> is different from the definition in IEEE-802.3 section 8.2.1.5 (10Base5) and section 10.3.1.4 (10Base2). These documents define <i>jabber</i> as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.</p> | <b>detail, extensive</b> |
| <b>Fragments received</b>   | <p>Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error.</p> <p>Fragment frames normally increment because both runs (which are normal occurrences caused by collisions) and noise hits are counted.</p>                                                                                                                                                                               | <b>detail, extensive</b> |
| <b>Ifout errors</b>         | Number of outbound packets received on the VCP that could not be transmitted because of errors.                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail, extensive</b> |
| <b>Packet drop events</b>   | Number of outbound packets received on the VCP that were dropped, rather than being encapsulated and sent out of the switch as fragments. The packet drop counter is incremented if a temporary shortage of packet memory causes packet fragmentation to fail.                                                                                                                                                                                                                                                 | <b>detail, extensive</b> |
| <b>64 octets frames</b>     | Number of packets received on the VCP (including invalid packets) that were 64 octets in length (excluding framing bits, but including FCS octets).                                                                                                                                                                                                                                                                                                                                                            | <b>detail, extensive</b> |

Table 598: show virtual-chassis vc-port statistics Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                 | Level of Output   |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| 65–127 octets frames    | Number of packets received on the VCP (including invalid packets) that were between 65 and 127 octets in length, inclusive (excluding framing bits, but including FCS octets).    | detail, extensive |
| 128–255 octets frames   | Number of packets received on the VCP (including invalid packets) that were between 128 and 255 octets in length, inclusive (excluding framing bits, but including FCS octets).   | detail, extensive |
| 256–511 octets frames   | Number of packets received on the VCP (including invalid packets) that were between 256 and 511 octets in length, inclusive (excluding framing bits, but including FCS octets).   | detail, extensive |
| 512–1023 octets frames  | Number of packets received on the VCP (including invalid packets) that were between 512 and 1023 octets in length, inclusive (excluding framing bits, but including FCS octets).  | detail, extensive |
| 1024–1518 octets frames | Number of packets received on the VCP (including invalid packets) that were between 1024 and 1518 octets in length, inclusive (excluding framing bits, but including FCS octets). | detail, extensive |
| Rate packets per second | Number of packets per second received and transmitted on the VCP.                                                                                                                 | detail, extensive |
| Rate bytes per second   | Number of bytes per second received and transmitted on the VCP.                                                                                                                   | detail, extensive |

## Sample Output

### show virtual-chassis vc-port statistics

```
user@switch> show virtual-chassis vc-port statistics
fpc0:
```

```
-----
Interface          Input  Octets/Packets      Output  Octets/Packets
internal-0/24       0      / 0                0      / 0
internal-0/25       0      / 0                0      / 0
internal-1/26       0      / 0                0      / 0
internal-1/27       0      / 0                0      / 0
vcp-0               0      / 0                0      / 0
vcp-1               0      / 0                0      / 0
internal-0/26       0      / 0                0      / 0
internal-0/27       0      / 0                0      / 0
internal-1/24       0      / 0                0      / 0
internal-1/25       0      / 0                0      / 0
```

```
{master:0}
```

**show virtual-chassis vc-port statistics (EX8200 Virtual Chassis)**

```
user@external-routing-engine> show virtual-chassis vc-port statistics
```

```
member0:
```

```
-----
Interface          Input Octets/Packets      Output Octets/Packets
vcp-4/0/4           43171238 / 48152          47687133 / 51891
vcp-4/0/7           0 / 0                     0 / 0
```

```
member1:
```

```
-----
Interface          Input Octets/Packets      Output Octets/Packets
vcp-3/0/0           0 / 0                     0 / 0
vcp-3/0/1           0 / 0                     0 / 0
vcp-3/0/4           47695376 / 51899         43180556 / 48160
```

```
member8:
```

```
-----
```

```
member9:
```

```
-----
```

**show virtual-chassis vc-port statistics brief**

```
user@switch> show virtual-chassis vc-port statistics brief
```

```
fpc0:
```

```
-----
Interface          Input Octets/Packets      Output Octets/Packets
internal-0/24       0 / 0                     0 / 0
internal-0/25       0 / 0                     0 / 0
internal-1/26       0 / 0                     0 / 0
internal-1/27       0 / 0                     0 / 0
vcp-0               0 / 0                     0 / 0
vcp-1               0 / 0                     0 / 0
internal-0/26       0 / 0                     0 / 0
internal-0/27       0 / 0                     0 / 0
internal-1/24       0 / 0                     0 / 0
internal-1/25       0 / 0                     0 / 0
```

```
{master:0}
```

**show virtual-chassis vc-port statistics extensive**

```
user@switch> show virtual-chassis vc-port statistics extensive
```

```
fpc0:
```

```
-----
RX          TX
Port: internal-0/24
Total octets: 0          0
Total packets: 0          0
Unicast packets: 0          0
Broadcast packets: 0          0
Multicast packets: 0          0
MAC control frames: 0          0
CRC alignment errors: 0
Oversize packets: 0
Undersize packets: 0
Jabber packets: 0
Fragments received: 0
```

```

Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127   octets frames: 0
128-255  octets frames: 0
256-511  octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second:   0          0

...

Port: vcp-0
Total octets:          0          0
Total packets:         0          0
Unicast packets:       0          0
Broadcast packets:     0          0
Multicast packets:     0          0
MAC control frames:    0          0
CRC alignment errors:  0
Oversize packets:     0
Undersize packets:     0
Jabber packets:        0
Fragments received:    0
Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127   octets frames: 0
128-255  octets frames: 0
256-511  octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second:   0          0

Port: vcp-1
Total octets:          0          0
Total packets:         0          0
Unicast packets:       0          0
Broadcast packets:     0          0
Multicast packets:     0          0
MAC control frames:    0          0
CRC alignment errors:  0
Oversize packets:     0
Undersize packets:     0
Jabber packets:        0
Fragments received:    0
Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127   octets frames: 0
128-255  octets frames: 0
256-511  octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second:   0          0

...

```

```
{master:0}
```

#### show virtual-chassis vc-port statistics member 0

```
user@switch>show virtual-chassis vc-port statistics member 0
fpc0:
```

```
-----
Interface           Input  Octets/Packets      Output  Octets/Packets
internal-0/24        0           / 0             0           / 0
internal-0/25        0           / 0             0           / 0
internal-1/26        0           / 0             0           / 0
internal-1/27        0           / 0             0           / 0
vcp-0                0           / 0             0           / 0
vcp-1                0           / 0             0           / 0
internal-0/26        0           / 0             0           / 0
internal-0/27        0           / 0             0           / 0
internal-1/24        0           / 0             0           / 0
internal-1/25        0           / 0             0           / 0
```

```
{master:0}
```

# Troubleshooting Procedures

- [Troubleshooting Virtual Chassis Fabric on page 5391](#)

## Troubleshooting Virtual Chassis Fabric

---

This topic describes some of the following common troubleshooting issues for a Virtual Chassis Fabric (VCF):

- [Virtual Chassis Port Link Does Not Form on page 5391](#)
- [QFX5100 Leaf Device Assumes Routing Engine Role on page 5392](#)

### Virtual Chassis Port Link Does Not Form

**Problem**    **Description:** You connect a 40-Gbps QSFP+ port or a 10-Gbps SFP+ port between a leaf device and a spine device in an autoprovisioned or preprovisioned VCF. You expect the automatic Virtual Chassis port (VCP) conversion feature to convert the link into a VCP link, but the conversion doesn't occur.

The [show virtual-chassis vc-port](#) output indicates that the status of the interface is **Absent** or one or both of interfaces don't appear in the [show virtual-chassis vc-port](#) output.

**Cause**    If one end of a link is configured as a VCP and the other is not configured as a VCP, the VCP link does not form.

The automatic VCP conversion feature, therefore, does not work in the following situations:

- a 40-Gbps QSFP+ or 10-Gbps SFP+ interface on one end of the link is already configured as a VCP.

If you have previously removed a device from a VCF but haven't used the **request virtual-chassis vc-port delete** command to convert the interface that was connected to the removed device out of VCP mode, the interface is still configured as a VCP.

If you have removed a device from one Virtual Chassis or VCF and not changed the VCP port setting, the device being added to the VCF might also be configured as a VCP.

- a 40-Gbps QSFP+ port on an EX4300 switch, which is configured as a VCP by default, is interconnecting to a spine device.

**Solution** Manually configure the interface that is not configured as a VCP into a VCP using the **request virtual-chassis vc-port set** command.

### QFX5100 Leaf Device Assumes Routing Engine Role

**Problem** **Description:** A QFX5100 device configured as a leaf device assumes the Routing Engine role during VCF setup. The **show virtual-chassis** output confirms the role.

**Solution** The device can assume the Routing Engine role for several minutes during setup before it receives the configuration from the master Routing Engine, but eventually returns to the linecard role with no user intervention.

**Related Documentation**

- [Virtual Chassis Fabric Overview on page 5241](#)